UNIVERSITÀ
DEGLI STUDI
FIRENZE

Forecasting patients' admissions in an ED: The case of the Meyer Hospital

Authors:
Maria Penades Ruiz
Ines Ros Prat

Tutor:
Filippo Visintin

Engineering School
Academic Year 2014-2015

INDEX

INDEX OF FIGURES

INDEX OF TABLES

GLOSSARY

ACF - Autocorrelation Function
AIC - Akaike Information Criterion
ANSA - Agenzia Nazionale Stampa Associata
AR - Autoregressive
ARIMA - Autoregressive Integrated Moving Average
CI – Confidence Interval
ED - Emergency Department
EM - Emergency Medicine
HW - Holt-Winters
IE - Industrial Engineering
IOM - Institute Of Medicine
LWBS - Left Without Being Seen
MA - Moving Average
MAE - Mean Absolute Error
MAPE - Mean Absolute Percentage Error
MSE - Mean Square Error
PACF - Partial Autocorrelation Function
RMSE - Root Mean Square Error
SIMEU - Società Italiana di Medicina d'Emergenza-Urgenza

1. INTRODUCTION

Overcrowding in the Emergency Department (ED) has become an increasingly significant worldwide public health problem in the last decade (Di Somma et al., 2014a; Anneveld et al., 2013). Therefore, an efficient management of patient flow (demand) in EDs has become an urgent issue for many hospital administrations.

The American College of Physicians defines crowding as occurring when "the identified need for emergency service exceeds available resources for patient care in the Emergency Department, hospital, or both" (Hoot and Aronski, 2008). This phenomenon is fuelled by rapidly growing numbers of ED visitors combined with declining numbers of ED facilities (Institute of Medicine, 2006). ED crowding is associated with adverse effects on patient outcomes and can have quite detrimental consequences; it diminishes the ability to provide immediate access and stabilization to those patients who have an emergent medical condition (Aliyas, 2012a).

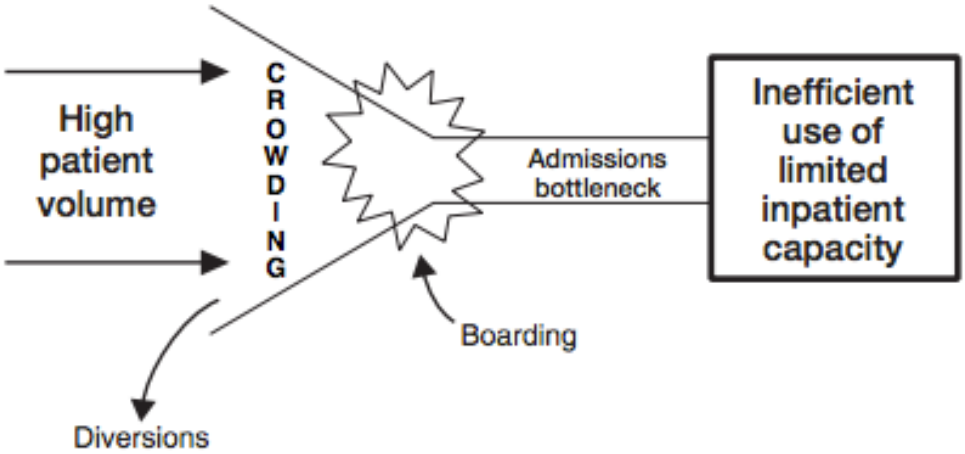

Figure 1: Visual summary of overcrowding's formation.

Overcrowding has become a major topic of discussion at Emergency Medicine (EM) conferences, such as those held annually by the Society for Academic Emergency Medicine and the American College of Emergency Physicians (Derlet and Richards, 2008).

As evidence to the magnitude of this problem, the Institute of Medicine (IOM) has estimated that over a 90% of the Emergency Departments are affected by overcrowding (Mahler et al., 2011). A survey of 250 EDs published in the Annals of Emergency Medicine in 2003 found that 11% of them regularly were on diversion, 73% had two or

more boarded patients, 59% used hallways for patients, 38% doubled up patients in rooms, and 47% used non-clinical space for patient care (Schneider et al., 2003). Difficulties quantifying crowding and providing solutions were highlighted in the recent IOM report calling for the application of advanced Industrial Engineering (IE) research techniques to evaluate ED crowding (Crane et al., 2014).

Several strategies have been tried with varying degrees of success but the problem still continues to affect hospitals across the world (Aliyas, 2012b). For example, Finamore and Turris recommended the creation of satellite clinics for reducing ED wait times (Finamore and Turris, 2009), Tanabe et al. 2008 suggested that inpatient flow could be improved by closing the waiting room and instead sending patients directly to a stretcher or a chair inside the ED, Miro et al. recommended the improvement of internal factors, such as the layout of the work environment, as a possible strategy for improving patient flow through the ED (Miro et al., 2003)...

Based on studies related to the topic, EDs in the United States have seen a near thirty million patient per year increasing volume (Di Somma et al., 2014b), receiving a total of 136.3 million visits during 2011 (National Hospital Ambulatory Medical Care Survey, 2011).

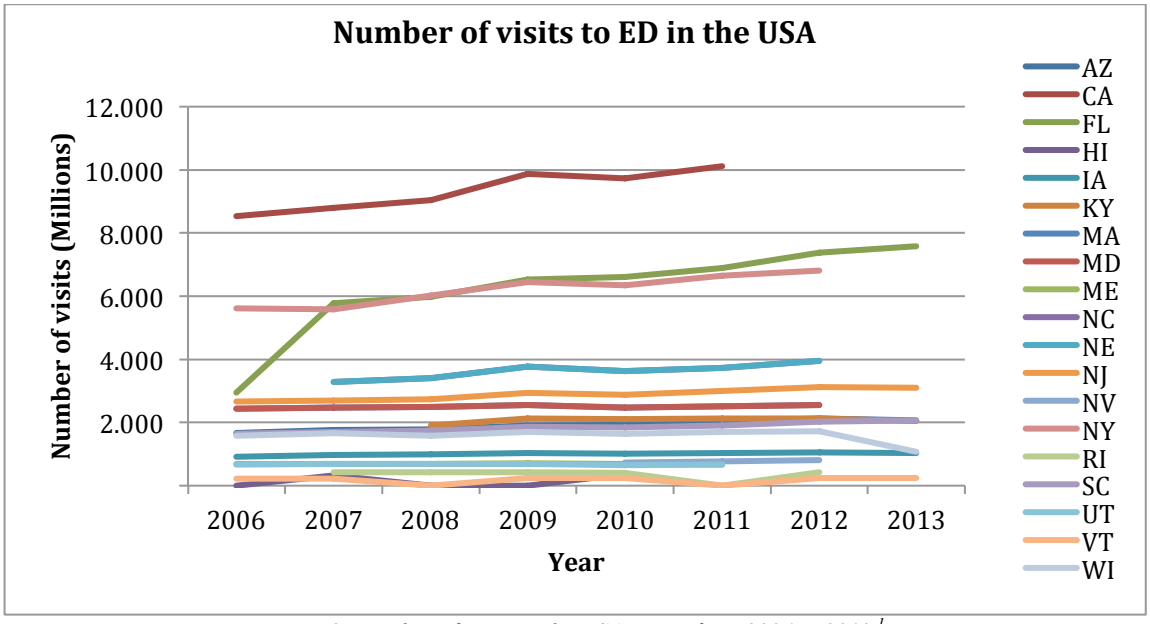

Figure 2: Number of visits to the USA's EDs from 2006 to 2013.[1]

[1] https://www.hcup-us.ahrq.gov/db/state/sedddist/sedddist_visits.jsp

Crowding is a major issue for Italian EDs as well. Unpublished data from SIMEU (Italian Society of Emergency Medicine) from July 2010 show that ED visits have grown by 5% to 6% per year over the past 5 years, with 30 million ED visits in 2009 (Pines et al., 2011). On the 23rd of January 2015, the Italian news agency ANSA reported that emergency room health workers across Italy took the streets of several cities such as Rome, Naples, Milan and Florence in protest to raise awareness of the issue.

In France much effort has been made over the past few decades to improve Emergency Department management. However, the number of visits to EDs in France has rapidly increased. Between 1996 and 1999, the annual number of visits increased by 5.8% and increased by 43% between 1990 and 1998 (Baubeau et al., 2000).

The ability to accurately forecast demand in Emergency Departments has considerable implications for hospitals to improve resource allocation and strategic planning. The aim of this study is to develop a model for forecasting monthly attendances of an ED in order to give health-care staff an opportunity to prepare for this demand and try to alleviate and mitigate problems related to overcrowding. The study has been inspired by a real context, the Meyer hospital in Florence, which is one of the most renowned children's hospitals in Europe. Data provided by the hospital itself has been analysed and the conclusions are exposed further on in this thesis.

1.1. Overall extent of the contribution

This thesis contributes to the body of knowledge in four ways:

- It shows how time-series analysis can be used to forecast demand for emergency services in a real hospital Emergency Department.
- It provides a systematic review of the literature. Articles regarding the matter, crowding, boarding and forecasting, have been classified and analysed in order to facilitate future investigations.
- A structural analysis of the Meyer hospital to describe the arrival process and the understanding of the ED procedure to optimize its performance.
-

1.2. Structure of the thesis

In section 2 the literature that has provided information for the development of the model has been analysed. The review protocol and classification criteria have been illustrated in detail. A flow chart providing evidence of the final election based in a meaningful way is included.

Section 3 gives details of the methodology. A justification for the chosen model is provided as well as the theory behind it.

Section 4 presents information about the selected hospital as well as its ED. The selected method is applied to the provided data and results are presented.

Section 5 concludes synthesizing the results in an understandable way as well as the extent of the contribution.

2. LITERATURE REVIEW

One of the most important parts of a scientific thesis is the analysis of the literature. Knowing where to look up for the right information, organising and examining it are key steps to perform a strong-based thesis. Therefore, a considerably amount of time has been spent in this part of the project.

The quantitative tools to support the analysis of the data collected from the hospitals constitute the main literary reference to this thesis.

Section 2.1 presents the classification criteria that have been followed to carry out the research in the three different topics: crowding, boarding and forecasting.

Subsequently, section 2.2, proposes a description and a structured classification of quantitative models used to forecast data from the hospitals.

Finally, the last section places our work on the basis of the proposed classification and explains the different techniques that have been carried out for the development of the thesis.

2.1. Classification criteria

Due to the importance of forecasting the number of patient arrivals in the hospital to maintain performance and to help enhance the management of hospitals establishments, several forecasting techniques have been developed. Consequently, a deep search has been carried out in various search engines, specifically in Science Direct, Emerald, IngentaConnect, Taylor and Francis, SAGE, Springer, Scopus, EBSCO and Informs. Different combinations of key words have been used to download different articles:

“Emergency Department” (Title/abstract/keywords) + Crowding (everywhere)
“Emergency Department” (Title/abstract/keywords) + Boarding (everywhere)
"Emergency Department" (Title/abstract/keywords) + Forecasting (everywhere)

In order to focus the research, a 5-level classification approach has been applied to the relevant articles aiming to grade the different aspects of its techniques.

Here we distinguished between qualitative and quantitative. Qualitative-based forecasting methods predict the future, usually using opinion and management judgment of experts in specific fields. Quantitative methods, on the other hand, rely on mathematical models.

The next performed categorisation consisted in differentiating among articles by their main theme: forecasting, boarding or crowding. Depending on how much the articles treated the subject, they were graded with a 0, a 1 or a 2.

Last but not least, for those subjects marked as quantitative, the methods they used were specified; simulation, optimization, queuing theory, Markov chains, and system dynamics. The option "Others" was used for those articles whose methods did not consist in one of the previously stated and the option "Not sure" for those whose methods were not clear.

To provide evidence of how the final sample have been selected, the next flow-chart has been created.

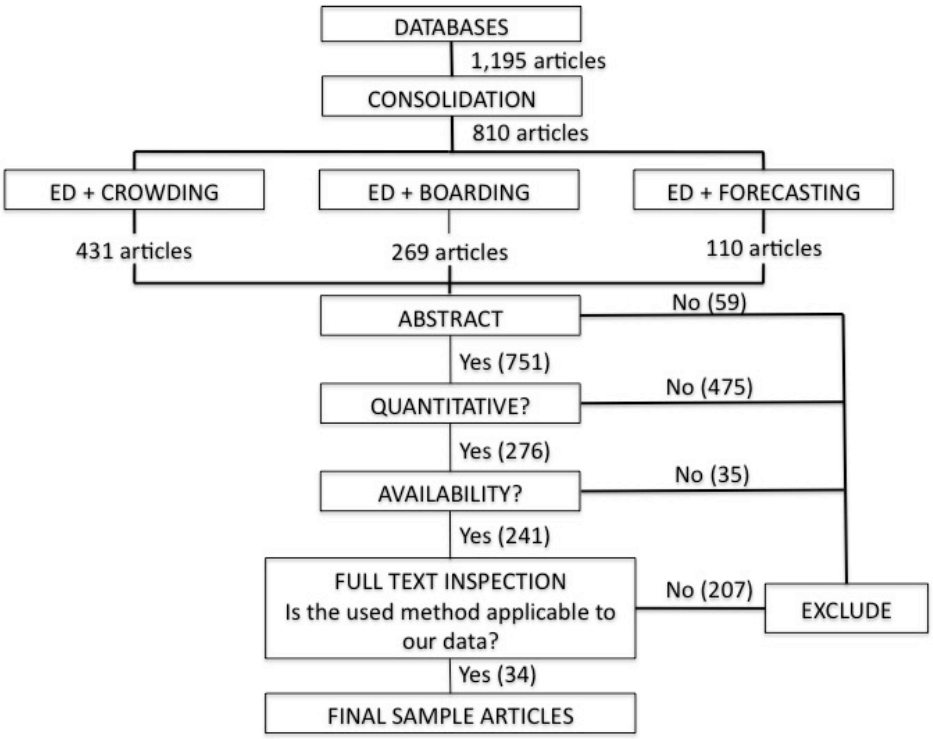

Figure 3: Flow chart illustrating the selection process of the final sample articles.

2.2. Literature analysis

Once this categorization has been done, it has been decided to focus the studio in those quantitative articles related to forecasting issues using a proposed classification by Wargon et al. It includes a table that contains articles evaluating patient-volume forecasting in walk-in centres or ED from 1981 to 2007 (Wargon et al., 2009). Table 1 completes the study integrating papers until 2015.

Paper Reference	Year	Objective	Type of setting	Location	Type of Patients	Data used	Prediction time horizon	Methods	Model parameters	Findings
Almenhdawe_2013	2013	Analyse the interaction between an EMS provider and multiple EDs in a region	Walk-in clinic, single centre	Ontario, Canada	Intermediate and acute care patients	. Data from the flow time of patients of a hospital.	-	Markow chain	**k**: number of hospitals. **N**: number of ambulances. **lambda subzero**: ambulance patient arrival rate to the EMS system. **Lambda subk**: arrival rate of walk-in patients at the kth ED. $\mathbf{p_k}$: probability that an ambulance patient is send to the k^{th}ED. $\mathbf{c_k}$: number of servers in the kth ED, which corresponds to the service capacity at the kth ED. **mu subk**: service rate per server in the kth ED.	Ambulances usually operate at around uA = 35% utilization (including transit time). For such a case, the loss probability is small. This indicates that the queuing network introduced in this paper is robust as long as the system of interest is working under normal operating conditions.
Araz_2014	2014	Develop forecasting models which may be used to assist hospitals and other health agencies in improving their resource use and quality of care while responding to disease outbreaks.	ED visits, single centre	Omaha, Nebraska	ILI-related patients	ED related ILI-visits	Weekly forecast	SARIMA model, Holt Winters exponential smoothing models and 3 linear regression models	ARIMA (1,1,1) (0,1,0) The exponential smoothing model is run with parameters of (0.45, 0, 0.86). Linear regression models: total influenza antigen testing data (x1), positive Ag test data (x2), total Respiratory Syncytial Virus infection (RSV) test data (x3), positive RSV test data (x4), Omaha GFT data (x5), Douglas County ED visits data (x6), and ILI Network data (x7), to forecast the NMC hospital ED visits (y). Because there is some time lag between the actual ED visits and GFT data (x5), Douglas County ED visits data (x6), and ILI Network data (x7).	A regression model constructed with the full data set used as the predictors of ILI-related ED visits performed better than other models (RMSE = 57.61)
Boyle_2008	2008	Forecast the time that admitted patients leave the ED	ED visits, two centres	Queensland, Australia	ED patients	Five years of ED presentation and admission data (1/7/02 – 30/6/07)	Forecasts for the six months Jan'07-Jun'07	Different regression models	y = response X=design matrix Beta= parameters E= random disturbances	The highest accuracy was linear regression with monthly variation modelled with 11 dummy variables.
Boyle_2011	2011	Describe the use of surveillance and forecasting models to predict and track epidemics (and, potentially, pandemics) of influenza.	ED visits, 27 centres	Queensland, Australia	Patients with influenza symptoms	5 years of historical data (2005–2009) on ED presentations and hospital admissions for influenza-like illnesses in 27 Queensland public hospitals.	**1.** One day ahead forecast **2**. Four week window	**1**. Surveillance monitoring, CUSUM plan. **2**. Historical data forecast. **3**. The correlation coefficients between internet search data for Queensland and statewide ED influenza presentations indicated an increase in correlation since 2006 when weekly influenza search data became available.	-	-
Burr_2006	2006	Evaluate the impact of inconsistent seasonal effects on performance assessments in the context of detecting anomalous counts in data that exhibit seasonal variation.	ED visits, single centre	Albuquerque NM, Mexico	Patients for which the chief complaint was mapped into the respiratory category	Respiratory syndrome daily counts in ED	-	Non-hierarchical and hierarchical models	-	For the data we analyze, the "one season fits all " assumption is violated, and Detection Probabilities performance claims based on simulated data that assume "one season fits all," for the forecast methods considered, except for moving average methods, tend to be optimistic. Moving average methods based on relatively short amounts of training data are competitive on all three data sets, but are particularly competitive on the real data and on data from the hierarchical model, which are the two data sets that violate the "one season fits all" assumption.
Chase_2012	2012	Determine which indicators could be used to accurately model the state of the system and determine how far in advance a significant increase in patient volume could be predicted to adequately plan and prepare to prevent a crowding situation.	ED visits, single centre	-	Adults	Daily visits from 12:00 AM July 1,2009, through 11:45 PM November 30, 2010.	-	Care Utilization Ratio (CRU), graphical analysis, binary logistic regression analysis	Care Utilization Ratio (CRU)	The CUR was a robust predictor of the state of the ED. Prediction intervals of 30 minutes, 8 hours, and 12 hours performed best of all models analyzed, with deviances of 1.000, 0.951, and 0.864, respectively. A 95% significance was used to validate the models against the July 2010 through November 2010 data set. Positive predictive values ranged from 0.738 to 0.872, true positives ranged from 74% to 94%, and true negatives ranged from 70% to 90% depending on the threshold used to determine the state of the ED withthe 30-minute prediction model.

Paper Reference	Year	Objective	Type of setting	Location	Type of Patients	Data used	Prediction time horizon	Methods	Model parameters	Findings
Chen_2011	2011	Analyse the effects of meteorological, clinical and economic factors on monthly ED revenue and visitor volume.	ED visits, single centre	Kaohsiung, Taiwan	Pedriatic, trauma and non trauma ED visits	Monthly data	Four-year (2005–2008) data set was used to construct the forecasting model, while the data for the first 9 months of the 5th year (2009) was used to test the forecasting capability of the model	Spearman correlation and cross-correlation analyses, ARIMA model	ARIMA model (1, 0, 0) MAPE Consumer Price Index (CPI)	Meteorological, clinical, and economic factors are associated with ED revenue and visitor volume. The good long-term forecasting capability of the model proposed in this study can help EDs to optimize departmental resources and manpower.
Eng_2007	2007	Describe the time demand patterns at the ED and apply systems status management to tailor ED manpower demand.	ED visits	Singapore	Patients of all ages	Demographic information, time of registration, waiting time and processing time	-	Observational study of all patients presenting to the ED at the Singapore General Hospital during a 3-year period and a time series analysis to determine time norms regarding physician activity for various severities of patients.	-	The yearly ED attendances increased from 113 387 (2004) to 120 764 (2005) and to 125 773 (2006). There was a progressive increase in severity of cases, with priority 1 (most severe) increasing from 6.7% (2004) to 9.1% (2006) and priority 2 from 33.7% (2004) to 35.1% (2006). Existing demand pattern, with seasonal peaks in June, weekly peaks on Mondays, and daily peaks at 11 to 12 AM.
Hoot_2011	2011	Forecast ED crowding at multiple institutions, and assess its generalizability for predicting the near-future waiting count, occupancy level, and boarding count.	ED visits, multiple centres	United States	All patients at each participating site during the study period (11/1/2005 – 1/31/2007)	Daily visits	-	The Forecast ED tool implements a computerized "virtual ED" through a discrete event simulation intended to mimic the operations of an actual ED.	MAE= 0.6-3.1%,occupancy level: MAE=9-14.5% and boarding count: MAE= 0.9-2.7%	The Forecast ED tool generated potentially useful forecasts of input and throughput measures of ED crowding at five external sites, without modifying the underlying assumptions
Jones_2008	2008	Study the temporal relationships between the demands for key resources in the emergency department (ED) and the inpatient hospital, and develop multivariate forecasting models.	ED visits, multicentre	Utah and southern Idaho.	ED patients and inpatients	ED arrivals, ED census, ED laboratory orders, ED radiography orders, ED computed tomography (CT) orders, Inpatient census, Inpatient laboratory orders, Inpatient radiography census, Inpatient CT orders.	Forecasts made from 1 to 24 hours ahead	Descriptive analysis and model fitting were carried out using graphical and multivariate time series methods. Multivariate models were compared to a univariate benchmark model in terms of their ability to provide out of-sample forecasts of ED census and the demands for diagnostic resources.	MAE (Figure), R2 (Table)	Descriptive analyses revealed little temporal interaction between the demand for inpatient resources and the demand for ED resources at the facilities considered. Multivariate more accurate forecasts of ED census and of the demands for diagnostic resources.
Kam_2010	2010	Develop and evaluate time series models to predict the daily number of patients visiting the Emergency Department (ED)	ED visits, single centre	Korea	Adults	Daily visits	3 months	Three forecasting models were established: 1) average; 2) univariate seasonal auto-regressive integrated moving average (SARIMA); and 3) multivariate SARIMA.	Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC) and Mean Absolute Percentage Error (MAPE)	The multivariate SARIMA model was the most appropriate for forecasting the daily number of patients visiting the ED. Because it's MAPE was 7.4%, this was the smallest among the models, and for this reason was selected as the final model.
Kecojevic_2011	2011	Increase our understandingof perceived benefits and disadvantages of HIV testing in EDs and to codify domains of public health and clinical care most affected by implementing HIV testing in EDs.	-	Baltimore, MD	-	Data were collected from the inaugural conference of the National Emergency Department HIV Testing Consortium.	-	Opinions were systematically collected from attendees of the 2007 National ED HIV TestingConsortium meeting. Structured evaluation of strengths, weaknesses, opportunities, and threats analysis was conducted to assess the impact of ED-based HIV testing on public health. A modified Delphi method was used to assess the impact of ED-based HIV testing on clinical care from both individual patient and individual provider perspectives.	-	Experts in ED-based HIV testing perceived expanded ED HIV testing to have beneficial impacts for both the public health and individual clinical care; however, limited resources were frequently cited as a possible impediment.
Kline_2010	2010	Attribute matching matches an explicit clinical profile of a patient to a reference database to estimate the numeric value for the pre-test probability of an acute disease	ED visits, 15 centres	3 countries	Adults	Time of clinical evaluation for suspected pulmonary embolism (PE).	6 year study	Wells Logistic Regression–based Model, PERC Rule, Wilcoxon method	-	Attribute matching categorizes over one-quarter of patients tested for PE as having a pre-test probability of <2.5%, and the observed rate of VTE within 45 days in this subset was <2.5%
Laker_2014	2014	Evaluate flexible partitioning between low- and high-acuity ED areas to identify the best operational strategy for subsequent implementation	ED visits, single centre	-	Adults	Daily visits (85000 visits/year)	1 year study	Discrete-event simulation (DES)	-	Adding some flexibility into bed allocation between low- and high-acuity can provide substantial reductions in overall patient waiting and a more efficient ED.

Paper Reference	Year	Objective	Type of setting	Location	Type of Patients	Data used	Prediction time horizon	Methods	Model parameters	Findings
Laskowski_2009	2009	Patient flow through ED	ED visits, multiple centres	Canada	-	-	-	Agent based modelling simulation (ABM), Queuing Model (QM)	-	The two modelling methodologies applied to investigating patient access and patient waiting times in hospital EDs, to develop tools that can help guide policy and practice improvements.
Marcilio_2013	2013	Develop different models to forecast the daily number of patients seeking ED care in a general hospital according to calendar variables and ambient temperature readings and to compare the models in terms of forecasting accuracy.	ED visits, single centre	Brazil	Adults	Daily visits (389 visits/day)	33 months to develop the ED patient visits forecasting models and last 3 months to measure each model's forecasting accuracy by the mean absolute percentage error (MAPE).	Three different time-series analysis methods: generalized linear models (GLM), generalized estimating equations (GEE), and seasonal autoregressive integrated moving average (SARIMA).	MAPE of each model (Table)	In this setting, GLM and GEE models showed better accuracy than SARIMA models. Including information about ambient temperature in the models did not improve forecasting accuracy. Forecasting models based on calendar variables alone did in general detect patterns of daily variability in ED volume and thus could be used for developing an automated system for better planning of personnel resources.
McNaugbton_2012	2012	Evaluate the relationship between ED bed assignment (traditional, hallway, or conference room bed) and mean ED evaluation time, defined as the time spent in an ED bed before admission or discharge	ED visits, single centre	-	Adults	Daily visits Monday-Friday 11 AM to 11 PM (19 239 visits/year)	1 year study	Multiple linear regression and marginal prediction	-	Patients assigned to non-traditional beds experience a small delay in ED disposition compared with non-traditional beds.
Peavey_2012	2012	Introduce simulation and mock-up research methods used to inform and optimize building design	-	-	-	-	-	Simulation and mock-up models	-	Modelling simulations using data to provide valid, compelling feedback to be implemented in design decision making
Rapold_2012	2012	Evaluategeneral and specific measures of community health as modifiers of risk for asthma and congestive heart failurefollowing an episode of acute exposure to wildfire smoke.	ED visits, multicentre	North Carolina, US	Adult patients with asthma and CHF patients over 44 years old.	Daily visits during 44 days	-	HYSPLIT model and Monte Carlo approximation	-	The results indicate that Socio-Economic Factors should be considered as modifying risk factors in airpollution studies and be evaluated in the assessment of air pollution impacts.
Reis _2004	2004	To study the effects of different syndrome grouping methods on model accuracy, a key factor in the outbreak-detection performance of syndromic surveillance systems	ED visits, two centres	-	Adults and Children	Daily visits during 1,680 consecutive days with 230,000 total patient visits classified as 1 relying on chief complaint, 1 on diagnostic codes, and 1 on a combination of the two	1680 days of study and forecasting models	3 models: Moving Average, Linear, Exponential	MAPE1 and MAPE2 (table)	The methods used to group input data into syndromic categories can have substantial effects on the overall performance of syndromic surveillance systems and can improve the modelling accuracy and its detection sensitivity.
Reis_2003	2003	Present a methodology for developing models of expected ED visit rates	ED visits, single centre	Boston, US	Children	Daily ED visits (137 visits/day)	Models constructed with 8 years of data and validated with the next 2 years data.	Time series methods, trimmed-mean seasonal models and ARIMA model.	ARIMA(2,0,1) model for ED volume with MAPE=9.37% and ARIMA(1,0,1) model for respiratory-related ED volume with MAPE=27.54%	Time series methods applied to historical ED utilization data are an important tool for syndromic surveillance and can be generalized to other healthcare settings to develop automated surveillance systems capable of detecting anomalies in disease patterns and healthcare utilization
Schweigler_2009	2009	Investigate whether models using time series methods can generate accurate short-term forecasts of emergency department (ED) bed occupancy	ED visits, 3 centres	-	Adults	Hourly ED bed occupancy values of three hospitals (98,199 patients/year in Site 1, 59,344 in Site 2 and 55,757 in Site 3)	-	**1.** Hourly historical average **2.** Seasonal autoregressive integrated moving average (ARIMA) **3.** Sinusoidal with an autoregression (AR)-structured error term. For each site, the accuracies of the three methods were compared with one-way analysis of variance (ANOVA), followed by post hoc comparisons with Tukey Kramer statistics.	Parameters of the models (Tables)	Both a sinusoidal model with AR-structured error term and a seasonal ARIMA model [(1,0,1)(0,1,1)] were found to robustly forecast ED bed occupancy 4 and 12 hours in advance at three different EDs, without needing data input beyond bed occupancy in the preceding hours. This forecasting method was found to work equally well at three different institutions with differing operational characteristics, without having to adjust any of the model input variables
Sun_2009	2009	Forecast emergency department (ED) attendances	ED visits, single centre	Singapore	Adults	Daily visits during 1005 days, classified as P1, P2 and P3 by the patient acuity category scale (400 visits/day)	30 months of study and prediction	ARIMA, Univariate analysis by t-tests and multivariate time series analysis	MAPE P1=16.8%,, MAPE P2=6.7%, MAPE P3=8.6% and MAPE TOTAL=4.8%	P1: ARIMA(0,1,1), P2: ARIMA(1,1,1)(1,0,1), P3: ARIMA(0,1,1)(1,0,1).Daily patient attendances at ED can be predicted with good accuracy using the modelling techniques in time series analysis.

Paper Reference	Year	Objective	Type of setting	Location	Type of Patients	Data used	Prediction time horizon	Methods	Model parameters	Findings
Tehrani_2013	2013	Estimating the annual national costs associated with ED visits for dizziness	ED visits	US	Dizziness and vertigo patients	Visits during 15-year to ED for dizziness (12,202 visits) and for no dizziness (360,424 visits) to predict 2011 year	-	Time-series forecasting analysis using a stepwise autoregressive method that combines time-trend regression with an autoregressive model and uses a stepwise method to select the lags to use for the autoregressive process	The estimated number of 2011 US ED visits for dizziness or vertigo was 3.9 million (95% confidence interval [CI] = 3.6 to 4.2 million) 95% confidence interval [CI] = 3.6 to 4.2 million	Evaluation and management of dizziness and vertigo in the ED is very costly ($4 billion/year, about 4% of total ED costs)
Vermeulen_2009	2009	Testing whether the balance between daily hospital admissions and discharges affects next-day ED length of stay	ED visits, multicentre regional	Toronto	Children and adults patients	Visits per 3 years (22,995 visits)	3 years study	Measuring daily ratio of admissions to discharges at each hospital and the next-day median ED length of stay in the same hospital by using linear regression.	**SD (**50th percentile ED length of stay) = 218 (51) minutes.	Disequilibrium between the number of admitted and discharged inpatients significantly affects next-day ED length of stay
Wargon_2010	2010	If creating mathematical models using calendar variables could identify the determinants of ED census over time and assessed the performance of long term forecasts	ED visits, multicentre regional	Paris	Patients greater than 15 years of age	Visits per two years (299743 patients in 2004-2005 and 322510 patients in 2006-2007)	4 years study	Univariate General Linear Model	**Mean, SD, max, min and sum** of the four hospitals (table), R2=0.5035, MAPE=4.45%	The forecast model explains 50% of the variability
Wathen_2007	2007	Child randomized trial compared nerve block tomorphine in the management of pain caused by femur fracture	ED visits, single centre	Ontario	Children aged 15 months to 18 years	Patients aged 15 months to 18 years and presenting to pediatric ED with an acute femur fracture	40 months study	Kaplan-Meier survival analysis method and 1-sided t-test	**Median duration of analgesi**a= 313 minutes [95% CI 154 to 360 minutes]; **Median duration of morphine**=60 minutes [95% CI 10 to 255 minutes]	The use of nerve blocks to provide analgesia offers several advantages compared with traditional systemic medications
Wu_2013	2013	Enhance patient flow and throughput as well as to preserve limited resources for the sickest patients	ED visits, single centre	Taiwan	Adults	ED patients between 7:30 to 11:30 AM from Wednesday toSunday (3305 visits in 8 months)	8 months study	X2 test, Mann Whitney U test, Student t test, logistic regression: relationship between discharge rates before and after intervention of total patients and different triage groups	**Mean+-SD** table of demographic factors, occupancy rate, 72h revisit rate before and after intervention	The intervention strategy used peer pressure enhanced by e-mailreminders to improve patient flow and throughput. More patientswere discharged at the end of shifts, particularly triage level IIIpatients. The ED LOS decreased significantly by a median of 17.9minutes for patients whose final disposition was discharge.

3. METHODOLOGY

3.1. Forecasting selected method

After all the research carried out and the analysis of the most relevant articles and books, for several reasons, it has been decided that two interesting methods to apply are the ARIMA model and exponential smoothing. These two widely used techniques provide complementary approaches to forecasting time series that will be discussed further on in this thesis.

Linear regression was discarded due to the lack of a predictor variable to base the predictions in.

3.2. Theory behind the methods

Exponential smoothing and ARIMA models are the two of the most commonly used methods to time series forecasting. While ARIMA models describe the autocorrelations in the data, the exponential smoothing aims to describe the trend and seasonality in the data.

3.2.1. Holt-Winters seasonal method

Forecast based on exponential smoothing methods assign the past observations with exponentially decreasing weights. In other words, recent observations are given relatively more weight in forecasting than older observations. Exponential smoothing is used to make short-term forecasts and it makes no assumptions about the correlations between successive values of the time series. It should also be mentioned that for forecasts using this method, the prediction intervals require a non-correlation between forecast errors.

The Holt-Winters seasonal method is considered an extension of the exponential smoothing, as it is able to deal with time series that contain both trend and seasonal variations. It has two versions, additive and multiplicative methods, the use of which depends on the characteristics of the particular time series. The former refers to models whose random fluctuations in the data are roughly constant in size over time and the

latter to series that present seasonal variations that change proportional to the level of the series itself.

Smoothing is determined by three forecast equations (level, trend and seasonal component) and three smoothing equations controlled by three smoothing parameters: alpha, beta, and gamma, all of which have values between 0 and 1. Values close to 0 mean that relatively little weight is placed on the most recent observations when making forecasts of future values. The parameter alpha (α) is the smoothing parameter of the level, beta (β) stands for the trend (slope) and gamma (γ) for the seasonal component.

3.2.3. Seasonal ARIMA model

A very important part of the process of fitting an ARIMA model is stationarizing. Therefore, before introducing the model itself, it is necessary to discuss the technique of differencing time series.

A stationary time series is one whose properties such as means, variances and correlations do not depend on the time at which the series is observed. For example, a time series that presents trends or seasonality is not considered stationary; its values will be affected by this trend or seasonality.

There are different ways to transform a non-stationary time series in a stationary one. A very useful way of doing so is differencing. A differenced series is the change between consecutive observations in the original series. This procedure can help stabilize the mean of a time series by removing changes in the level of a time series, and so eliminating trend and seasonality. Occasionally, to obtain a stationary series, more than one difference needs to be done.

To identify a non-stationary time series the ACF and PACF, as well as the time plot of the data, are very helpful.

Apart from ordinary difference or first differences, a series can also present a seasonal difference. It is understood as the difference between an observation and the corresponding observation from the previous year.

Once the data subject of study is considered to be stationary, it is time to develop the ARIMA model. ARIMA stands for Auto Regressive Integrated Moving Average model. From its name it is easy to deduce that it combines differencing with autoregression and

a moving average model. In an autoregression model, the output variable depends linearly on its own previous values; it is a regression of the variable against itself. A moving average model is defined as a linear regression of the current value of the series against past forecast errors.

In this thesis' case study, a seasonal ARIMA model will be developed, which is formed by including additional seasonal terms to the non-seasonal model. It is written as follows:

$$ARIMA\ \underbrace{(p,d,q)}_{\substack{Non-seasonal\\ part}}\ \underbrace{(P,D,Q)_m}_{\substack{Seasonal\\ part}} \qquad m = \text{number of periods per season}$$

To distinguish between seasonal and no seasonal parts, uppercase notation and lowercase notations are used respectively. The different variables of the seasonal part stand for:

- p = order of the autoregressive part
- d = degree of first differencing involved
- q = order of the moving average part

The seasonal part of the model consists of terms that are very similar to the non-seasonal components of the model, but they involve backshifts of the seasonal period.

Selecting appropriate values for p, d and q can be difficult. Telling its value from a time plot is not usually possible so the ACF and PACF are again very useful tools. The ACF shows the correlations between values of the process at different times. The PACF measures the relationship between y_t and y_{t-k} after removing the effects of other time lags.

To confirm that the chosen model is the adequate one for the data, it is necessary to check the residuals by plotting its ACF and confirming that it behaves as white noise; they are not autocorrelated and have a zero mean.

Once the residuals have been analysed it is time to start forecasting.

4. CASE STUDY

4.1. The Meyer Hospital

The project is focused in the Meyer Children Hospital, a pediatric hospital located in Florence, Italy. The hospital is an official member of the European Network of Health Promoting Hospitals of the World Health Organization and the personnel are involved in prevention and health promotion programs for the Regional and National Health Departments.

It was founded in 1884 by the Marquis Giovanni Meyer in memory of his wife Anna. The Meyer Pediatric Hospital was one of the firsts hospital institutions in Italy exclusively devoted to the problems of child health care from birth to adolescence.
Its fame reached a peak just after the Second World War when tubercular meningitis was treated successfully for the first time.
In 1995, the Meyer Hospital, with the Department of Pediatrics of the University of Florence, became an independent health institute of the National Health System, due to recognition of its role as a highly specialized pediatric institution.

The original structure faces onto Via Luca Giordano and it has subsequently been enlarged to house numerous additional services, which have opened over the years.
In 2007 the historic seat was dismissed and established in a new structure near Florence central hospital, Careggi. With this transfer Meyer registered a significant quality jump, confirming its position among the most firm and innovative realities of Italian pedriatics. With the new Meyer came new objectives; 210 beds, 7 operation rooms, 9 diagnostics rooms, 5.000m^2 of gardens and terraces and on the roof a total area of 32.000m^2. Everything is immersed in a park of 72.000m^2. It is a hospital that for its biocompatible solutions, its innovative ideas in the blueprints, the use of simple and non-toxic materials and the vast use of the colour and light, represents an absolute innovation panorama of the Italian sanitary building.

Nowadays, the company Anna Meyer Hospital provides expertise and dedicated health treatment and services to infants, child, adolescents and new born. Sanitary excellences, modern technologies and the elevated quality of the acceptance are designed in a way to have the little patient and his family in a centre of everything. There is not a space or a

tiny particular detail that has not been thought in favour of a child and his family, the real protagonists of the Children's Hospital. For this reason, the hospital has tried to create a lively atmosphere; clowns and entertainers in the wards, recreation corners for games or study, a video library and other services: these are all elements that ensure that the children admitted have a hospital life as serene as possible. Those forced to spend long periods of time in the hospital can keep up with their education thanks to the support of schoolteachers.

It is a public hospital funded by private donations: companies or individuals. The body responsible for the management and administration of these grants is the Meyer Foundation. It was established as an operational tool synergistic to the Hospital Meyer for activities and assistance to the fundraising Meyer. Keeping with the policy choices of the Hospital, the Foundation is also a way to help and support the realization of actions, which give an "added value" to Meyer, as it becomes increasingly skilled in technical and scientific activities that are appreciated to the public.

The aims of the Foundation are in fact both: the development of initiatives in the territory of Meyer and the research on the issue of child specialization with particular reference to the psycho-pedagogical problems. Particular attention is paid by the Foundation in support of Meyer's commitment to improve the quality of living of children and families in the hospital and also by supporting the training of personnel and scientific research carried out by Meyer.

4.2. The Meyer Emergency Department

The pedriatic Emergency Department of the Hospital Meyer is committed to ensuring the best care for children in the most fair and optimistic way. It is open 24 hours and it works closely with the Territorial Emergency Service. 20 to 24 doctors, 35 nurses and 6 auxiliary nurses integrate it.

The most urgent cases are treated first, following the strict rules of the internationally recognized "triage" system, used in crowded emergency rooms and walk-in clinics. It is used to prioritize the use of space and equipment such as operating rooms in crowded medical facilities. This advanced triage system involves a colour-coding scheme using

red, yellow, green, white, and black tags assigned to patients depending on their critical condition which defines the priority to access to medical examination. Trained nurses assess patients and assign them a colour.

- Red code: immediate access to medical examination.
- Yellow code: access within 15 minutes.
- Green code: access within 60 minutes to medical examination.
- White code: access within 180 minutes and in any case after the red, yellow and green codes.

It is guaranteed that the patient will not spend more that 4 hours from their triage to their return home, unless of course clinical situations which result in hospitalization or observation.

For the less urgent cases (green and white codes) waiting times for medical examination can exceed the established limits in case of large crowds or to the presence of particularly severe cases. In this case, the medical staff will inform the patient. It is recognised that approximately a 5% of the annual patients leave the ED without been seen (LWBS). The vast majority of these cases that return home are patients who have been assigned a white triage code. According to data provided by the hospital, a 4.89% of the patients LWBS in 2013 out of a total of 42,722 visits and in 2014, of 44,800 visits a 5.10% LWBS.

4.2.1. ED arrival process and procedure

There are two types of patients that enter the ED; walk-in patients and those brought in by ambulances that access directly to medical examination.

The first step for walk-in patients is the assignment of a triage colour; patients may have to wait in this early step if the ED is too crowded. Once they have been assigned a colour, they are moved to a waiting room that corresponds to their triage code. The procedure continues with the patients' access to medical examination, which takes place in equipped boxes. After the doctor's diagnosis, the patient can either go home or stay in for observation. There are two observation rooms; one with chairs for those patients who have to wait for several hours and the other one with beds, for patients who may

have to stay in up to 2 days before going home or been transferred to a particular department where they will receive treatment.

4.2.2. ED admissions analysis

A structural analysis of the Meyer Hospital to describe the arrival process has been done with R, a widely used statistical software. Data provided by the hospital itself, including daily visits to the ED from January 2009 to May 2015, has been used to better understand its performance.

The information related to each patient's visit is: Admission Year, Admission Triage Score, Admission Triage Code, Discharge Triage Score, Discharge Triage code, ID of the Patient, Admission time, Time when a Patient is seen by a doctor for the first time, Short Intensive Observation, Diagnosis Code, Discharge Time, Discharge Code and Discharge Code Description.

First of all, to obtain a general vision of the number of admission to the Meyer ED it is interesting to take a look at the next graph.

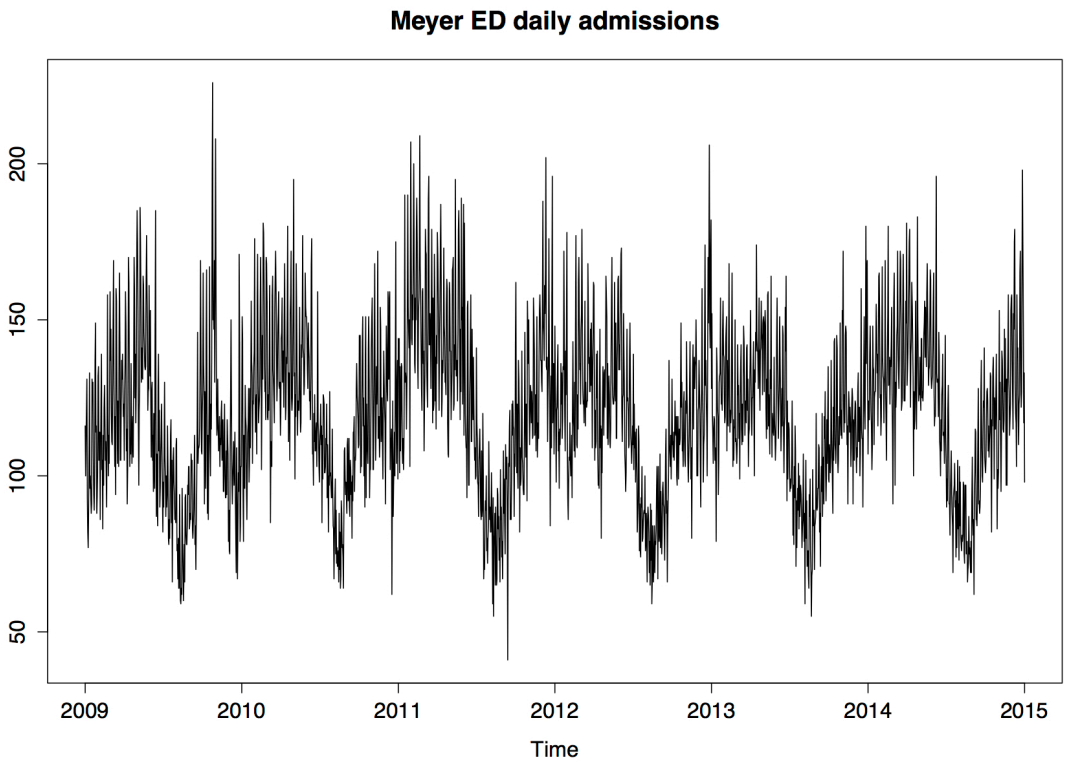

Figure 4: Daily admissions to the Meyer's ED from January 2009 to December 2014.

Figure 4 helps to get an idea of the number of daily patients during the six years of study. It is easy to observe that there is a high demand during all this period, fact that supports the relevancy of the topic. Its behaviour and tendencies will be detailed as following.

Monthly admission's plot can be used to detect if there is any kind of trend or seasonality.

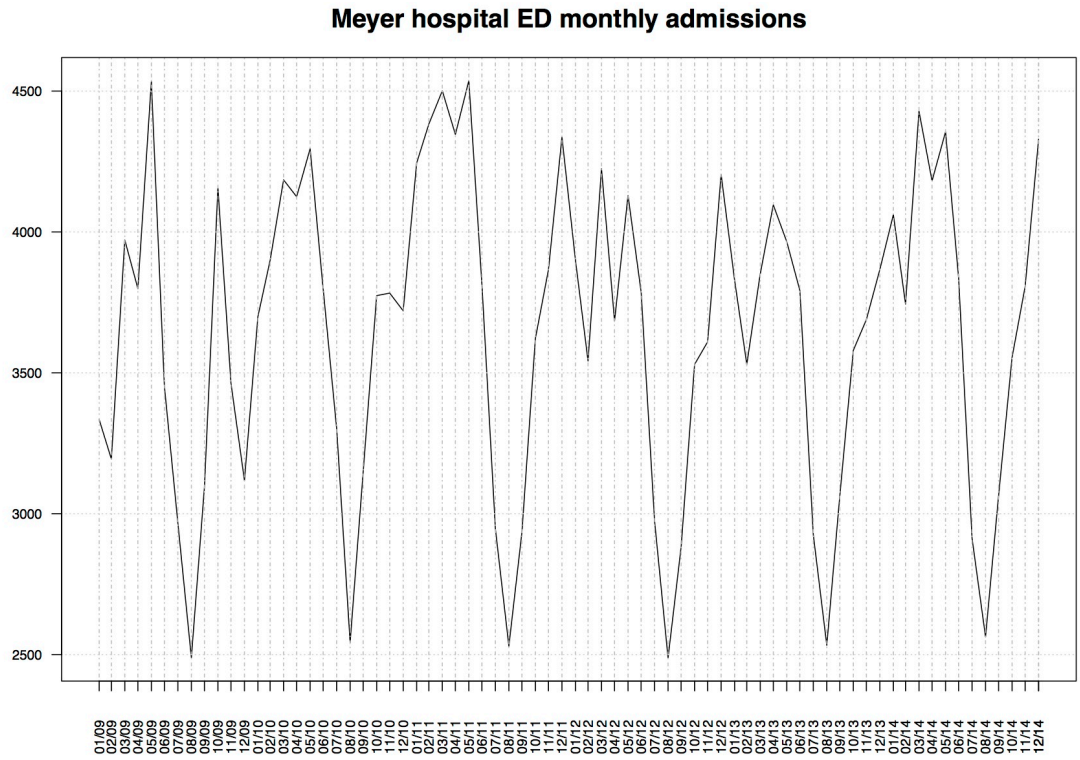

Figure 5: Monthly admissions to the Meyer's ED from January 2009 to December 2014.

It can be perceived that each year, after the ED receives its maximum number of visits, a very accentuated fall takes place, presenting its minimum during the month of August. This behaviour can be explained by the fact that it is the holiday season and families are out of town. The fact that during the summer season is when fewer cases of flu take place can also support this fall.

Figure 6 illustrates the yearly admissions to the ED. It can be observed that the number of patients visiting the ED of the Meyer's Hospital differs a lot from one year to another. The graph presents a peak in the year 2011, when more than 460.000 patients visited the ED.

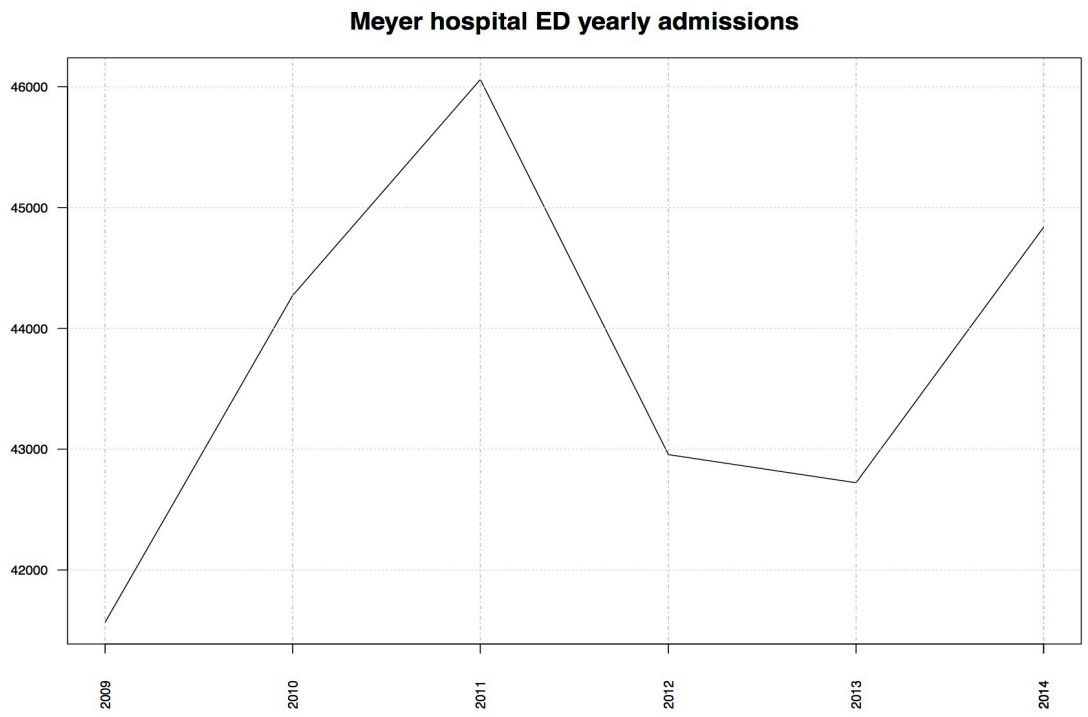

Figure 6: Yearly admissions to the Meyer's ED from 2009 to 2014.

As it has previously been explained, the use of the triage system results very interesting and so does its behaviour. The following graph illustrates the arrivals to the ED stratified by the colour code.

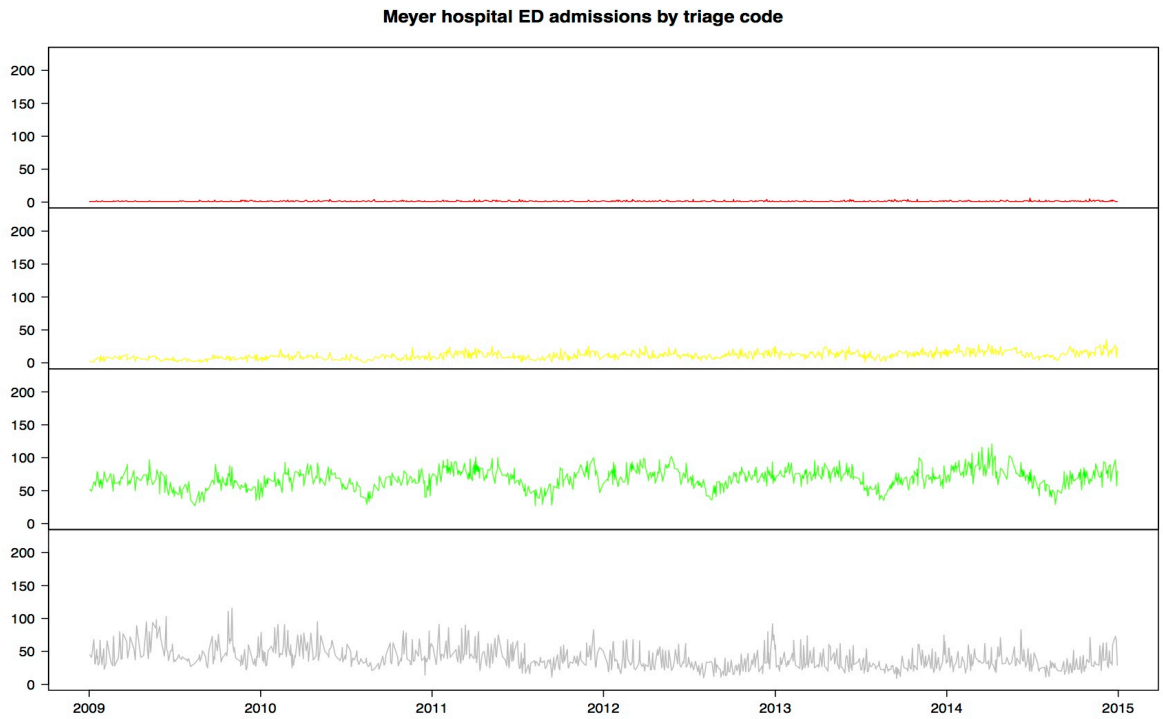

Figure 7: Admissions to the Meyer's ED stratified by triage code.

The most urgent cases, red and yellow code, present a very similar cycle that differs from the less urgent ones, green and white, that behave in a similar way. However, the most outstanding fact is the difference in number of patients. It is easy to appreciate that the green and white codes present a considerably higher number of visits than the urgent cases. The elevated volume of these patients contributes in a major part to crowding, as these numbers represent that the vast majority of patients that show up in the ED should probably not be there as they are classified as non-urgent.

To illustrate this point, the following graph superposes the four colours.

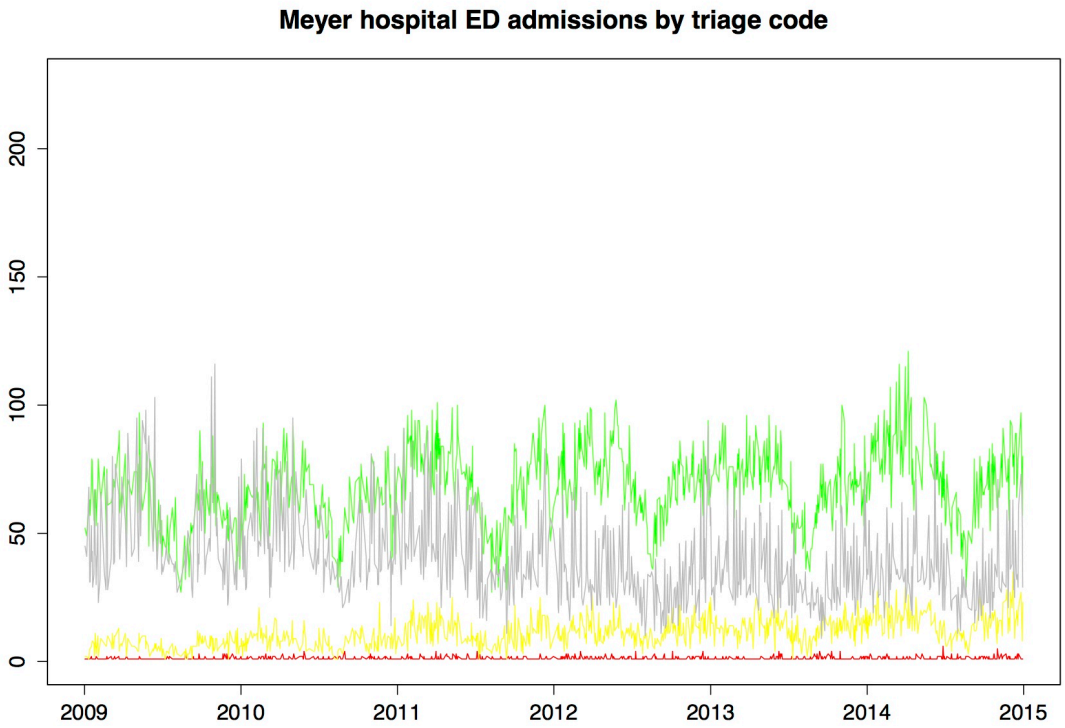

Figure 8: Superposition of patients admissions to the Meyer's ED by triage colour.

The next graph represents the ED visits by hour of the day. Every hour is represented by a boxplot that allows us to perceive the mean of the number of visits for the 6 years. Anomalies are easy to observe due to the fact that the represented data includes the admission of a lot of patients. It is from 10a.m. until 10p.m. that visits are more frequent. In contrast, from midnight until 7a.m. visits are rare.

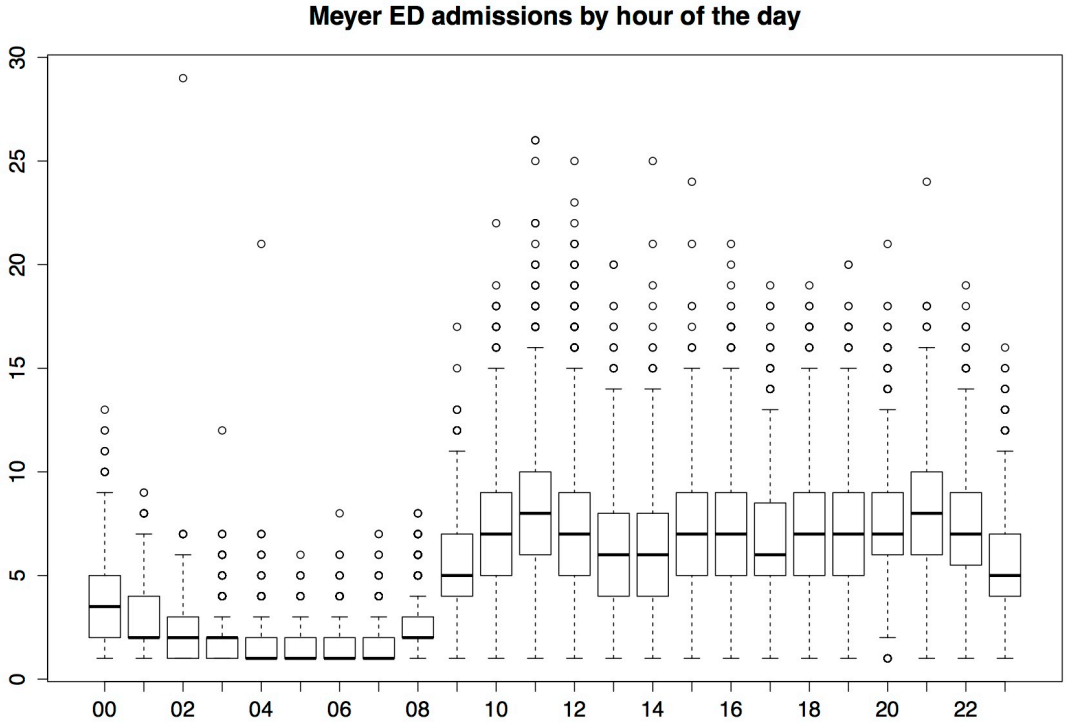

Figure 9: Patients' Admissions to the Meyer's ED by hour of the day.

Another interesting way of approaching the arrivals is taking a look at them by day of the week. As expected, it is during the weekends when visits are more frequent. Monday presents an imperceptible higher patient volume than the rest of the week, when visits seem to be nearly constant.

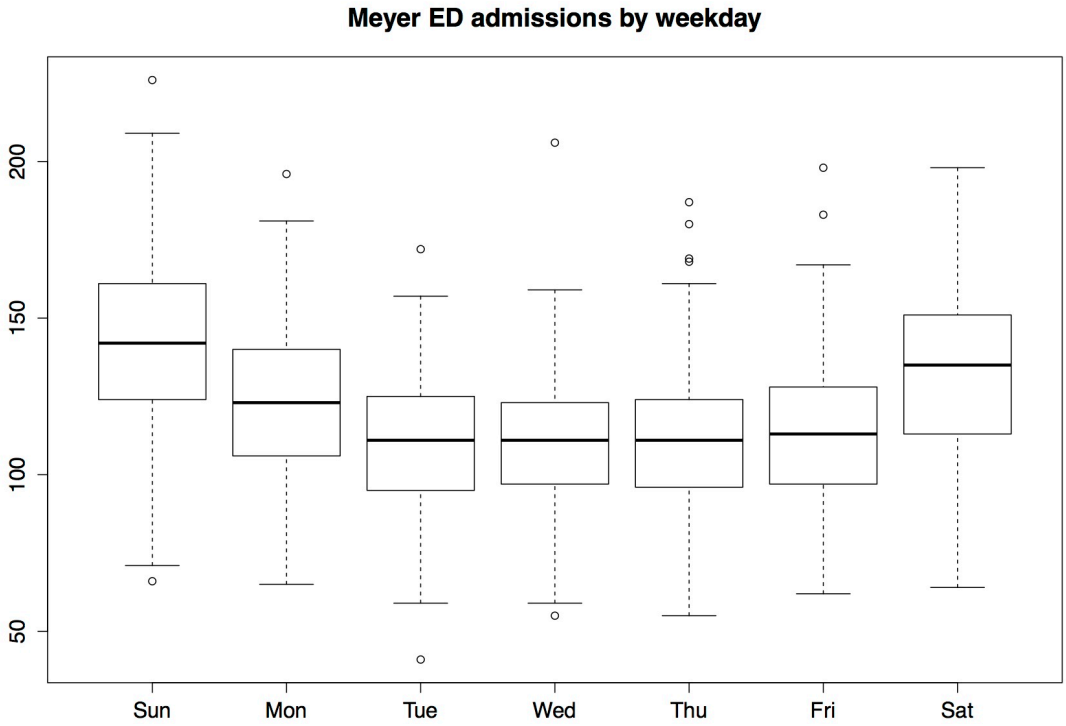

Figure 10: Patients' Admissions to the Meyer's ED by weekday.

4.3. Results of applying the methods to the Meyer's ED data

Once the performance of the hospital has been understood, it is necessary to continue the analysis with the aim of finding a model to forecast the next year's arrivals to the Meyer's ED. Data including arrivals from January 2009 to December 2014 has been used to develop the forecasting model and the remaining data, the first five months of 2015, has been used to evaluate the obtained results.

Instead of analysing data referring to patients' daily admissions, it has been considered more appropriate to work with monthly admissions. Using monthly data is enough to justify significantly the study that is being carried out. Regarding the daily data, the larger number of observations could cause difficulties throughout the analysis.

As it has already been mentioned, the graph representing monthly admissions can be easily used to detect if there is any kind of trend or seasonality depending on the season of the year.

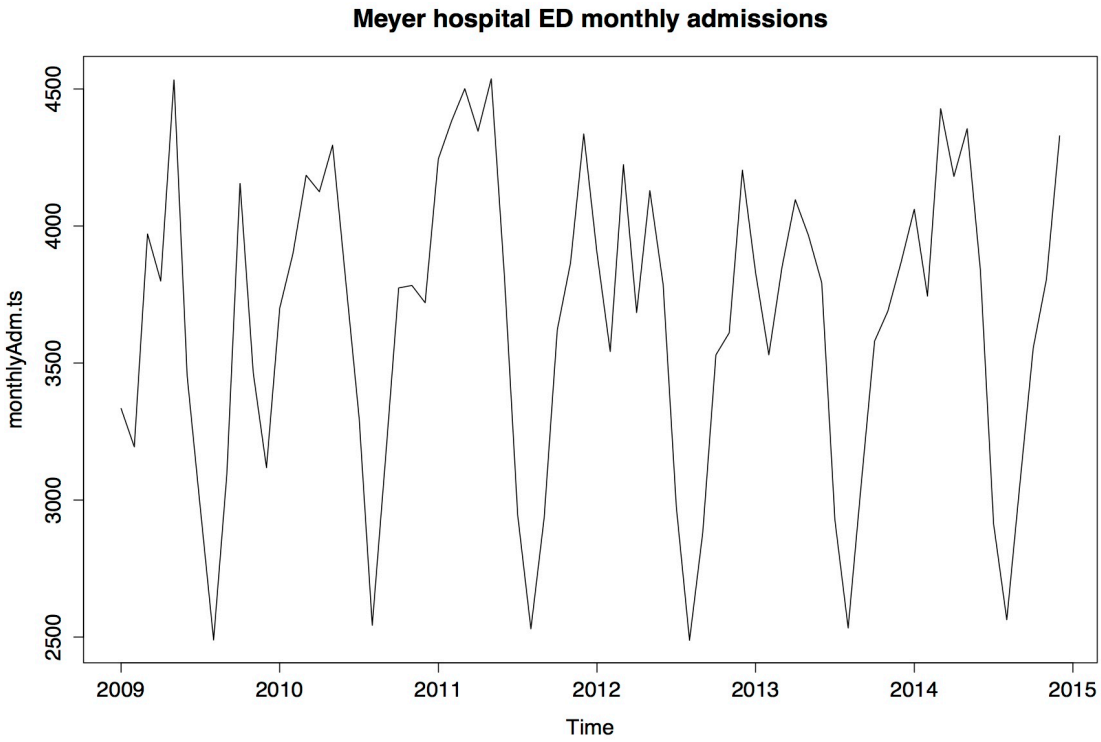

Figure 11: Monthly admissions to the Meyer's ED from January 2009 to December 2014.

From this time series plot it can be appreciated the existence of a seasonal variation in the number of visits per month: there is a peak every spring, and it falls to its minimum on August. This time series could probably be described using an additive model, as the seasonal fluctuations are roughly constant in size over time and do not appear to depend

on the level of the time series, the random fluctuations also seem to be roughly constant in size over time.

In a seasonal time series there is a trend component, a seasonal component and an irregular component. Decomposing the time series means separating the time series into these three components, which is estimating them.

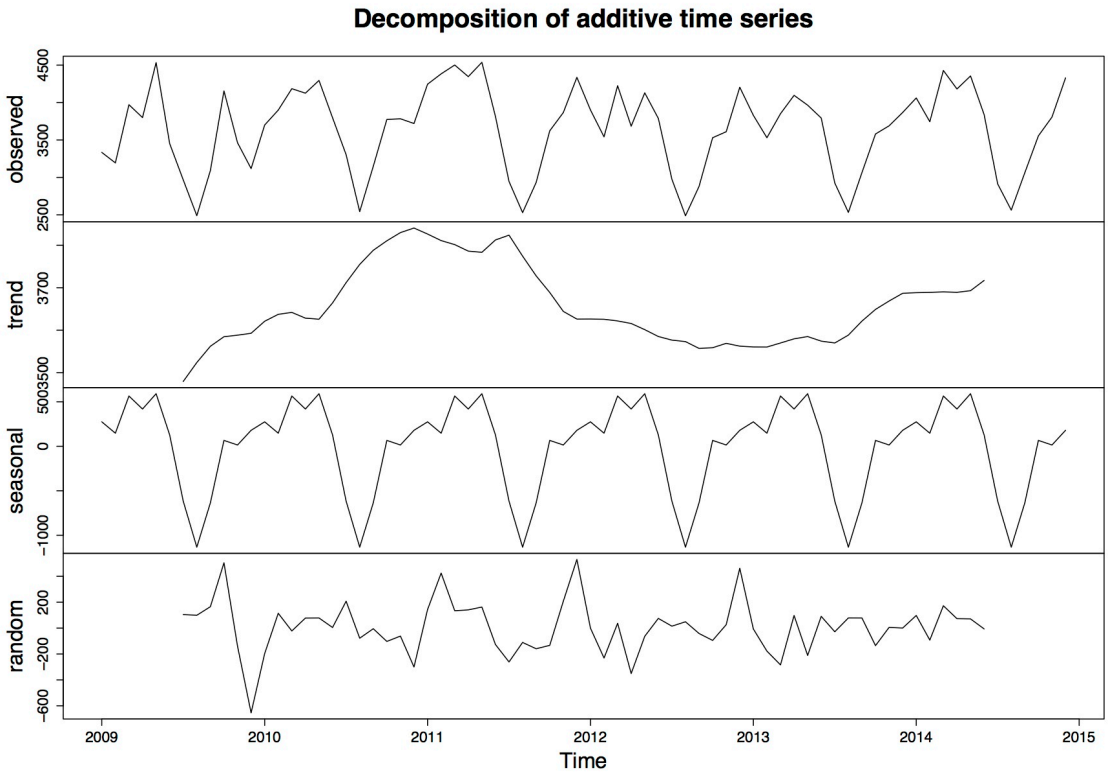

Figure 12: Decomposition of additive time series: Observed, Trend, Seasonal and random component.

The plot above presents the initial time series (observed), the estimated trend component (trend), the estimated seasonal component (seasonal), and the estimated irregular component (random). This graph states in a more clear way the fact of the existing seasonality.

4.3.1. Holt-Winters seasonal method

Once the Holt-Winters model has been done, the estimated values of alpha, beta and gamma are 0.22, 0.00, and 0.89, respectively.

Referring to the value of the parameters, alpha (0.22) is appreciably low meaning that the estimation of the level at the current time point is based in two recent observations and some others in the far away past.

Beta's value is 0.00, which indicates that the initial value of the trend is not updated through the time series. This could have been expected as the level changes a bit in the time series but the slope of the trend factor is almost the same.

Contrarily, the high value of gamma (0.89) is an indicator that the estimation of the seasonal component at the current time point is based in very recent observations.

For the Holt's exponential smoothing, the original time series (black line) is plotted against the forecasted values (red line):

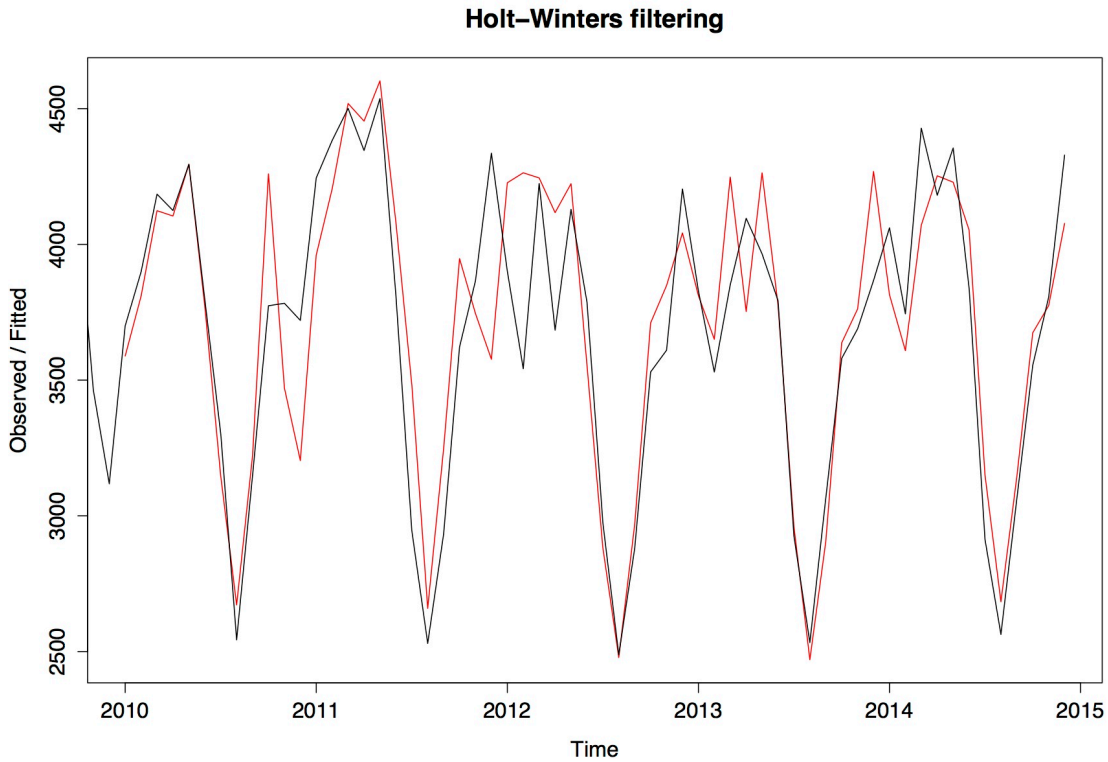

Figure 13: Holt-Winters - Filtering: Observed values vs. Forecasted values.

From figure 13 it can be highlighted that this method is very successful in predicting the seasonal falls, which occur in August every year.

The forecast of the next two years corresponds to the following plot:

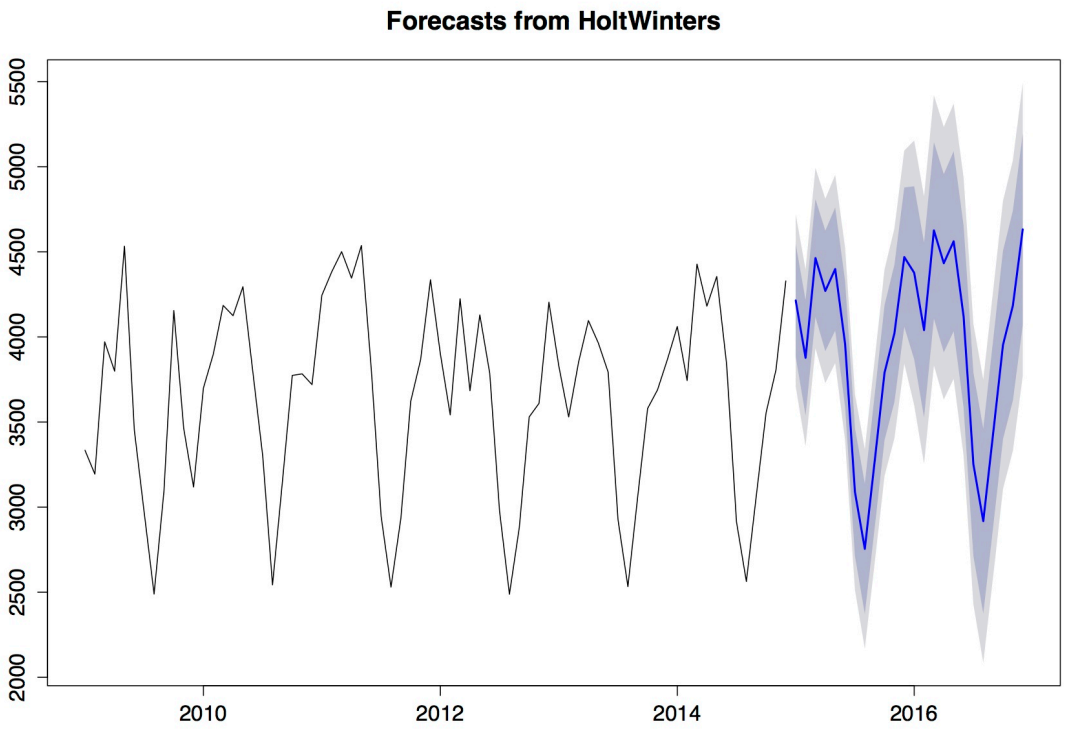

Figure 14: Holt-Winters - Forecasted values 2015-2016.

Forecasts are represented with a blue line, and the dark grey and light grey shaded areas correspond to the 80% and 95% prediction intervals, respectively.

The following table contains the exact predicted value, the observed value and the low and high value of confidence interval.

Month	Forecasted Value	Observed Value	Lo 95	Hi 95	|Forecast-Observed|	Error (%)
January 2015	4,214	4,400	3,708	4,720	186	6.77
February 2015	3,877	3,741	3,359	4,395	136	3.64
March 2015	4,463	4,449	3,933	4,993	14	0.31
April 2015	4,270	4,085	3,729	4,811	185	4.53
May 2015	4,399	4,232	3,846	4,951	167	3.95

Table 2: HW model - Forecasted value, observed value, CI 95% and forecasting error from the first five months of 2015.

The fact that the observed value of the admissions in 2015 is inside the confidence interval states a good prediction by the Holt-Winters model.

To validate the predicting model and following the recommendation of the book "Forecasting principles and practice", the correlogram and the Ljung-Box have been used to confirm if the forecast errors present a non-zero correlation at lags 1-24.

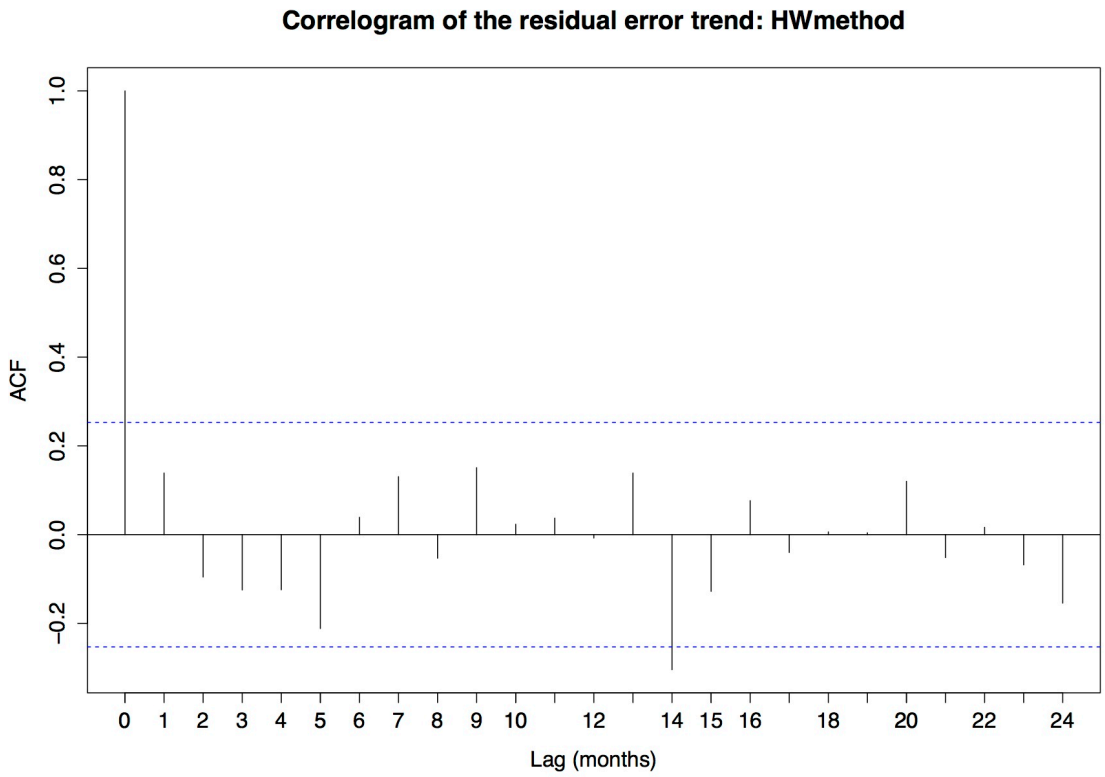

Figure 15: Holt-Winters correlogram of the residual error trend.

It can be observed that only at lag 14 the forecast errors exceed the significance bounds and the p-value of the Ljung-Box test is 0.3663. Both of these facts indicate that there is little evidence of non-zero autocorrelations at lags 1-24.

As suggested by the book "A little book of r for time series" by Coghlan, by making a time plot of the forecast errors and a histogram it can be checked if the forecast errors have constant variance over time, and are normally distributed with zero mean.

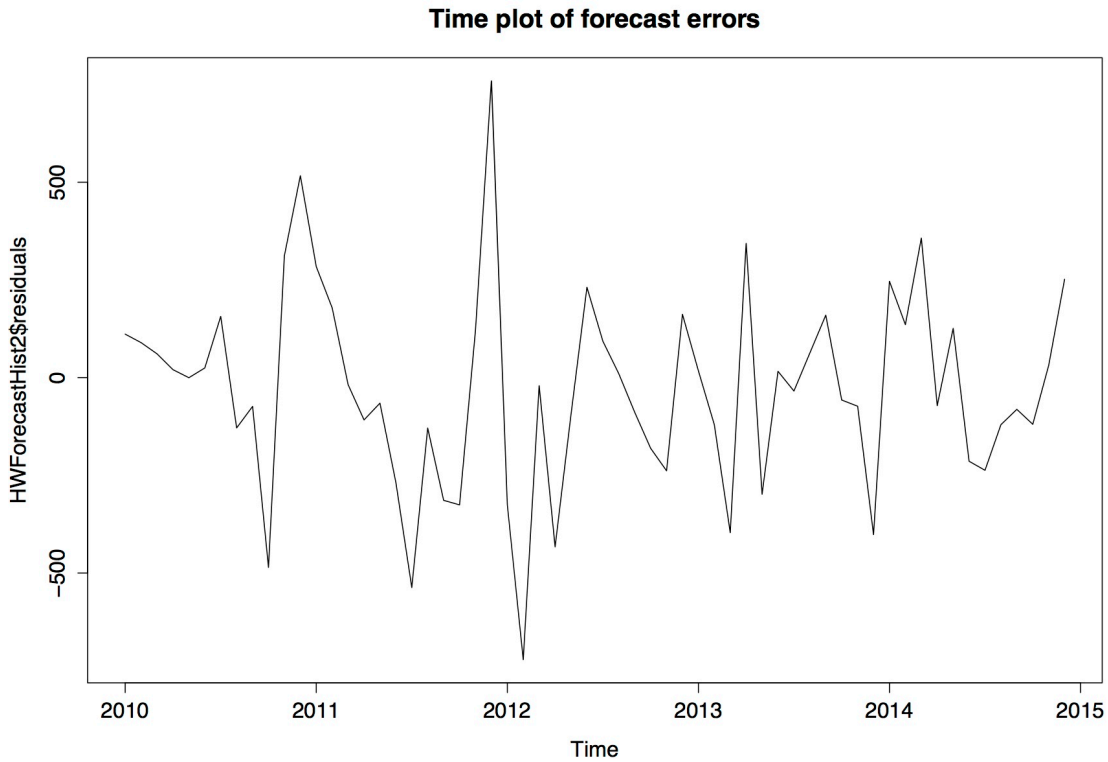

Figure 16: Holt-Winters - Time Plot of forecasted errors.

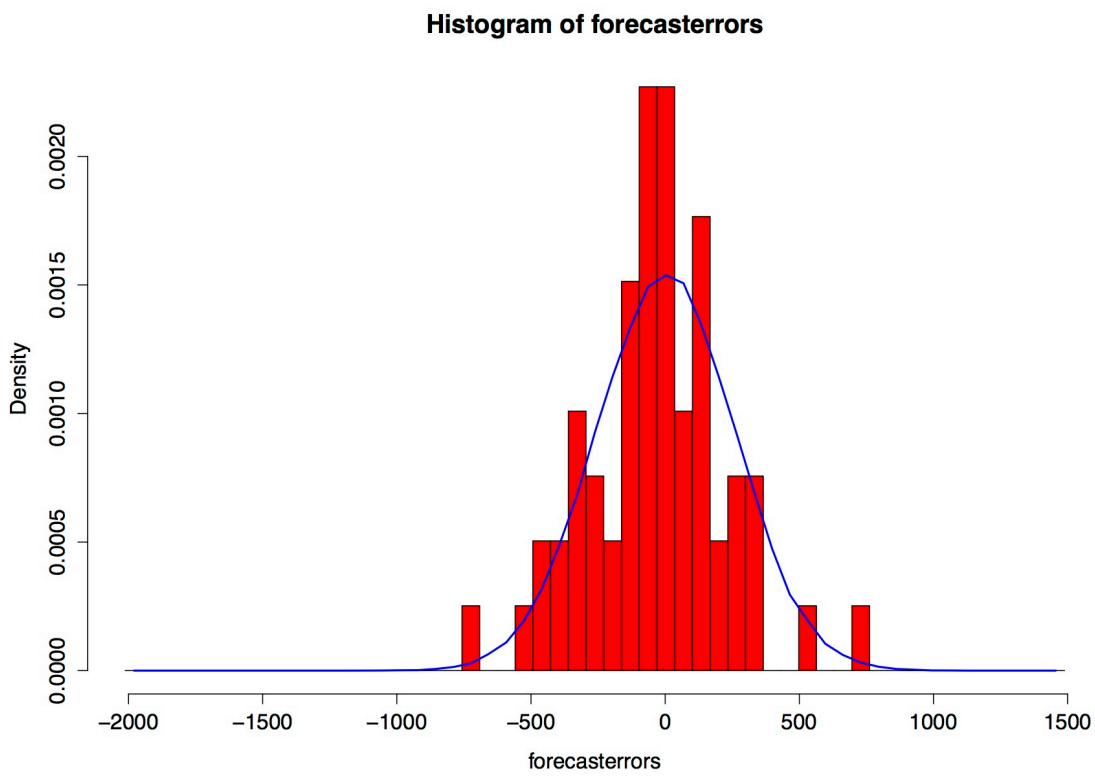

Figure 17: Holt-Winters - Histogram of forecasted errors.

Figure 16 shows that the forecast errors have constant variance over time. As for the histogram, the forecast errors follow a normal distribution with a very close to zero mean, with an exact value of -31.77. Although it is not exactly zero, this small difference (compared to the high number of observations) can be corrected by subtracting 31.77 to the limit bounds of the CI. Even though after doing so, the real values are still inside the CI.

These facts lead us to the conclusion that there is little evidence of autocorrelation at lags 1-24 for the forecast errors, they appear to be normally distributed with almost zero mean and constant variance over time. For this reason, it is considered that the Holt-Winters exponential smoothing provides a good predictive model of the monthly admissions to the Meyer ED. The assumptions in which the prediction intervals were base are considered valid.

4.3.2. Seasonal ARIMA model

According to the theory behind the model, ARIMA models are defined for stationary time series so the first step that needs to be done before performing the model is discuss whether the data is stationary or not.

From Figure 12, the 'Decomposition of additive time series', it can be noticed that the original data has a clear seasonal pattern which means that cannot be assumed to be stationary. Consequently, it is required to difference the time series to reach stationarity.

As well as looking at the time plot of the data, the ACF plot is also useful for identifying non-stationary time series. For a stationary time series, the ACF drops to zero relatively quickly, while the ACF of non-stationary data decreases slowly or never drops to zero.

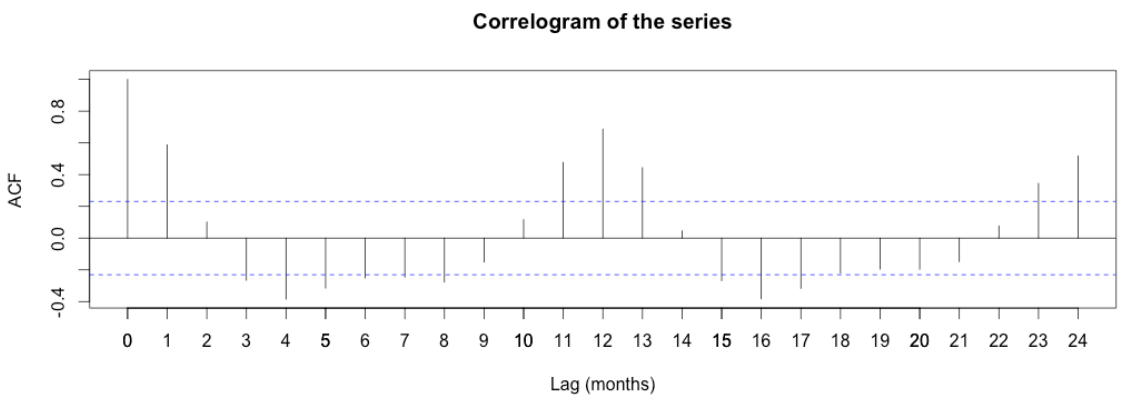

Figure 18: Correlogram of the monthly admissions to the Meyer's ED.

To make the data stationary, it is necessary to differentiate it. Nevertheless, if the data has a strong seasonal pattern, seasonal differencing is recommended to be done before first difference as the resulting series could already become stationary. If first differencing is done first, seasonality can still be present.

The next plot shows the data subject of study with a seasonal differencing applied.

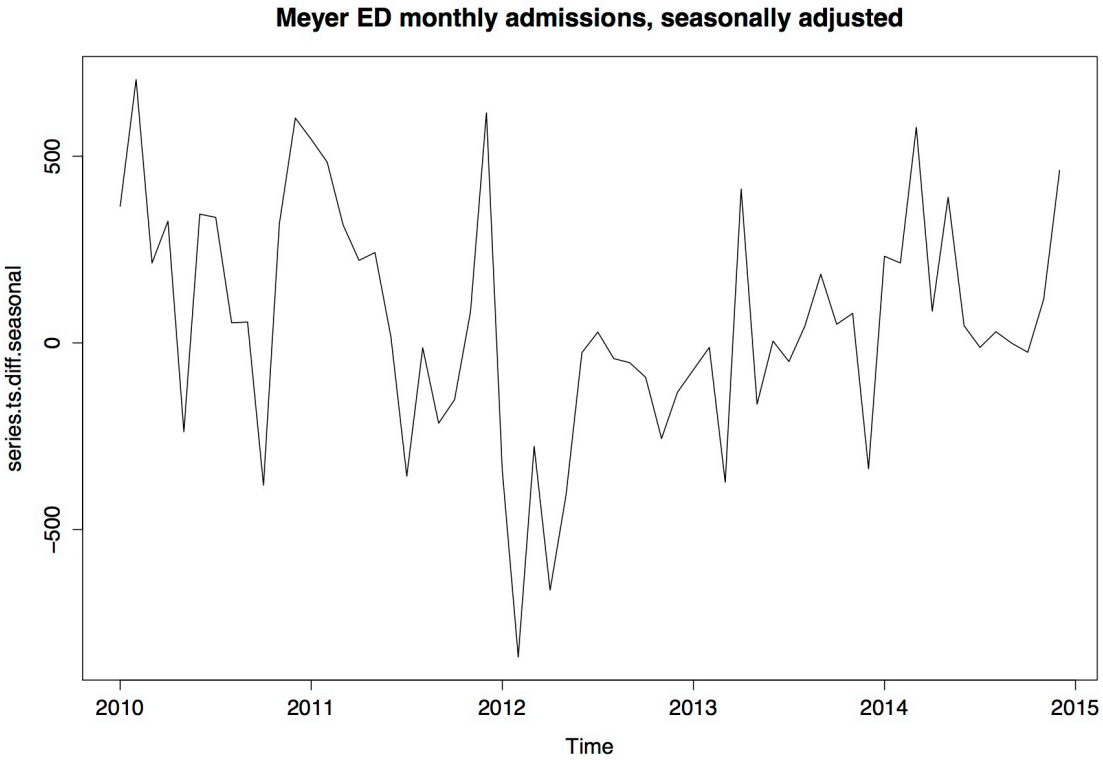

Figure 19: Monthly seasonally adjusted admissions to the Meyer's ED.

The time series of the seasonal difference appears to be stationary in mean and variance, so an ARIMA (P,1,Q) model is probably appropriate for the monthly admissions. By eliminating the seasonal component of the time series, we are left with an irregular component.

It can be now examined whether there are correlations between successive terms of this irregular component; if so, this could help to make a predictive model for the monthly admissions to ED.

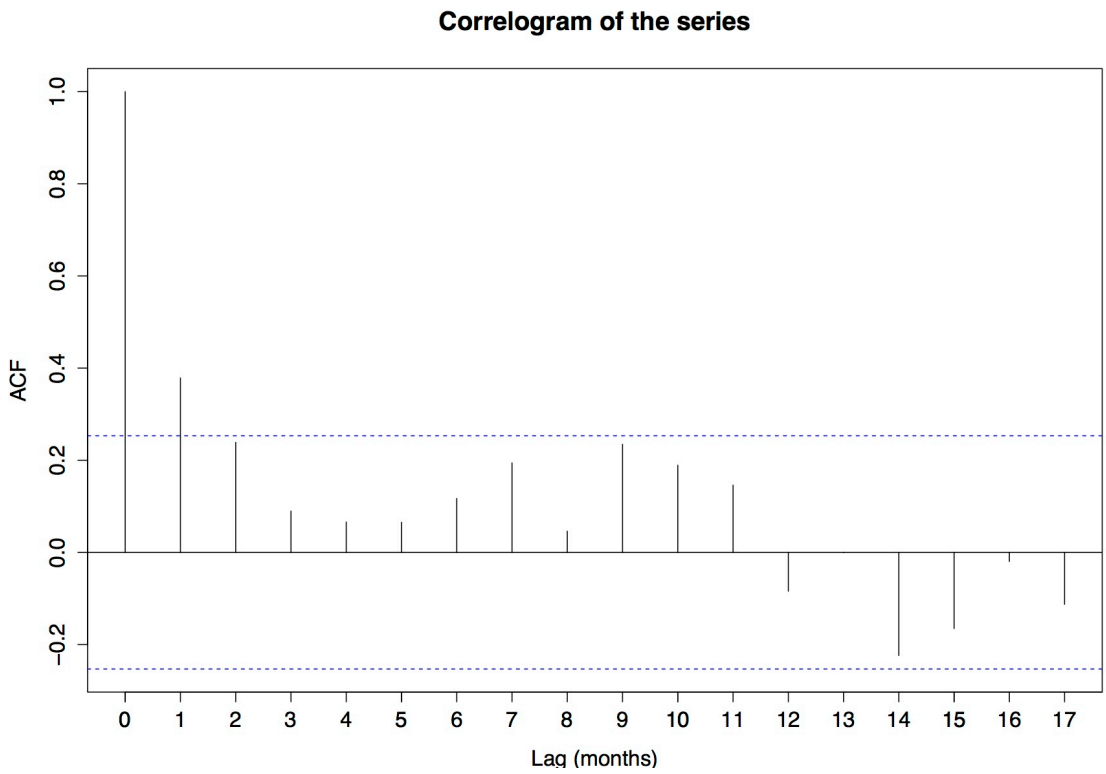

Figure 20: Correlogram of the monthly seasonally adjusted admissions to the Meyer's ED.

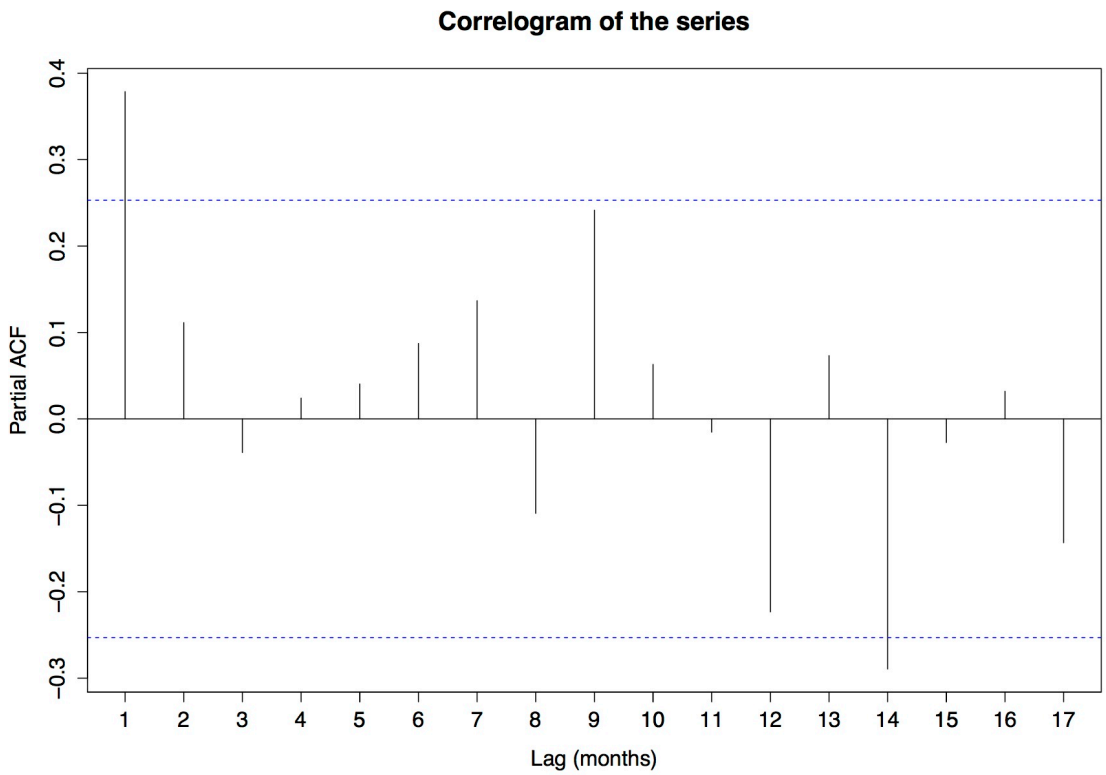

Figure 21: Partial correlogram of the monthly seasonally adjusted admissions to the Meyer's ED.

The current aim is to find an appropriate ARIMA model based on the graphics shown above: the ACF and the PACF. The significant spike at lag 1 in the ACF suggests a non-seasonal MA(1) component. Consequently, we begin with an ARIMA(0,0,1)$(0,1,0)_{12}$ model, indicating a seasonal difference, and non-seasonal MA(1) component.

The PACF shows a significant spike at lag 1 in the non-seasonal component and at lag 2 in the seasonal component suggesting a seasonal AR(2) component. Taking into account the PACF graph, the model should be ARIMA(1,0,0)$(2,1,0)_{12}$.

A very useful tool to compare models with the same order of difference is the AICc, a low value of this parameter indicates a better predictive model. In the case of the ARIMA(0,0,1)$(0,1,0)_{12}$ model, the AICc is 687.75, while in the ARIMA(1,0,0)$(2,1,0)_{12}$ model is 680.7. Other models with AR terms have also been tried, but none of them present a smaller AICc value. Consequently, the chosen model is ARIMA(1,0,0)$(2,1,0)_{12}$. Its residuals are plotted in Figure 22. All the spikes are now within the significance limits, and so the residuals appear to be white noise.
Since the correlogram shows that none of the sample autocorrelations for lags 1-24 exceed the significance bounds, and the p-value for the Ljung-Box test is 0.0975, it can be concluded that there is very little evidence for non-zero autocorrelations of the error terms at lags 1-24 in this model.

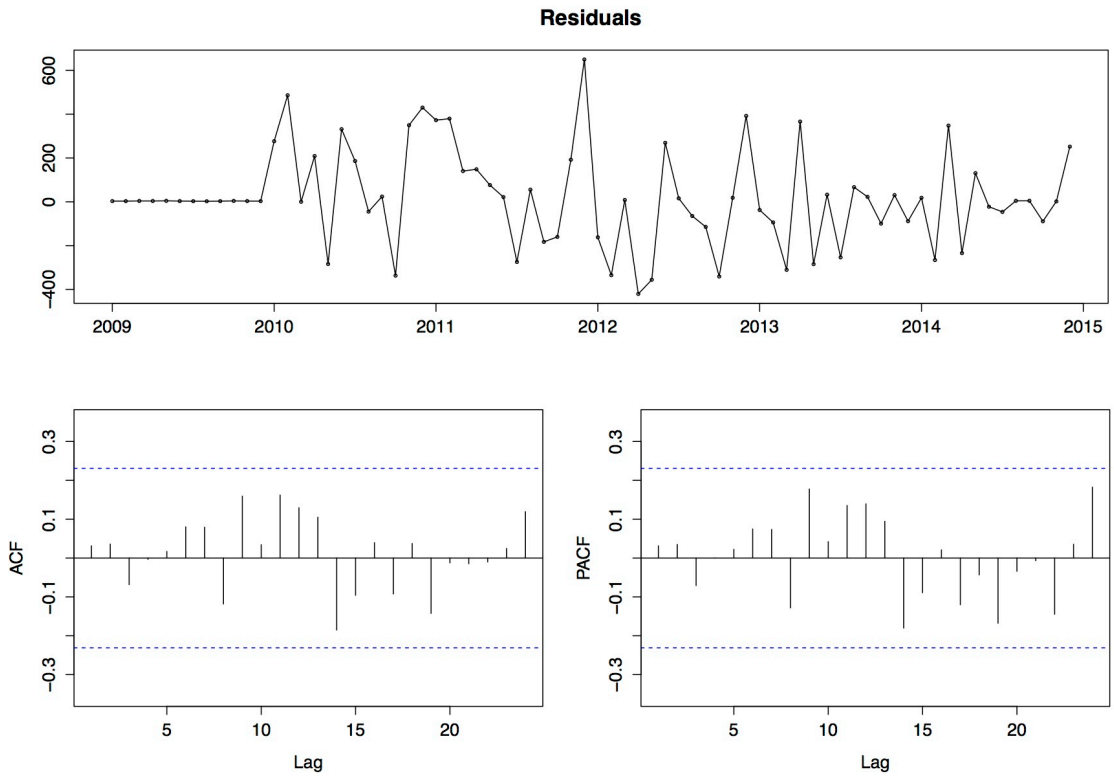

Figure 22: Plot, correlogram and partial correlogram of the forecasting residuals.

The auto.arima() function could have been used to do most of this work as it selects the best ARIMA model according to a specific data. In this case, it gives the same result.

At this point, the seasonal ARIMA model $(1,0,0)(2,1,0)_{12}$ passes the required checks and is ready to forecast the data of the monthly admissions to the Meyer's ED. The predictions will then be checked with the first five months of 2015 admissions.

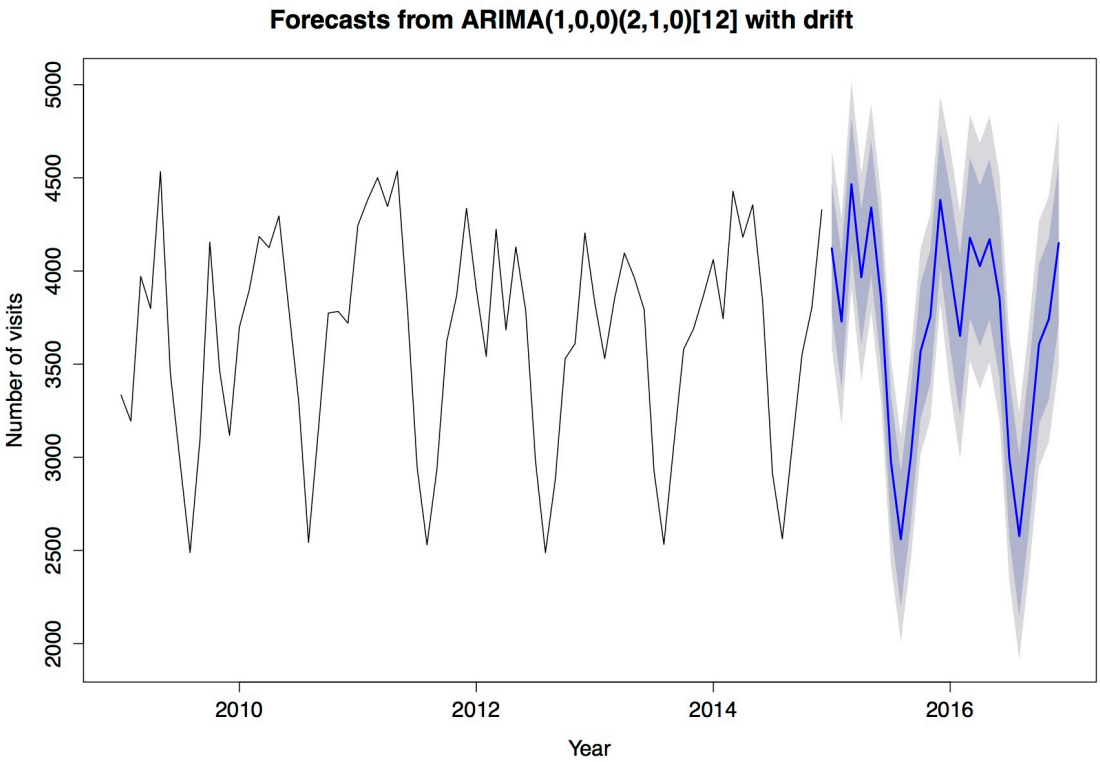

Figure 23: ARIMA model - Forecasted values 2015-2016.

Forecasts from the model for the next two years are shown in Figure 23.

As it has formerly been suggested, the forecasting needs to be checked. The real values of the visits for the five first months of 2015 are included in their respective CI. It can be concluded that this model has performed a good forecasting.

As follows the exact values of the forecasted points, the observed values of 2015 admissions and the prediction intervals with 95% of confidence are shown.

Month	Forecasted Value	Observed Value	Lo 95	Hi 95	|Forecast-Observed|	Error (%)
January 2015	4,121	4,400	3,593	4,649	279	6.34
February 2015	3,728	3,741	3,179	4,277	12	0.32
March 2015	4,465	4,449	3,914	5,016	17	0.38
April 2015	3,966	4,085	3,415	4,517	118	2.89
May 2015	4,340	4,232	3,789	4,891	109	2.58

Table 3: ARIMA model - Forecasted value, observed value, IC 95% and forecasting error from the first five months of 2015.

The graph below plots the standardized residuals, the autocorrelation function of the residuals and the p-values of the Ljung–Box test.

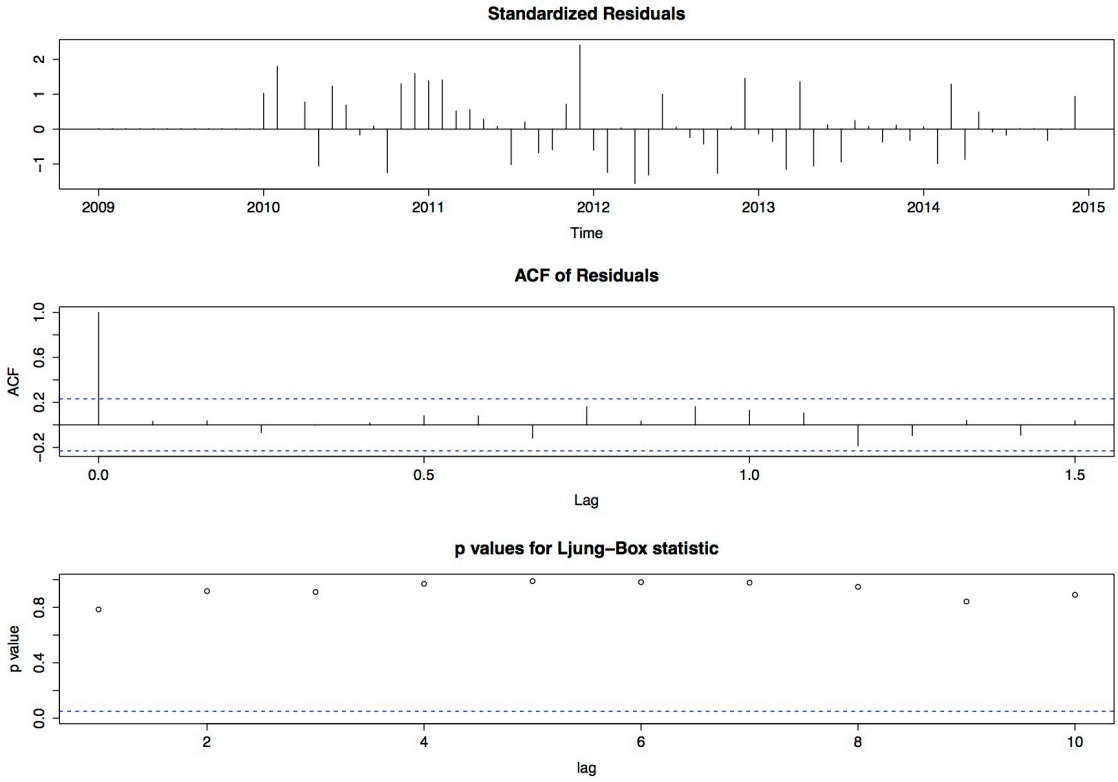

Figure 24: ARIMA model - Standardized residuals, correlogram of residuals and p-values for the Ljung-Box statistic.

The correlogram shows that the autocorrelations for the forecast errors do not exceed the significance bounds for lags 1-24. Furthermore, the p-value for Ljung-Box test is 0.0975, indicating that there is little evidence of non-zero autocorrelations.

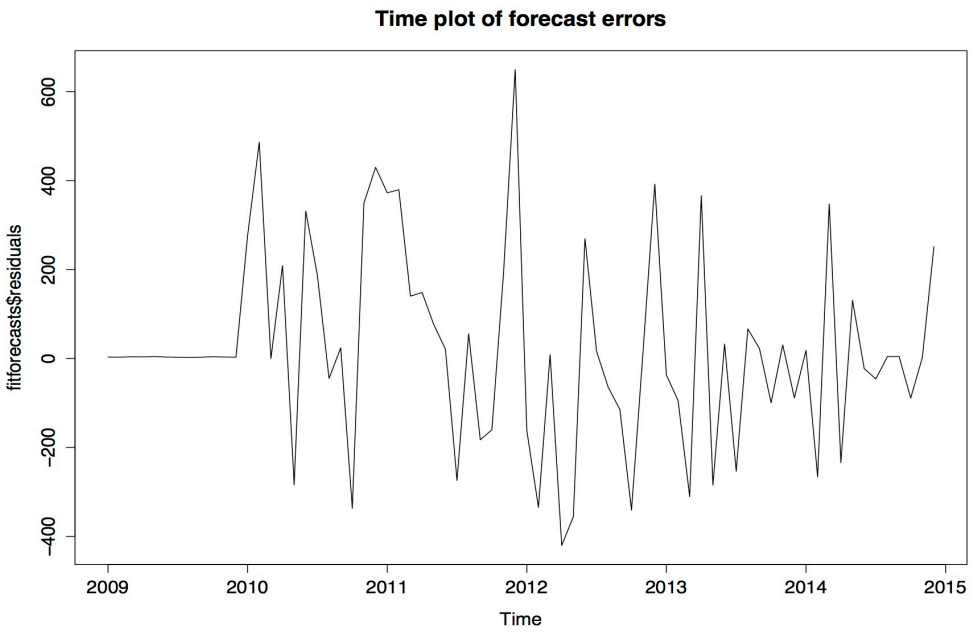

Figure 25: ARIMA model - Time plot of the forecast errors.

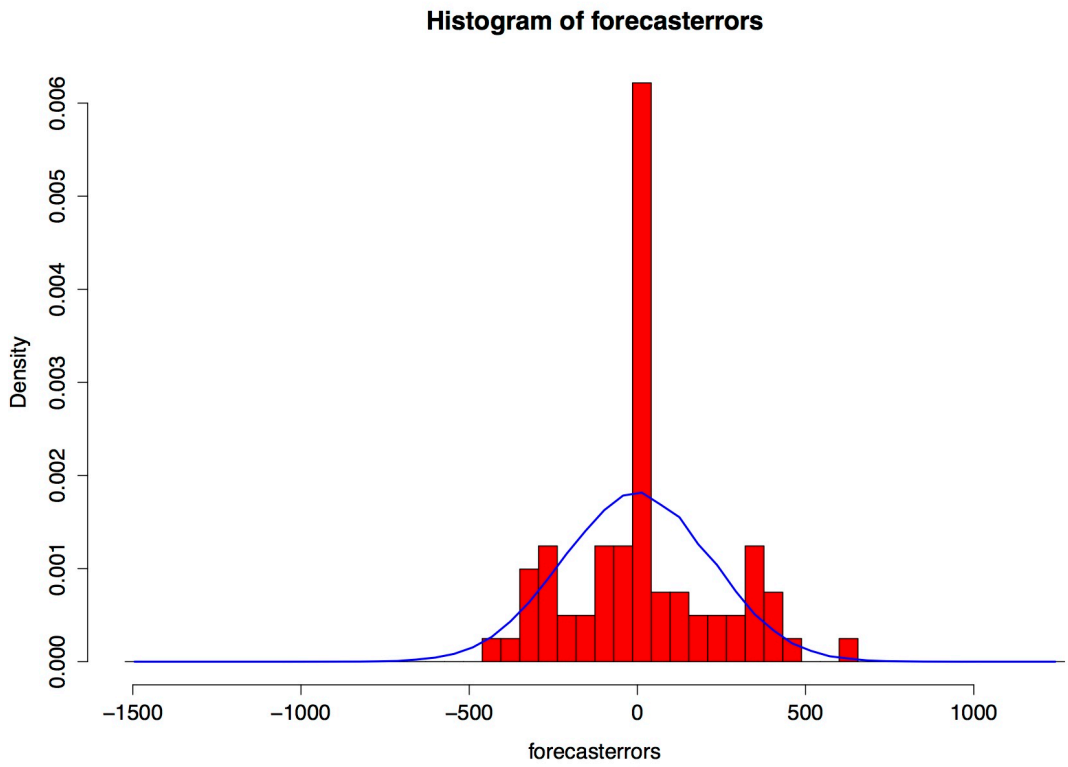

Figure 26: ARIMA model - Histogram of the forecast errors.

The time plot demonstrates that the variance of the forecast errors seems to be roughly constant over time (though perhaps there is slightly lower variance on the second half of the time series). From the histogram of the time series, it is deduced that the forecast errors are approximately normally distributed with a small mean of 20.20.

Since successive forecast errors do not seem to be correlated and are normally distributed with mean almost zero and constant variance, the ARIMA(1,0,0)(2,1,0)$_{12}$ does seem to provide an adequate predictive model for the monthly arrivals to the ED of the hospital.

The values of the residuals are shown in Table 4.

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sept	Oct	Nov	Dec
2009	3.33	3.19	3.97	3.79	4.53	3.44	2.96	2.48	3.08	4.14	3.45	3.1
2010	276.50	486.17	0.21	208.88	-283.73	331.48	186.2	-44.65	23.97	-336.68	349.5	429.96
2011	372.69	379.53	140.38	148.46	76.3	21.59	-274.22	55.49	-182.65	-159.99	191.98	649.38
2012	-161.8	-334.64	8.36	-420.23	-355.55	269.32	15.88	-64.65	-114.3	-340.78	18.5	391.96
2013	-37.17	-94.1	-310.43	366.03	-284.44	32.76	-253.3	66.83	22.58	-99.42	30.64	-88.28
2014	18.41	-265.59	347.51	-234.03	130.84	-22.24	-45.97	4.54	4.72	-88.87	2.02	251.81

Table 4: ARIMA model - Residuals from 2009 to 2014.

Choosing this model has been preceded by a wide analysis of different models, not only considering different ARIMA parameters but also considering different data.

As it has already been explain, the final consideration has been fitting the model with the 6 complete years of data and checking the results of the forecast with the first five months of 2015.

To choose the model, a comparison of some relevant parameters that will be explained below was carried out.

The following table summarizes all the models taken into account as well as their characteristic parameters that led to selection of the model.

	ARIMA MODEL	**AICc**	**p-value Ljung-Box test**	**Mean of the forecast errors**	**RMSE**
· Model developed with data from 2009 to 2014. · Forecasted values compared with January-March of 2015	**$(1,0,0)(2,1,0)_{12}$**	**680.70**	**0.0975**	**20**	**219.86**
	$(0,0,1)(0,1,0)_{12}$	687.75	1.775e-05	0.072	266.74
	$(0,0,1)(2,1,0)_{12}$	680.85	0.06355	23.10023	218.14
	$(1,0,0)(0,1,0)_{12}$	685.62	0.00026	-1.412431	260.72

Table 5: ARIMA model - Different developed models with its AICc, p-value Ljung-box test, mean of the forecast errors and RMSE.

The Root Mean Squared Error (RMSE) represents the sample standard deviation of the difference between the predicted values and the observed ones. It is considered a good measure of accuracy, but only to compare forecasting errors of different models for a particular variable. A low RMSE indicates a better model.

The Ljung –Box test that tells the level of correlation between the residuals, is not considered as important as the AICc or the RMSE.

Looking at the models that passed the Ljung-Box test, those with a p-value lower than 0.05 considering a 95% of confidence, it can be detected that they present the highest RMSE so are directly discarded.

As it has already been explained, to compare ARIMA models by the AICc values, the order of differencing must be the equal. The table includes models with only seasonal

differencing as well as models with both first and seasonal differencing. Nevertheless, the models developed with data from 2009 to 2014 all include a seasonal differencing as the original data had clearly a seasonal pattern.

It is difficult to find a model that passes all the residual tests, in fact, none of the models taken into consideration does. In practice, the chosen model will be the best model found, even if it does no pass all the tests.

Bearing all the parameters on mind, the chosen model, ARIMA(1,0,0)$(2,1,0)_{12}$, turns out to be the best one as its AICc and RMSE are lowest amongst models with only seasonal differencing. Even though the Ljung-Box test shows a little evidence of correlated errors, this has not affected the decision.
It should also be highlighted that there is another model, the ARIMA(0,0,1)$(2,1,0)_{12}$, which presents very similar RMSE and AICc values to the chosen one. This model has been discarded by looking at the lags of the ACF and PACF and the auto.arima() function, which fitted better with the other model.

It has also been considered interesting to develop ARIMA models for the data stratified by triage colour with the objective of trying to optimize the forecast.
The auto.arima() function has been applied and the following models have been obtained:

	ARIMA Model	**AICc**	**p-value Ljung-Box test**	**Mean of the forecast errors**	**RMSE**
Red	$(1,1,1)(2,0,0)_{12}$	426.44	0.0063	-0.004	4.51
Yellow	$(0,1,1)(2,0,0)_{12}$	785.88	0.0022	-0.67	57.63
Green	$(0,1,1)(2,1,0)_{12}$	756.11	0.00052	-15.15	113.27
White	$(0,1,2)(1,0,0)_{12}$	925.16	0.05162	0.17	154.99

Table 6: ARIMA model - ARIMA models for different triage colours.

As expected, the obtained models not only differ from the selected ARIMA model of the original data but they also differ between one and other. This behaviour is normal as the data considered for each model is different.

Comparing the AICc values has no sense as the order of differencing of each model is different and they all refer to different data.

Although it has been explained that a lower RMSE means a better forecast, in this case, the comparison is not meaningful as this parameter depends on the number of observation taken into consideration. For example, the red colour's RMSE, with a very low value, demonstrates that it is a good forecasting model itself. This low value is justified by the fact that it is the colour that presents fewer patients and the easiest to forecast. In contrast, the white code presents the highest RMSE.

Finally, to test the accuracy of the developed models, the real values of 2015 by triage code have been compared with the CI forecasted intervals resulting in an exit of all of the colours.

5. CONCLUSION

The present work has carried out a study based on data provided by a real hospital, the Meyer Hospital in Florence, including patients' visits from January 2009 to May 2015. The aim of this thesis was to understand the data's behaviour in order to develop a forecasting model to predict future visits and facilitate the hospital to deal with overcrowding.

After the analysis of the information and the necessary transformations of the data, two methods have been applied: the ARIMA and the Holt-Winters models. Even though each one of them provides a different approach, both of them have turned out to be good predictors of future visits.

Comparing the obtained results of both of them, the perception is that the ARIMA model developed a more accurate prediction. In fact, the ARIMA model presents a mean error of a 2.5% while the Holt-Winters experiences an error of the 3.84%

The following table summarizes all the relevant results obtained after applying the ARIMA model.

Month	Forecasted Value	Observed Value	Lo 95	Hi 95	|Forecast-Observed|	Error (%)
January 2015	4,121	4,400	3,593	4,649	279	6.34
February 2015	3,728	3,741	3,179	4,277	12	0.32
March 2015	4,465	4,449	3,914	5,016	17	0.38
April 2015	3,966	4,085	3,415	4,517	118	2.89
May 2015	4,340	4,232	3,789	4,891	109	2.58

Table 7: ARIMA model - Forecasted value, observed value, IC 95% and forecasting error from the first five months of 2015

The same is exposed for the Holt-Winters model.

Month	Forecasted Value	Observed Value	Lo 95	Hi 95	|Forecast-Observed|	Error (%)
January 2015	4,214	4,400	3,708	4,720	186	6.77
February 2015	3,877	3,741	3,359	4,395	136	3.64
March 2015	4,463	4,449	3,933	4,993	14	0.31
April 2015	4,270	4,085	3,729	4,811	185	4.53
May 2015	4,399	4,232	3,846	4,951	167	3.95

Table 8: HW model - Forecasted value, observed value, IC 95% and forecasting error from the first five months of 2015.

The accuracy of the forecasted values of the Holt Winters and ARIMA models have been calculated, tested and compared by means of MSE, MAE and MAPE.

Error measures			
	MSE	**MAE**	**MAPE**
Holt-Winters	66,539.01	194.37	1.01
ARIMA	48,340.5	156.26	0.99

Table 9: Forecast accuracy measures.

As it has previously been commented, the results of the error measures demonstrate that both models are effective and that the ARIMA model seems to be a more precise and accurate as it presents lower values of MSE, MAE and MAPE.

Figure 28 compares and contrasts the observed values for the five first months of 2015 with the forecasted values by both of the ARIMA and Holt-Winters models.

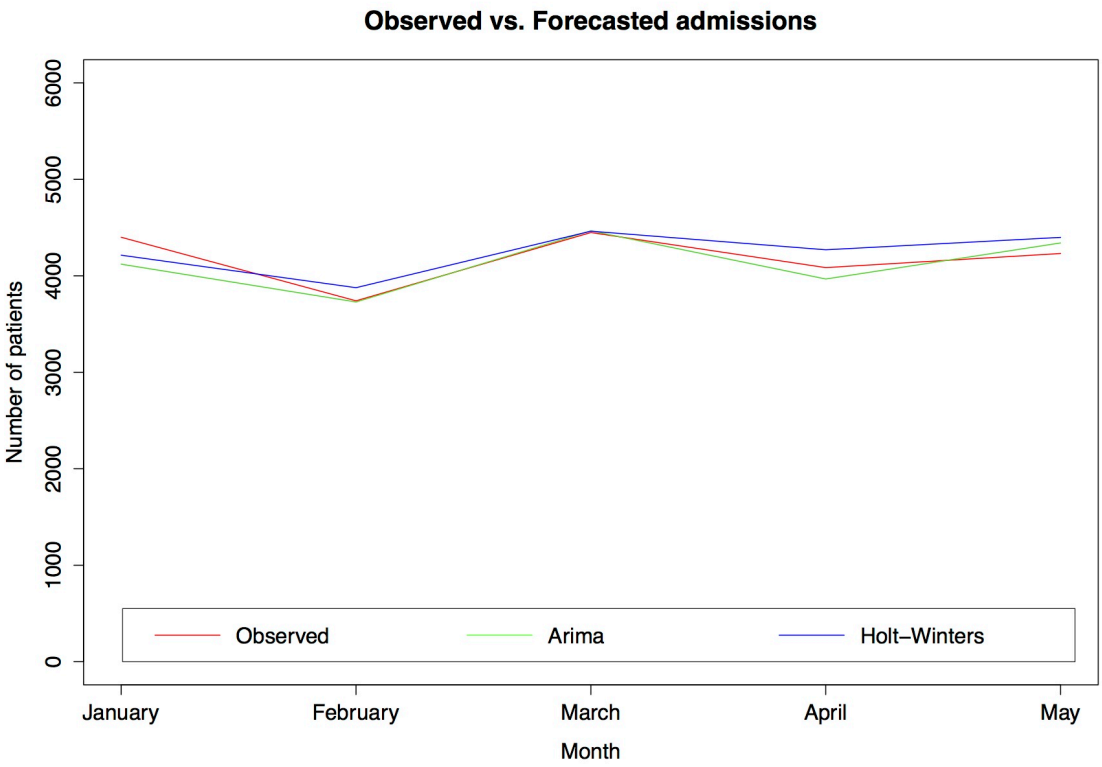

Figure 27: Comparative graph - Observed values 2015, forecasted values ARIMA and Holt-Winters model.

It is easy to observe that both of the models develop a very close prediction to the real values.

To conclude, as these two developed methods elaborate an accurate forecasting of the future visits, the aim of this thesis has been accomplished. The hospital has been provided a useful tool to try to deal with overcrowding. It is important to comment that these models have a clear limitation; they have been developed under a very specific data so they can only be applied to the Meyer Hospital.

REFERENCES

- Almehdawe, E., Jewkes, B., & He, Q. M. (2013). A Markovian queueing model for ambulance offload delays. *European Journal of Operational Research*, *226*(3), 602-614.

- Anneveld, M., van der Linden, C., Grootendorst, D., & Galli-Leslie, M. (2013). Measuring emergency department crowding in an inner city hospital in The Netherlands. Int J Emerg Med, 6, 21.

- Araz, O. M., Bentley, D., & Muelleman, R. L. (2014). Using Google Flu Trends data in forecasting influenza-like–illness related ED visits in Omaha, Nebraska. *The American journal of emergency medicine*, *32*(9), 1016-1023.

- Baubeau, D., Deville, A., Joubert, M., Fivaz, C., Girard, I., & Le Laidier, S. (2000). Les passages aux urgences de 1990 à 1998: une demande croissante de soins non programmés. DREES, Etudes et Résultats. Juillet, 72.

- Boyle, J. R., Sparks, R. S., Keijzers, G. B., Crilly, J. L., Lind, J. F., & Ryan, L. M. (2011). Prediction and surveillance of influenza epidemics. *Medical journal of Australia*, *194*(4), S28.

- Boyle, J., Wallis, M., Jessup, M., Crilly, J., Lind, J., Miller, P., & Fitzgerald, G. (2008, August). Regression forecasting of patient admission data. In *Engineering in Medicine and Biology Society, 2008. EMBS 2008. 30th Annual International Conference of the IEEE* (pp. 3819-3822). IEEE.

- Burr, T., Graves, T., Klamann, R., Michalak, S., Picard, R., & Hengartner, N. (2006). Accounting for seasonal patterns in syndromic surveillance data for outbreak detection. *BMC Medical Informatics and Decision Making*, *6*(1), 40.

- Chase, V. J., Cohn, A. E., Peterson, T. A., & Lavieri, M. S. (2012). Predicting emergency department volume using forecasting methods to create a "surge response" for noncrisis events. *Academic Emergency Medicine*, *19*(5), 569-576.

- Chen, C. F., Ho, W. H., Chou, H. Y., Yang, S. M., Chen, I. T., & Shi, H. Y. (2011). Long-term prediction of emergency department revenue and visitor volume using autoregressive integrated moving average model. *Computational and mathematical methods in medicine*, *2011*.

- Coghlan, A. (2014). A Little Book of R For Time Series. Wellcome Trust Sanger Institute: Cambridge, UK.

- Crane, P. W., Zhou, Y., Sun, Y., Lin, L., & Schneider, S. M. (2014). Entropy: A Conceptual Approach to Measuring Situation-level Workload Within Emergency Care and its Relationship to Emergency Department Crowding. The Journal of emergency medicine, 46(4), 551-559.

- Derlet, R. W., & Richards, J. R. (2008). Ten solutions for emergency department crowding. Western Journal of Emergency Medicine, 9(1), 24.

- Di Somma, S., Paladino, L., Vaughan, L., Lalle, I., Magrini, L., & Magnanti, M. (2014). Overcrowding in emergency department: an international issue. Internal and emergency medicine, 10(2), 171-175.

- Finamore, S. R., & Turris, S. A. (2009). Shortening the wait: a strategy to reduce waiting times in the emergency department. Journal of emergency nursing, 35(6), 509-514.

- Hoot, N. R., & Aronsky, D. (2008). Systematic review of emergency department crowding: causes, effects, and solutions. Annals of emergency medicine, 52(2), 126-136.

- Hoot, N. R., Epstein, S. K., Allen, T. L., Jones, S. S., Baumlin, K. M., Chawla, N., ... & Aronsky, D. (2009). Forecasting emergency department crowding: an external, multicenter evaluation. *Annals of emergency medicine*, *54*(4), 514-522

- Hyndman, R. J., & Athanasopoulos, G. (2014). Forecasting: principles and practice. OTexts.

- Institute of Medicine Committee on the Future of Emergency Care in the US Health System. (2006). Hospital-based emergency care: at the breaking point.

- Jones, S. S., Evans, R. S., Allen, T. L., Thomas, A., Haug, P. J., Welch, S. J., & Snow, G. L. (2009). A multivariate time series approach to modeling and forecasting demand in the emergency department. *Journal of Biomedical Informatics*, *42*(1), 123-139.

- 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

- Kecojevic, A., Lindsell, C. J., Lyons, M. S., Holtgrave, D., Torres, G., Heffelfinger, J., ... & National Emergency Department HIV Testing Consortium. (2011). Public Health and Clinical Impact of Increasing Emergency Department–Based HIV Testing: Perspectives From the 2007 Conference of the National Emergency Department HIV Testing Consortium. *Annals of emergency medicine*, *58*(1), S151-S159.

- Kline, J. A., Mark Courtney, D., Than, M. P., Hogg, K., Miller, C. D., Johnson, C. L., & Smithline, H. A. (2010). Accuracy of very low pretest probability estimates for pulmonary embolism using the method of attribute matching compared with the Wells score. *Academic Emergency Medicine*, *17*(2), 133-141.

- Laker, L. F., Froehle, C. M., Lindsell, C. J., & Ward, M. J. (2014). The flex track: flexible partitioning between low-and high-acuity areas of an emergency department.

- Laskowski, M., McLeod, R. D., Friesen, M. R., Podaima, B. W., & Alfa, A. S. (2009). Models of emergency departments for reducing patient waiting times.*PloS one*, *4*(7), e6127.

- Mahler, S. A., McCartney, J. R., Swoboda, T. K., Yorek, L., & Arnold, T. C. (2012). The impact of emergency department overcrowding on resident education. The Journal of emergency medicine, 42(1), 69-73.

- 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.

- McNaughton, C., Self, W. H., Jones, I. D., Arbogast, P. G., Chen, N., Dittus, R. S., & Russ, S. (2012). ED crowding and the use of nontraditional beds. *The American journal of emergency medicine*, *30*(8), 1474-1480.

- Miro, O., Sanchez, M., Espinosa, G., Coll-Vinent, B., Bragulat, E., & Milla, J. (2003). Analysis of patient flow in the emergency department and the effect of an extensive reorganisation. Emergency Medicine Journal, 20(2), 143-148.

- National Hospital Ambulatory Medical Care Survey: 2011 Emergency Department Summary Tables.

- Ong, M. E. H., Ho, K. K., Tan, T. P., Koh, S. K., Almuthar, Z., Overton, J., & Lim, S. H. (2009). Using demand analysis and system status management for predicting ED attendances and rostering. *The American journal of emergency medicine*, *27*(1), 16-22.

- Paul, J. A., & Lin, L. (2012). Models for improving patient throughput and waiting at hospital emergency departments. The Journal of emergency medicine, 43(6), 1119-1126.

- Peavey, E. K., Zoss, J., & Watkins, N. (2012). Simulation and mock-up research methods to enhance design decision making. *HERD: Health Environments Research & Design Journal*, *5*(3), 133-144.

- Pines, J. M., Hilton, J. A., Weber, E. J., Alkemade, A. J., Al Shabanah, H., Anderson, P. D., ... & Schull, M. J. (2011). International perspectives on emergency department crowding. Academic Emergency Medicine, 18(12), 1358-1370.

- Rappold, A. G., Cascio, W. E., Kilaru, V. J., Stone, S. L., Neas, L. M., Devlin, R. B., & Diaz-Sanchez, D. (2012). Cardio-respiratory outcomes associated with exposure to wildfire smoke are modified by measures of community health.*Environ Health*, *11*(71), 10-1186.

- Reis, B. Y., & Mandl, K. D. (2003). Time series modeling for syndromic surveillance. *BMC Medical Informatics and Decision Making*, *3*(1), 2.

- Reis, B. Y., & Mandl, K. D. (2004). Syndromic surveillance: the effects of syndrome grouping on model accuracy and outbreak detection. *Annals of emergency medicine*, *44*(3), 235-241.

- Saber Tehrani, A. S., Coughlan, D., Hsieh, Y. H., Mantokoudis, G., Korley, F. K.,

- Kerber, K. A., ... & Newman-Toker, D. E. (2013). Rising annual costs of dizziness presentations to US emergency departments. *Academic Emergency Medicine*, *20*(7), 689-696.

- Schneider, S. M., Gallery, M. E., Schafermeyer, R., & Zwemer, F. L. (2003). Emergency department crowding: a point in time. Annals of emergency medicine, 42(2), 167-172.

- 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.

- Tanabe, P., Gisondi, M. A., Medendorp, S., Engeldinger, L., Graham, L. J., & Lucenti, M. J. (2008). Should you close your waiting room? Addressing ED overcrowding through education and staff-based participatory research. Journal of Emergency Nursing, 34(4), 285-289.

- Vermeulen, M. J., Ray, J. G., Bell, C., Cayen, B., Stukel, T. A., & Schull, M. J. (2009). Disequilibrium between admitted and discharged hospitalized patients affects emergency department length of stay. *Annals of emergency medicine*,*54*(6), 794-804.

- Wargon, M., Casalino, E., & Guidet, B. (2010). From model to forecasting: a multicenter study in emergency departments. *Academic Emergency Medicine*,*17*(9), 970-978.

- Wargon, M., Guidet, B., Hoang, T. D., & Hejblum, G. (2009). A systematic review of models for forecasting the number of emergency department visits.Emergency Medicine Journal, 26(6), 395-399.

- Wathen, J. E., Gao, D., Merritt, G., Georgopoulos, G., & Battan, F. K. (2007). A randomized controlled trial comparing a fascia iliaca compartment nerve block to a traditional systemic analgesic for femur fractures in a pediatric emergency department. *Annals of emergency medicine*, *50*(2), 162-171.

- Wu, K. H., Cheng, F. J., Li, C. J., Cheng, H. H., Lee, W. H., & Lee, C. W. (2013). Evaluation of the effectiveness of peer pressure to change disposition decisions and patient throughput by emergency physician. *The American journal of emergency medicine*, *31*(3), 535-539.

- Websites:

    - https://www.hcup-us.ahrq.gov/db/state/sedddist/sedddist_visits.jsp

ANNEX

Packages used

```
install.packages("RODBC")
install.packages("zoo")
install.packages("xts")
install.packages("plyr")
install.packages("gplots")
install.packages("forecast")
install.packages("tseries")
library(RODBC)      #to access to DB
library(plyr)       #to make contingency tables
library(zoo)        #for time series
library(xts)        #for time series
library(gplots)     #to plot confidence intervals
library(forecast)
library(tseries)
```

Code for Data import from DB access and creation of the initial dataset

# Original field name	**Given name**	**Meaning**
#AnnoAccettazione	amdYear	Admission Year
#D01_COD_GRAVITA_ACCETTAZIONE	admTriaScore	Admission Triage Score
#TriageAmmissione	admTriaCode	Admission Triage Code
#D01_COD_GRAVITA_DIMISSIONE	disTriaScore	Discharge Triage Score
#TriageDimissione	disTriaCode	Discharge Triage Code
#D01_ID_ACCESSO	idPat	ID Patient
#D01_DATAORA_ACCETTAZIONE	admTime	Admission Time

```
#D01_DATAORA_MED_PRESA_CARICO      docTime        Patient seen by for first time
#OBI                               obi            Osservazione Breve Intensiva
#D01_COD_DIAGNOSI                  diaCode        Diagnosis Code
#D01_DATAORA_DIMISSIONE            disTime        Discharge Time
#D01_COD_MODALITA_DIMISSIONE       disCode        Discharge Code
#D27_DESC_MODALITA_DIMISSIONE      disCodeDesc    Discharge Code Description

#It is considered the access to the ED before 2015 and after 2008
  edData<-tblVisintin_PS[tblVisintin_PS$AnnoAccettazione!=2015 & tblVisintin_PS$AnnoAccettazione!=2008  ,]

#The columns of the dataframe are renamed
  colnames(edData)[which(names(edData) == "AnnoAccettazione")] <- "admYear"
  colnames(edData)[which(names(edData) == "D01_COD_GRAVITA_ACCETTAZIONE")] <- "admTriaScore"
  colnames(edData)[which(names(edData) == "TriageAmmissione")] <- "admTriaCode"
  colnames(edData)[which(names(edData) == "D01_COD_GRAVITA_DIMISSIONE")] <- "disTriaScore"
  colnames(edData)[which(names(edData) == "TriageDimissione")] <- "disTriaCode"
  colnames(edData)[which(names(edData) == "D01_ID_ACCESSO")] <- "idPat"
  colnames(edData)[which(names(edData) == "D01_DATAORA_ACCETTAZIONE")] <- "admTime"
  colnames(edData)[which(names(edData) == "D01_DATAORA_MED_PRESA_CARICO")] <- "docTime"
  colnames(edData)[which(names(edData) == "OBI")] <- "obi"
  colnames(edData)[which(names(edData) == "D01_COD_DIAGNOSI")] <- "diaCode"
  colnames(edData)[which(names(edData) == "D01_DATAORA_DIMISSIONE")] <- "disTime"
  colnames(edData)[which(names(edData) == "D01_COD_MODALITA_DIMISSIONE")] <- "disCode"
  colnames(edData)[which(names(edData) == "D27_DESC_MODALITA_DIMISSIONE")] <- "disCodeDesc"

#Formats of dates and hours are changed
  edData$admTime<-as.character(edData$admTime)
  edData$docTime<-as.character(edData$docTime)
  edData$disTime<-as.character(edData$disTime)
  edData$admTime=as.POSIXct(strptime(edData$admTime, "%d/%m/%Y %H:%M:%S"))
  edData$docTime=as.POSIXct(strptime(edData$docTime, "%d/%m/%Y %H:%M:%S"))
```

```
edData$disTime=as.POSIXct(strptime(edData$disTime, "%d/%m/%Y %H:%M:%S"))

edData$admDate<-as.POSIXct(round(edData$admTime,"days"))  #Admission time rounded to the day
edData$admDateAndHour<-as.POSIXct(round(edData$admTime,"hours")) #Admission time rounded to the hour
edData$admDay<-as.numeric(round(difftime(edData$admTime,as.POSIXct("2009-01-01"), units="days") , digits = 0))
admTime<-edData$admTime

edData$disCode<-as.factor(edData$disCode) #The discharge code is transformed in a factor
edData$admDate<-as.Date(edData$admTime, tz="Europe/Berlin") #Conversion of a POSIXct object in a date, necessary to specify time zone
edData$admDateAndHour<-as.POSIXct(round(edData$admTime,"hours")) #Admission time rounded to the hour
edData$admDay<-as.numeric(round(difftime(edData$admTime,as.POSIXct("2009-01-01"), units="days") , digits = 0))
edData$admWeekDay<-format(edData$admTime, "%a")
edData$admWday<-format(edData$admTime, "%w")
edData$admMonth<-format(edData$admTime, "%m")
edData$admHour<-format(edData$admTime, "%H")
```

Code for creating daily and monthly time series using the default package + decomposition

```
dailyAdm<- ddply(edData, c("admDate"),summarise,n=length(idPat))
monthlyAdm<- ddply(edData,c("admYear","admMonth"),summarise,n=length(idPat))
arrange(monthlyAdm,monthlyAdm$admYear,monthlyAdm$admMonth)

startdate=c(2009,1,1)
dailyAdm.ts<-ts(dailyAdm$n,start=startdate,frequency=365+1* (!startdate[1]%%400 || ((startdate[1]%%100)&&!startdate[1]%%4) ))
#Take into account the leap years
monthlyAdm.ts<-ts(monthlyAdm$n,start=startdate,frequency=12)

#Time series decomposition
dailyAdm.ts.decom<-decompose(dailyAdm.ts)
dailyAdm.ts.trend<-dailyAdm.ts.decom$trend
```

```
  dailyAdm.ts.seasonal<-dailyAdm.ts.decom$seasonal
  dailyAdm.ts.random<-dailyAdm.ts.decom$random
  dailyAdm.ts.plot<-plot(dailyAdm.ts,main="Meyer ED daily admissions", ylab="")
  dailyAdm.ts.decom.plot<-plot(dailyAdm.ts.decom)
  dailyadm.boxplot<-boxplot(dailyAdm.ts ~ cycle(dailyAdm.ts))

#Seasonally adjusted time-series and plot
  dailyAdm.ts.seasonalAdjusted<-dailyAdm.ts-dailyAdm.ts.seasonal
  dailyAdm.ts.seasonalAdjusted.plot<-plot(dailyAdm.ts.seasonalAdjusted, main="Meyer ED daily admissions, seasonally adjusted")

###### Code for creating daily, monthly and hourly time series (single and multiple) using the xts and zoo packages

  n=length(edData$admTime)
  v=rep(1,each=n)
  adm.xts<-xts(v, edData$admTime)
  rm(n,v)

  dailyAdm.xts<-apply.daily(adm.xts, FUN=sum) #Automatically calculates the daily admissions
  monthlyAdm.xts<-apply.monthly(adm.xts, FUN=sum)
  yearlyAdm.xts<-apply.yearly(adm.xts, FUN=sum)

#Total admission time-series plots
  dailyAdm.xts.plot<-plot(dailyAdm.xts, lty="solid", main="Meyer hospital ED daily admissions", ylab="",major.ticks="months",
major.format="%m/%y", cex.axis=0.6, las=2)
  monthlyAdm.xts.plot<-plot(monthlyAdm.xts, lty="solid", main="Meyer hospital ED monthly admissions", ylab="",major.ticks="months",
major.format="%m/%y", cex.axis=0.6, las=2)
  yearlyAdm.xts.plot<-plot(yearlyAdm.xts, lty="solid", main="Meyer hospital ED yearly admissions", ylab="", major.ticks="years",
major.format="%Y", cex.axis=0.6, las=2)

#Different time series according to the triage code
  dailyAdmRed<- ddply(edData[edData$admTriaCode=="ROSSO",], "admDate",summarise,nRed=length(idPat))
```

```
dailyAdmYellow<- ddply(edData[edData$admTriaCode=="GIALLO",], "admDate", summarise,nYellow=length(idPat))
dailyAdmGreen<- ddply(edData[edData$admTriaCode=="VERDE",], "admDate",summarise,nGreen=length(idPat))
dailyAdmWhite<- ddply(edData[edData$admTriaCode=="BIANCO",], "admDate",summarise,nWhite=length(idPat))


#An object xts(dataframe,date) is created with more time series where data frame contains columns corresponding to each time series
dailyAdmByTriaCode<-merge(dailyAdmRed,dailyAdmYellow)
dailyAdmByTriaCode<-merge(dailyAdmByTriaCode,dailyAdmGreen)
dailyAdmByTriaCode<-merge(dailyAdmByTriaCode,dailyAdmWhite)
dailyAdmTotData<-subset(dailyAdmByTriaCode, select = -admDate)
dailyAdmTotDate<-subset(dailyAdmByTriaCode, select = admDate)
dailyAdmByTriaCode.xts<-xts(dailyAdmTotData, dailyAdmTotDate$admDate)
dailyAdmByTriaCode.zoo<-as.zoo(dailyAdmByTriaCode.xts) #It is converted in a zoo object because it has moore graphic tools

dailyAdmByTriaCode.zoo
graphCol=c("red", "yellow", "green","grey")
ylimMin=0
ylimMax=max(dailyAdm$n)
main="Meyer hospital ED admissions by triage code"
xlab=""
ylab=""
dailyAdmByTriaCode.zoo.plot1<-plot(dailyAdmByTriaCode.zoo, col=graphCol,ylim=c(ylimMin,ylimMax),xaxp=c(ylimMin,ylimMax,3),
cex.axis=0.9, las=1, main=main,xlab=xlab, ylab=ylab)
dailyAdmByTriaCode.zoo.plot2<-plot(dailyAdmByTriaCode.zoo, screens=1,
col=graphCol,ylim=c(ylimMin,ylimMax),xaxp=c(ylimMin,ylimMax,3),xlab=xlab, main=main,ylab=ylab)
rm(ylimMin,ylimMax,main,xlab,ylab)
```

Code for creating Boxplot, Barplot, Stripcharts, Interval plots

```
#Hourly Admissions (Contingency tables are created with various descriptive statistics)
hourlyAdm<- ddply(edData, "admDateAndHour",summarise,n=length(idPat))
```

```
  hourlyAdm$h<-format(hourlyAdm$admDateAndHour, "%H")
  hourlyAdm.M<-tapply(hourlyAdm$n, hourlyAdm$h,mean)
  hourlyAdm.sd<-tapply(hourlyAdm$n, hourlyAdm$h, sd)
  hourlyAdm.le<-tapply(hourlyAdm$n, hourlyAdm$h, length)
  hourlyAdm.se<-hourlyAdm.sd/sqrt(hourlyAdm.le)
  hourlyAdm.cilb <- tapply(hourlyAdm$n, hourlyAdm$h, function(v) t.test(v)$conf.int[1]) #confidence interval lower bound
  hourlyAdm.ciub <- tapply(hourlyAdm$n, hourlyAdm$h, function(v) t.test(v)$conf.int[2]) #confidence interval upper bound

#Boxplot hourly Admissions
  hourlyAdm.boxplot<-boxplot(hourlyAdm$n ~ hourlyAdm$h, main="Meyer ED admissions by hour of the day")

#Barplot hourly Admissions
  hourlyAdm.barplot<-barplot(hourlyAdm.M, ylim=c(0,max(hourlyAdm.M)+1.5*max(hourlyAdm.sd)), main="Meyer ED admissions by hour of
the day")
  arrows(hourlyAdm.barplot,hourlyAdm.M,hourlyAdm.barplot,hourlyAdm.M+hourlyAdm.sd,lwd=1.5, angle=90, length=0.1)

#Stripchart (SE is small because of many observations)
  hourlyAdm.stripchart<-stripchart(hourlyAdm$n ~ hourlyAdm$h ,vert = TRUE, pch = 1, method = "jitter", jit = 0.05, xlab = "", ylab = "")
  points(1:24, hourlyAdm.M, pch = 16, cex = 1.5)
  arrows(1:24, hourlyAdm.M,1:24, hourlyAdm.M + hourlyAdm.se, lwd = 1.5, angle = 90, length = 0.1)
  arrows(1:24, hourlyAdm.M, 1:24, hourlyAdm.M - hourlyAdm.se, lwd = 1.5, angle = 90, length = 0.1)

#Intervalplot (CI is small because of many observations)
  hourlyAdm.intervalplot<-plot(y=hourlyAdm.M, x=rep(0:23,each=1), main="Meyer ED admissions by hour of the day: 95% Confidence
intervals for the mean", ylab="",xlab="", xaxp= c(0, 24, 24))
  arrows(0:23, hourlyAdm.M,0:23, hourlyAdm.ciub, lwd = 1.5, angle = 90, length = 0.1)
  arrows(0:23, hourlyAdm.M, 0:23, hourlyAdm.cilb, lwd = 1.5, angle = 90, length = 0.1)

#Weekday Admissions
  weekDayAdm<- ddply(edData,c("admWday", "admWeekDay","admDate"),summarise,n=length(idPat))
  weekDayAdm.M<-tapply(weekDayAdm$n, weekDayAdm$admWeekDay,mean)
```

```
  weekDayAdm.sd<-tapply(weekDayAdm$n, weekDayAdm$admWeekDay,sd)
  weekDayAdm.le<-tapply(weekDayAdm$n, weekDayAdm$admWeekDay,length)
  weekDayAdm.se<-weekDayAdm.sd/sqrt(weekDayAdm.le)
  weekDayAdm.cilb <- tapply(weekDayAdm$n, weekDayAdm$admWeekDay, function(v) t.test(v)$conf.int[1]) #IC lower bound
  weekDayAdm.ciub <- tapply(weekDayAdm$n, weekDayAdm$admWeekDay, function(v) t.test(v)$conf.int[2]) #IC upper bound

#Boxplot weekDay Admissions
  weekDayAdm.boxplot<-boxplot(weekDayAdm$n ~ weekDayAdm$admWday,main="Meyer ED admissions by weekday", xaxt = 'n' )
  axis(1,1:7,c("Sun", "Mon","Tue","Wed", "Thu", "Fri","Sat" ))

#Stripchart (SE is small because of many observations)
  weekDayAdm.stripchart<-stripchart(weekDayAdm$n ~ weekDayAdm$admWday ,vert = TRUE, pch = 1, method = "jitter", jit = 0.05, xlab =
"", ylab = "", xaxt = 'n',main="Meyer ED admissions by weekday" )
  axis(1,1:7,c("Sun", "Mon","Tue","Wed", "Thu", "Fri","Sat" ))
  points(1:24, weekDayAdm.M, pch = 16, cex = 1.5)
  arrows(1:24, weekDayAdm.M,1:24, weekDayAdm.M + weekDayAdm.se, lwd = 1.5, angle = 90, length = 0.1)
  arrows(1:24, weekDayAdm.M, 1:24, weekDayAdm.M - weekDayAdm.se, lwd = 1.5, angle = 90, length = 0.1)

#Boxplot Monthly Admissions
  monthlyAdm<- ddply(edData,c("admYear","admMonth"),summarise,n=length(idPat))
  arrange(monthlyAdm,monthlyAdm$admYear,monthlyAdm$admMonth)
  monthlyAdm.boxplot<-boxplot(monthlyAdm$n ~ monthlyAdm$admMonth, main="Meyer ED admissions by month")

###### Code for analysing of the time series with the forecast package

#Fuction to make a residuals histogram
plotForecastErrors <- function(forecasterrors)
{
  #Histogram of the forecast errors
  mybinsize <- IQR(forecasterrors)/4
  mysd <- sd(forecasterrors)
```

```
  mymin <- min(forecasterrors) - mysd*5
  mymax <- max(forecasterrors) + mysd*3
  #Normally distributed data with mean 0 and standard deviation mysd is generated
  mynorm <- rnorm(100000, mean=0, sd=mysd)
  mymin2 <- min(mynorm)
  mymax2 <- max(mynorm)
  if (mymin2 < mymin) { mymin <- mymin2 }
  if (mymax2 > mymax) { mymax <- mymax2 }
  #Red histogram of the forecast errors, with the normally distributed data overlaid
  mybins <- seq(mymin, mymax, mybinsize)
  hist(forecasterrors, col="red", freq=FALSE, breaks=mybins) #freq=FALSE ensures the area under the histogram = 1
  myhist <- hist(mynorm, plot=FALSE, breaks=mybins)
  #The normal curve as a blue line on top of the histogram of forecast errors is plotted
  points(myhist$mids, myhist$density, type="l", col="blue", lwd=2)
}

#Monthly Admissions are the series to analyse
  series.ts<-monthlyAdm.ts
  monthlyAdm.ts.decom<-decompose(monthlyAdm.ts)
  plot(monthlyAdm.ts.decom)

  lag.max=24 #Lags to observed period
  for.max=2  #Forecasting period

  acfpli<-acf(series.ts)
  acfpli<-acf(series.ts,lag.max=lag.max)
  acfpli$lag <- acfpli$lag * 12
  plot(acfpli, xlab="Lag (months)", main= paste("Correlogram of the series" ))
  axis(1, 0:lag.max)
```

Holt Winters's exponential smoothing (additive model with trend and seasonality)

```
HWForecastHist<-HoltWinters(series.ts)
HWForecastHist
plot(HWForecastHist)

#The forecast on historical data
HWForecastHist2<-forecast.HoltWinters(HWForecastHist, h=24)
plot.forecast(HWForecastHist2)

#ACF to verify if the residuals are correlated
acfpl<-acf(HWForecastHist2$residuals, lag.max=lag.max)
acfpl$lag <- acfpl$lag * 12
plot(acfpl, xlab="Lag (months)", main= paste("Correlogram of the residual error trend: ", method, "method",sep="" ))
axis(1, 0:lag.max)

#The significance of the correlations is calculated
Box.test(HWForecastHist2$residuals, lag=lag.max, type="Ljung-Box")

#It is verified if the residuals are normally distributed with zero mean
plot.ts(HWForecastHist2$residuals, main="Time plot of forecast errors") #Make a time plot
plotForecastErrors(HWForecastHist2$residuals) #Make an histogram
mean(HWForecastHist2$residuals) #Mean of the residuals, it is close to zero
```

ARIMA model

```
series.ts.diff.seasonal<-diff(series.ts,12)
series.ts.diff.seasonal.diff1<-diff(series.ts.diff.seasonal,differences=1)
plot.ts(series.ts)
plot.ts(series.ts.diff.seasonal,main="Meyer ED monthly admissions, seasonally adjusted")
series.ts.diff<-series.ts.diff.seasonal
```

```
#ACF and PACF to verify correlation
acfpli<-acf(series.ts.diff)
nacfpli<-acf(series.ts.diff,lag.max=lag.max)
acfpli$lag <- acfpli$lag * 12
plot(acfpli, xlab="Lag (months)", main= paste("Correlogram of the series" ))
axis(1, 0:lag.max)

pacfpli<-pacf(series.ts.diff)
pnacfpli<-pacf(series.ts.diff,lag.max=lag.max)
pacfpli$lag <- pacfpli$lag * 12
plot(pacfpli, xlab="Lag (months)", main= paste("Correlogram of the series" ))
axis(1, 0:lag.max)

#With the ACF and PACF the ARIMA model is (1,0,0)(2,1,0)12, it is verified with the auto.arima function
auto.arima(series.ts,stepwise=FALSE, approximation=FALSE)
fit <- Arima (series.ts, order=c(1,0,0), seasonal=c(2,1,0), include.drift=TRUE)

#Residuals of the model
Residuals <- residuals(fit)
tsdisplay(Residuals)

#The significance of the correlations between the residuals is calculated
Box.test(res, lag=24, fitdf=12, type="Ljung")

#The forecast on historical data
fitforecasts <- forecast.Arima(fit)
plot(fitforecasts, ylab="Number of visits", xlab="Year")
tsdiag(fit)
plot(forecast(fit,h=12,level=c(95)), plot.conf=TRUE, shaded=FALSE, shadebars=FALSE, pi.col=1, pi.lty=3, flty=4, type="o", ylab="Number
of visits", xlab="Year")
```

```
#Residuals of the forecast done with the model
fitforecasts$residuals
plot.ts(fitforecasts$residuals, main="Time plot of forecast errors") #Time plot of forecast errors

#It is verified if the residuals are normally distributed with zero mean
plotForecastErrors(fitforecasts$residuals) #Make an histogram
mean(fitforecasts$residuals) #Mean of the residuals, it is close to zero
```

ARIMA models by triage code (Same procedure but stratified by colours)

```
#RED
monthlyred<- ddply(edData[edData$admTriaCode=="ROSSO",],c("admYear","admMonth"),summarise,n=length(idPat))
monthlyred.ts<-ts(monthlyred$n,start=startdate,frequency=12)
monthlyred.ts
fitred <-auto.arima(monthlyred.ts,stepwise=FALSE, approximation=FALSE) #ARIMA(1,1,1)(2,0,0)12
resred <- residuals(fitred)
tsdisplay(resred)
Box.test(resred, lag=24, fitdf=12, type="Ljung")
fitforecastsred <- forecast.Arima(fitred)
plot(fitforecastsred, ylab="Number of visits", xlab="Month")
plot.ts(fitforecastsred$residuals)
plotForecastErrors(fitforecastsred$residuals)
mean(fitforecastsred$residuals)

#YELLOW
monthlyyellow<- ddply(edData[edData$admTriaCode=="GIALLO",],c("admYear","admMonth"),summarise,n=length(idPat))
monthlyyellow.ts<-ts(monthlyyellow$n,start=startdate,frequency=12)
monthlyyellow.ts
fityellow<- auto.arima(monthlyyellow.ts,stepwise=FALSE, approximation=FALSE) #ARIMA(0,1,1)(2,0,0)12
resyellow <- residuals(fityellow)
```

```
tsdisplay(resyellow)
Box.test(resyellow, lag=24, fitdf=12, type="Ljung")
fitforecastsyellow <- forecast.Arima(fityellow)
plot(fitforecastsyellow, ylab="Number of visits", xlab="Month")
plot.ts(fitforecastsyellow$residuals)
plotForecastErrors(fitforecastsyellow$residuals)
mean(fitforecastsyellow$residuals)

#GREEN
monthlygreen<- ddply(edData[edData$admTriaCode=="VERDE",],c("admYear","admMonth"),summarise,n=length(idPat))
monthlygreen.ts<-ts(monthlygreen$n,start=startdate,frequency=12)
monthlygreen.ts
fitgreen<-auto.arima(monthlygreen.ts,stepwise=FALSE, approximation=FALSE) #ARIMA(0,1,1)(2,1,0)12
resgreen <- residuals(fitgreen)
tsdisplay(resgreen)
Box.test(resgreen, lag=24, fitdf=12, type="Ljung")
fitforecastsgreen <- forecast.Arima(fitgreen)
plot(fitforecastsgreen, ylab="Number of visits", xlab="Month")
plot.ts(fitforecastsgreen$residuals)
plotForecastErrors(fitforecastsgreen$residuals)
mean(fitforecastsgreen$residuals)

#WHITE
monthlywhite<-ddply(edData[edData$admTriaCode=="BIANCO",],c("admYear","admMonth"),summarise,n=length(idPat))
monthlywhite.ts<-ts(monthlywhite$n,start=startdate,frequency=12)
monthlywhite.ts
fitwhite<-auto.arima(monthlywhite.ts,stepwise=FALSE, approximation=FALSE) #ARIMA (0,1,2)(1,0,0)12
reswhite <- residuals(fitwhite)
tsdisplay(reswhite)
Box.test(reswhite, lag=24, fitdf=12, type="Ljung")
fitforecastswhite <- forecast.Arima(fitwhite)
```

```
plot(fitforecastswhite, ylab="Number of visits", xlab="Month")
plot.ts(fitforecastswhite$residuals)
plotForecastErrors(fitforecastswhite$residuals)
mean(fitforecastswhite$residuals)
```

Verification of the Holt Winter's and ARIMA models

```
#Function that returns Root Mean Squared Error (RMSE)
rmse <- function(error)
{
  sqrt(mean(error^2))
}

#Function that returns Mean Squared Error (MSE)
mse <- function(error)
{
  mean(error^2)
}

#Function that returns Mean Absolute Error (MAE)
mae <- function(error)
{
  mean(abs(error))
}

#Function that returns Mean Absolute Percentage Error (MAPE)
mape <- function(y, yhat)
{
  mean(abs((y - yhat)/y))
}
```

```
#Parametres of the HW model
  mae(HWForecastHist2$residuals)
  mse(HWForecastHist2$residuals)
  mape(monthlyAdm.ts,HWForecastHist2$residuals)
  rmse(HWForecastHist2$residuals)

#Parametres of the ARIMA models
  mae(fitforecasts$residuals)
  mse(fitforecasts$residuals)
  mape(monthlyAdm.ts,fitforecasts$residuals)
  rmse(fitforecasts$residuals)

  rmse(fitforecastsred$residuals)
  rmse(fitforecastsyellow$residuals)
  rmse(fitforecastsgreen$residuals)
  rmse(fitforecastswhite$residuals)
```