NHS Bed Occupancy Forecasting

Group 7 - NHS

Mario Alvarez Junquera, Adam Morris, Eden Page, Kajol Vaitha and Max Williams

EMAT30005 Mathematical and Data Modelling

Department of Engineering Mathematics, University of Bristol

June 23, 2023

1 Introduction

Each year, the UK National Health Service (NHS) undergoes immense pressure to provide care for its population. However, bed occupancy rates in NHS hospitals continues to increase, especially due to the ongoing COVID-19 pandemic. As of November 2022, the average occupancy rate for beds across England reached 88.1%, a 1.8% increase since the beginning of 2020 [1]. This not only puts pressure on hospital staff and resources, but increases the risk of delayed treatment for patients. This may result in a reduction in the quality of care received due to a backlog in bed availability. High bed occupancy can result in premature discharge, which can lead to re-admissions in the future. Causes of these increases can largely be put down to a combination of too many patients and inefficient discharging processes.

The University Hospital Ayr of NHS Ayrshire and Arran in Scotland provides for approximately 100,000 people in the area and contains 300 beds. As of January 9th 2023, the hospital reached full capacity after opening all available beds, meaning only urgent day cases and paediatric actions are allowed to take place [2]. It shows the consequences of an absence of staff, delayed transfers of care and a high volume of patients requiring intensive and complex care. Between 2021 and 2022, there has been a large number of delayed discharges which has cost NHS Ayrshire and Arran over £12 million. In September 2022, 5,975 bed days, the number of days where beds are occupied by patients, were lost across NHS Ayrshire and Arran. This was due to the number of weekly commissioned hours, being the number of care hours provided by NHS Scotland, falling by 45% since April 2021 [3]. Due to this, patients who are medically fit to leave the hospital are not able to do so.

The aim of this report is to produce a model which allows the forecasting of bed occupancy within the University Hospital Ayr of Ayrshire and Arran. The model generates a three-month forecast, which allows analysis to be completed on the scenarios under which the hospital is likely to see increases in admissions and a decrease in discharges. Ideally, bed occupancy should be kept at 85% or below. Through the evaluation of factors which affect the length of stay (LOS) of a patient, the model simulates the admission and discharge of generated patients. The model parameters are adapted to simulate varying situations and their effect on the overall bed occupancy.

2 Agent-Based Model

The chosen method to simulate the admission and discharge of patients is an agent-based model approach. Individual, autonomous agents are generated with an array of attributes that determines how they interact within the created environment. In this case, the agents are the patients, who interact with the hospital. Patients are assigned a LOS upon generation, based on the factors of condition, gender, age and discharge probability.

The provided data from NHS Ayrshire and Arran consists of the average LOS for male and female patients with eight categories of conditions in 2021; myocardial infarction, stroke, COVID-19, fracture of the femur, atrial fibrillation, pneumonia, COPD and endocrine, nutritional and metabolic diseases [4]. In addition, the NHS publish annual admissions data which contains the number of admissions, gender, mean wait time, age and the number of bed days for each International Classification of Diseases (ICD) category [5]. This is a much larger and more comprehensive dataset about each condition under investigation. Using the large NHS dataset for 2021-2022, the attributes for the provided categories of conditions are extracted. This is used to conduct further studies into how the factors will affect the LOS of the admitted patients. The results of the analysis are used to inform the model on how to generate incoming patients representatively.

2.1 Condition

By extracting the relevant data from the large NHS dataset, assuming that patients only have conditions from the predefined categories, the overall distribution of admissions can be seen in Figure 1. The largest proportion of incoming patients had endocrine, nutritional and metabolic diseases and pneumonia at 23.5% and 20% respectively. This is partially due to this category encompassing many conditions, including Type 2 Diabetes, which are among the most common conditions for incoming patients [6]. Femur fracture is the most uncommon condition, which has the lowest percentage of admissions at 6.2%. The proportion of COVID-19 patients is relatively high, due to the data being captured in 2021-2022, where there were over 17.5 million cases, and over 210,000 admitted patients across the UK [7, 8].

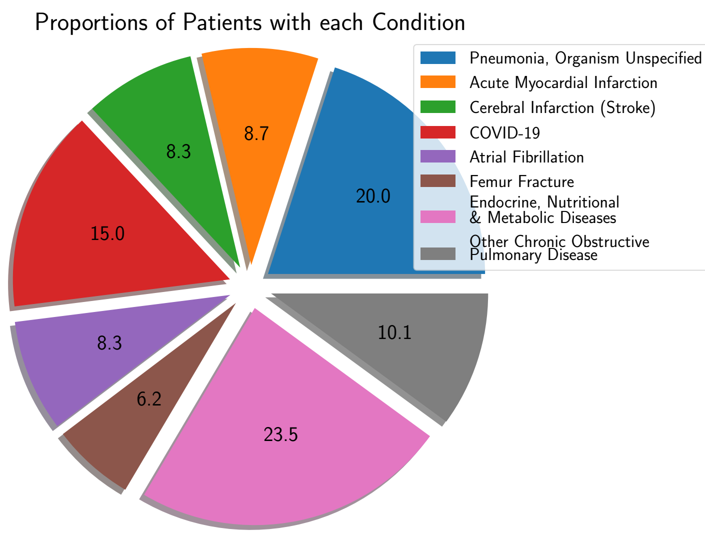

Figure 1: Proportion (%) of NHS patients with each given category of conditions in 2021-2022.

2.2 Gender

Previous studies have indicated that gender has a significant impact on patient LOS [9]. Depending on the condition, it is more likely for one gender to not only have the condition but also to have a longer average LOS [10]. An examination of the distribution of male and female patients allows patient generation to be more realistic within the model. Figure 2a shows that most of the conditions have an equal split of male and female patients, however, there is some variation with the type of condition. The significant differences are that females are more likely to have a femur fracture than males, and males are more likely to have myocardial infarction than females [4].

The provided data gives an insight into how the average LOS varies between male and female patients for each condition. Figure 2b verifies that gender influences the average LOS, as the variation is seen to depend on the condition. For example, females with a femur fracture stay in hospital for around four more days than males. However, male patients with pneumonia and stroke have a higher average LOS by two days. Therefore, it cannot be assumed that the length of stay is only influenced by the type of condition, so the factor of gender is taken into account during patient generation.

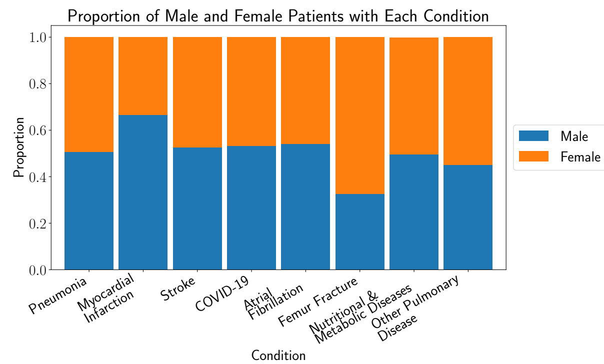

(a) Admitted male and female NHS patients as a proportion of total admissions for each condition in 2021-2022 [5].

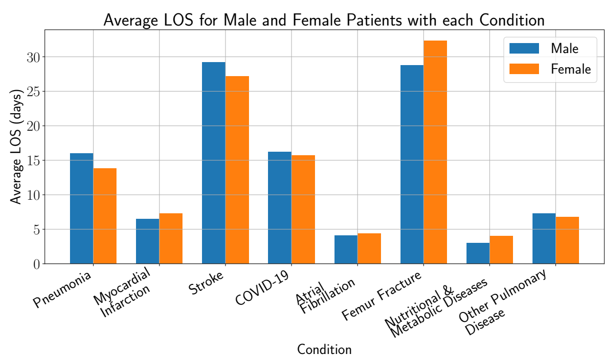

(b) Average LOS for male and female patients with each condition, extracted from the provided data [4].

Figure 2: Graphs representing the effect of gender on patient length of stay.

2.3 Age

The age of patients is intuitively a factor with significant variation and impact on how long a patient takes to recover in hospital. The number of patients from each ten-year age band for the selected conditions in the large NHS dataset is extracted, illustrated in Figure 3. It shows that 70-90 year-old patients account for nearly 50% of patients with each condition, with COVID-19 and nutritional and metabolic diseases having the most variation. Femur fracture and pulmonary disease have a significantly high proportion of older patients, as bones typically become more brittle with age. For younger patients, injuries such as fractures and breaks often heal much quicker and naturally than older patients, who may need surgery and longer physiotherapy. An important observation is that although femur fracture accounts for the smallest proportion of admissions, seen in Section 2.1, it has the highest average LOS of the eight conditions, at around 30 days.

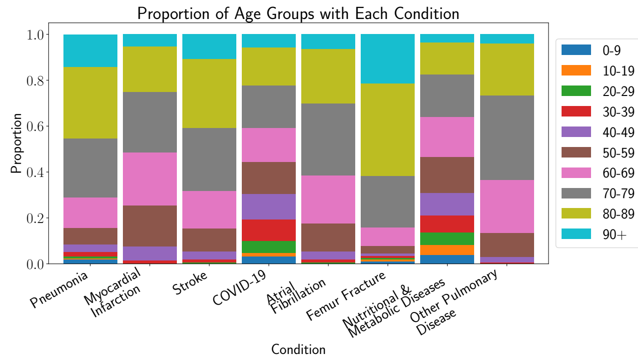

Figure 3: Age distribution of admitted patients across the eight categories of conditions.

Age Group	**Average LOS (days)**
0 to 9	0.9
10 to 19	7.5
20 to 29	8.7
30 to 39	9.5
40 to 49	10.6
50 to 59	11.6
60 to 69	14.5
70 to 79	14.6
80 to 89	16.3
90+	15

Table 1: Average LOS for each 10 year age group.

In order to accurately inform the model of how to generate patient ages, data from the NHS is obtained which shows the total number of admissions per age group and gender for the UK during the Full Course Equivalent of 2019-2020 [11]. From this, the probability of an agent having a certain age and condition can be estimated. To accurately define these probabilities, additional data is obtained, which contains the average LOS for each age group and condition [12]. Table 1 summarises this information.

Each agent is assigned an age group using the probabilities in Figure 3. An age factor is found by taking the mean of the average LOS and calculating the number of standard deviations the average LOS is away from the mean. The LOS of the agent is increased or decreased by 10% for each standard deviation above or below the mean respectively.

2.4 Discharge Probability

It is important that patients are discharged at rates that reflect true hospital admissions. To achieve this, rather than setting a proposed average discharge rate, a general trend for patient LOS should be determined. Figure 4 shows the provided data, showing the patient count at a given LOS within the University Hospital Ayr.

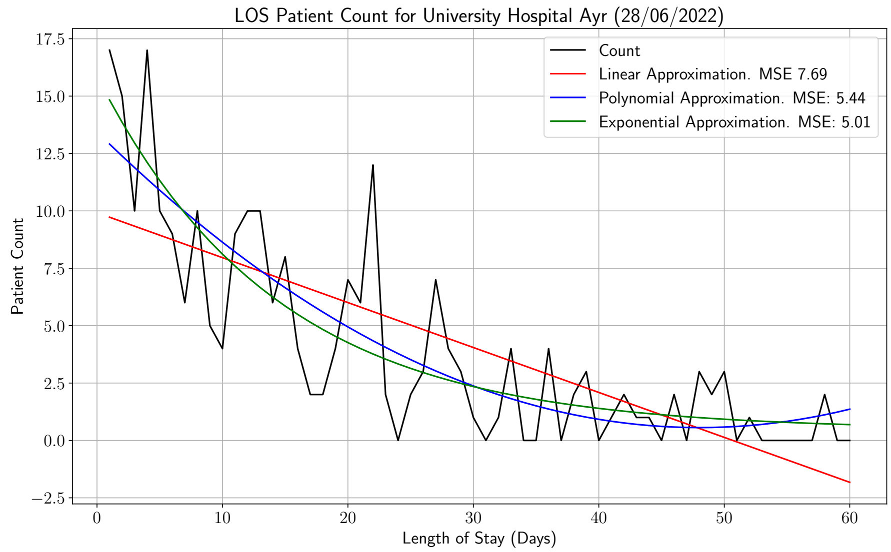

Figure 4: Provided LOS data for one day in June 2022 along with approximations used to estimate typical LOS trend.

Initial observations suggest that LOS distributions follow an exponential decay with increasing days. Three methods of curve fitting have been employed to further investigate this trend; linear, second-order polynomial and exponential approximations. Both the linear and polynomial approximations produce a higher mean squared error than the hypothesised exponential. Furthermore, as the number of days becomes sufficiently large, the approximations deviate from reality. An exponential approximation best describes the patient LOS distributions, as it avoids negative LOS and tends to zero.

The probability of a patient being discharged on a given day, x, is estimated using the cumulative distribution function of an exponential, given in Equation (1). λ represents the rate parameter, being the rate of increase of probability that the patient is discharged. The value of λ is calculated using the patients' generated average LOS, with smaller values increasing the rate parameter.

$$F(x;\lambda) = \begin{cases} 1 - e^{-\lambda x} & x \geq 0, \\ 0 & x < 0. \end{cases} \qquad \lambda = \frac{1}{\text{Average LOS}} > 0 \tag{1}$$

2.5 Model Parameters

The hospital is fixed to 300 beds, as this cannot be practically changed within the real hospital. Additionally, staff and resources are not modelled individually, but are assumed to have manually prescribed effects on patient LOS. For example, reduced staff availability is reflected in the model as an increase in LOS for all patients. If the bed occupancy is at 100% capacity, any new admissions are not admitted, nor added to a waiting queue.

Patients are assigned one of the eight categories of aforementioned conditions, and an age value based on national averages, which the hospital catchment area is assumed to follow. This age value is categorised into one of ten bands, with all ages within a single band behaving identically.

Unless otherwise stated, the admission rate is constant and set to 55 patients per day, based on the provided data. Discharges include both patients who have been successfully treated and patient deaths. For non-elective admissions in England between 2021 to 2022, the number of patient deaths as a proportion of discharges was 3.3% [13]. Due to this small proportion, deaths are included within the total discharges in the model.

The model begins by generating 55 new patients on Day 1. Each individual agent has a discrete probability distribution used to assign one of the eight possible conditions. Once a condition has been assigned, patients are given a gender based on the data from Section 2.2, again through a discrete random variable. Using both of these attributes, an initial average LOS is assigned, following the data in Figure 2b. Data from Section 2.3 is then used to assign the age bracket for a patient, based on their given condition. This is used to calculate the final average LOS of a patient, by multiplying the initial LOS value with the age factor.

All admitted patients are added to the hospital patient list, where they are treated in order. Using Equation (1), patient discharge is modelled in conjunction with a Bernoulli random variable

$$P(X_i) = \begin{cases} p_i & 1 \\ 1 - p_i & 0, \end{cases} \tag{2}$$

where $p_i = F(x_i; \lambda_i)$ and x_i is the current LOS of patient i. A value of 1 means the patient is discharged and 0 means the patient is not discharged.

This continues until all patients have been treated. The program outputs the number of admitted and discharged patients, as well as the percentage of bed occupancy for that day. The day is then incremented and the process repeats for 120 days. Although the targeted forecasting period is three months, an additional 30 days of simulation are added to initialise the hospital, allowing a steady state to be reached before recording results. See Appendix A for flowcharts of the processes involved per day of the simulation.

2.6 Verification using Prophet

A forecasting software, Prophet, is used to verify the model. Prophet is an additive regression model that uses a combination of linear and logistic growth trends to make predictions, incorporating both yearly and weekly seasonal patterns. These are modelled using Fourier series and dummy variables [14]. Its use is recommended by the NHS, amongst other advanced forecasting methods [15]. This software forecasts time series data, where nonlinear trends are fit with yearly, weekly and daily changes [16]. This is a computationally inexpensive tool which allows for comparison between the output from the simulation and the predicted results from the forecast.

This software requires at least one year of daily occupancy data. Due to the unavailability of viable Scottish data, a dataset from the national statistics repository for Wales is used [17]. This contains the total number of occupied and available beds for the past three years. Using this, Prophet produces a prediction for the percentage occupancy in 2023. Given that this data is not sourced from the studied area, overall and daily predictions are not suitable for performing direct comparisons. However, Prophet outputs yearly and weekly trends which can be used to verify that the trends generated by the model are realistic through statistical testing. Prophet is therefore treated as a ground truth to compare to the model.

3 Results

Figure 5a shows the initial results produced by the model, as well as its average over the course of ten individually run simulations. This base case shows the results when the parameters are set using the provided data and the probability distributions defined in Section 2. A single iteration contains large fluctuations on a day-to-day basis due to the inherent randomness of the system. However, averaging across multiple runs reduces this variance, allowing insight into trends that may be expected.

In this scenario, bed occupancy tends towards roughly 80% capacity. The model gives an average discharge rate of 52.5 patients per day, which falls close to the true recorded average of 56 patients per day. This result suggests that the simulated hospital is not in fact meeting its demand, due to the set admission rate of 55 patients per day. Although the average occupancy falls below the targeted 85%, this result only accounts for those admitted in an emergency with the given conditions. This leaves less than 5%, or 15 beds, for any other admissions, including those with planned appointments. It is within interest to further investigate scenarios and their given results from the model, with the goal of managing occupancy. This allows conclusions to be drawn as to when the hospital is expected to see its busiest periods. From this, suitable strategies can be suggested, such as the reallocation of required staff and resources.

3.1 Varying Length of Stay

The key contributing factor towards bed occupancy is the distribution of admitted lengths of stay, as previously discussed in Section 2.1. Applying an increase in all patient LOS simulates the scenario of staff and resource shortages. This is because the result of reduced staff would mean the staff-to-patient ratio increases, hence the time taken for a patient to be treated increases.

To model staff and resource shortages, all patient LOS are modified gradually from 50% to 150% of their original distributions, corresponding to the LOS factor seen in Figure 5b. The most severe changes in discharge rates occur between roughly 80-130%. As a result, small increases in LOS can have detrimental effects on both patients and the hospital, indicating the importance and effect of the strain on hospital staff. If staff shortages occur, this means patients are in hospital for longer than required, as the hospital simply cannot attend to patients as efficiently. However, if the hospital can reduce their patient LOS by 15%, bed occupancy falls to nearly 76%. This indicates that when applying possible strategies to improve resource allocation and planning, small improvements can have great benefits.

Small but significant changes in the strategy of the hospital can make the patient discharge process more efficient. For example, the NHS Outpatient Parenteral Antimicrobial Therapy (OPAT) service allows patients to be discharged early and receive antibiotic treatment at home [18]. Similarly, the Hospital at Home scheme lets patients be visited by nurses to receive care in their homes [19]. This would reduce admissions demand and ensures that hospital beds are occupied only by patients who urgently need them. In addition, the Discharge without Delay (DwD) scheme could improve the discharge rate of the hospital [20]. This plan involves optimising the discharge process by implementing minor strategies within the hospital. For example, having a default discharge time before 12pm to avoid delays, and holding staff meetings each morning and afternoon to plan the order of ward rounds that would optimise discharge.

3.2 COVID-19 Pandemic

Hospitals may experience a surge in a particular type of condition and must be able to cope with this using the available staff and resources. A recent event which illustrates this is the COVID-19 pandemic. To observe how the model reacts in this situation, the probabilities of the agent having each condition are updated. The patients have a 75% chance of having COVID-19 which resembles COVID-19 patients in England between June and December 2021. The other conditions are then given an even division of the remaining 25% [21]. Figure 5c shows the percentage bed occupancy and average occupancy forecast for a single iteration of the model and the average after ten iterations. The single iteration shows the fluctuation of the occupancy between 79% and 88% with an average occupancy of just under 84% over the three months. After more iterations, this turns out to average between 82% and 85%.

This means that overall, the occupancy is barely below 85%, indicating that this scenario would result in a relevant strain for the hospital. In addition, all other admissions are reduced to around 3.5% per condition, which means very few patients would receive treatment for other emergency conditions. Subsequently, they would need to be sent away to other hospitals or refused treatment until a later date. For example, stroke patients can be referred to specialist centres and COPD patients can use community-based initiatives [22]. Once the effects of the pandemic dissipate, the backlog formed by those patients would not represent as much of a problem for the hospital.

3.3 Dynamic Admission Rate

The chosen method to evaluate the created model is to compare weekly trends to those seen in actual data. Daily admissions and discharges are difficult to compare directly due to the randomness of these processes. It has been documented that admission rates follow weekly patterns, with clear decreases on the weekends [23]. The model has therefore been modified to include varying admission rates which follow this trend. The average admission rate has not been altered, only how it is distributed over a typical week.

Figure 5d shows how this affects occupancy rates in the model. Average occupancy rates follow a similar level to the initial model results, however, there is seen to be large increases mid-week, followed by sharp decreases towards the end of the week.

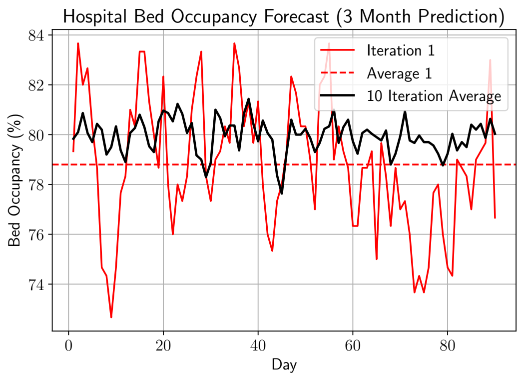

(a)

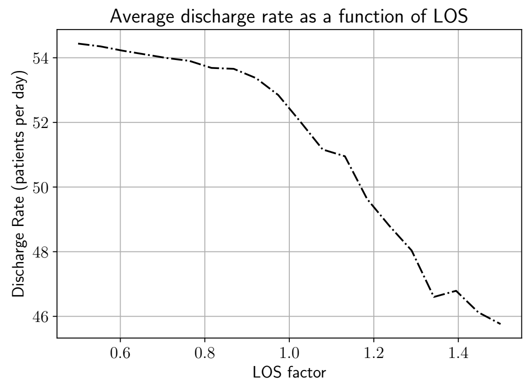

(b)

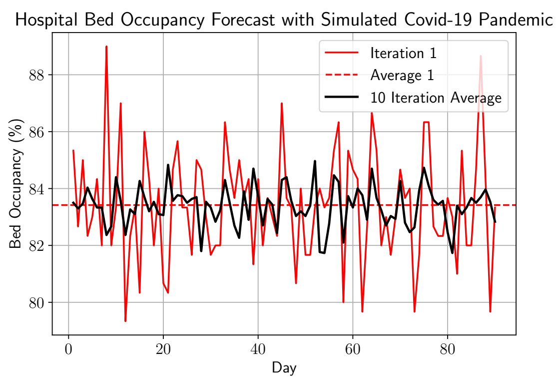

(c)

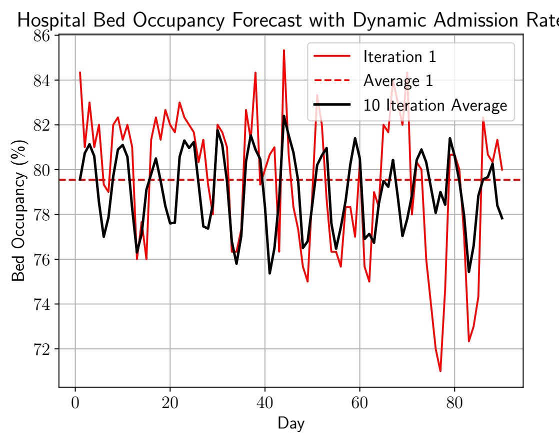

(d)

Figure 5: Output results from the agent-based model. (a) Initial model result showing three month future prediction of bed occupancy. The first iteration and its overall average, as well as daily 10 iteration average plotted. Average discharge rate = 52.5 patients per day. (b) Average discharge rate against increasing LOS distribution, where 1.0 factor indicates no change, 1.5 indicates 150% of original distribution. (c) Model prediction with COVID-19 pandemic scenario. Average discharge rate after ten iterations = 46.8 patients per day. (d) Model prediction with weekly trend admission rates.

3.4 Comparison to Prophet

The Prophet algorithm allows for the extraction of weekly trends, analysed from the Welsh national statistics repository data. The output, shown in Figure 6, provides this factor for each day.

The trend indicates decreases in bed occupancy on weekends as opposed to weekdays, similar to that outlined in Section 3.3. To evaluate the agent-based model, three months of future Prophet prediction are extracted to compare to the results of Section 3.3. To measure potential similarities of the two time-series, three different techniques are considered; dynamic time warping (DTW), the Pearson correlation coefficient (PCC) and Spearman's rank correlation coefficient (SRCC). DTW is used to create an optimal match between two time-series which vary in speed. The result

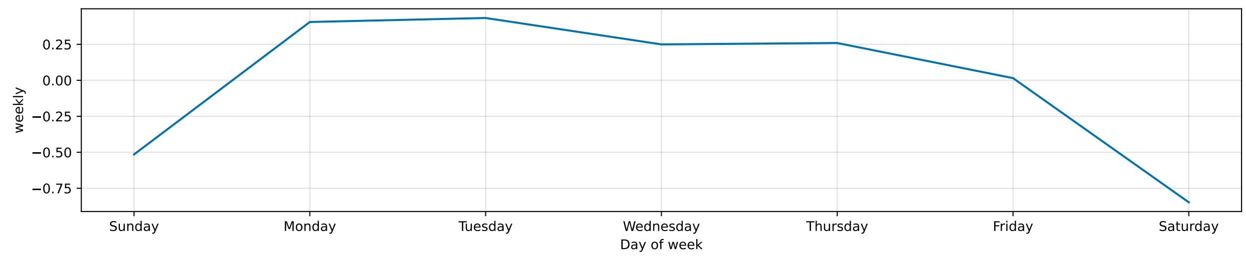

Figure 6: Weekly trend output from the Welsh data using Prophet, showing how the number of beds occupied varies depending on the day of the week.

from DTW represents the sum of the Euclidean distances between the warped time-series. The advantage of this method is that it is robust when measuring relations that have time-shifting. The results can therefore be compared to Prophet without aligning the days of the week. PCC is a measurement of covariance to find linear correlations between two variables. SRCC measures the statistical dependence between the rankings of two variables to find monotonic relationships. Both coefficients are measured between -1 and 1, where 1 is a perfect correlation. A value is considered as statistically significant if the calculated p-value is less than 0.05.

Test Statistic	**Base Case**	**Weekly Trends**
DTW	0.1294	0.1778
PCC	0.1344	0.3510
SRCC	0.2480	0.3750

Table 2: Results of evaluation techniques comparing the base case of the agent-based model to Prophet weekly trends.

Table 2 shows that PCC and SRCC produce similar, statistically significant values for the weekly trend data. These values indicate a small level of linear and monotonic relation between the two time-series. This shows that bed occupancy in either model tends to increase or decrease fairly similarly, and that the relationships are linear. The DTW values show the weekly trend distance is greater than the base case. This is due to the larger daily variation of bed occupancy seen as an impact of weekly trends. Appendix B shows how the distance measure changes over the course of the time series, for all possible paths of the weekly trend results. It is seen that as both time series progress, this value increases.

4 Discussion

The model is based on the University Hospital Ayr of NHS Ayrshire and Arran, meaning some of the parameters are already known. According to the information provided, the hospital has 300 available beds and the average admission rate is 55 patients per day. Both of these parameters are set to be constant, with the exception of the dynamic admission rate scenario, in which the admission rate varies depending on the day of the week. These fixed parameters represent the restrictions in the strategy of the hospital which cannot be changed.

In the model, staff and resources are not modelled individually, but are represented through an overall change in the LOS. Whenever the overall LOS is increased by a certain percentage, it is representing a reduction of staff availability and accessible resources and vice versa. However, the model does not consider the potential effects this has on the admission rate, which is assumed to remain constant. An improved model would incorporate an element of monitoring staff numbers and adapting the admission rate accordingly. This would make the model more realistic as reduced staff means that fewer patients are able to be admitted per day due to a slower admissions process.

The provided data from NHS Ayrshire and Arran gave the LOS information for eight categories of condition, all within the emergency department. Realistically, there are various conditions outside of these categories with which patients can be admitted. However, the model only considers these particular conditions for the entire hospital, assuming that all patients are emergency cases. This causes a significant difference in the average LOS because the chosen conditions are obtained predominantly by people in older age groups who use more bed days than younger age groups [24]. Despite this, these eight categories encompass the majority of causes of patient admission, hence it is feasible that the model represents the admissions process fairly similarly to reality. An improved model would therefore consider the full range of conditions supplied in the large NHS dataset, but this would make the model computationally expensive.

The results show that the bed occupancy fluctuates around 85%, which is the desired statistic for this paper. However, occupancy is affected by scheduled procedures. Even though the LOS for scheduled procedures has less variance and is smaller than emergency cases, there is a probability that all beds get occupied. The regular procedure of a hospital would be to have a waiting list, which gives priority to more urgent cases. In the model, when 100% of the beds get occupied, new admissions get turned away which is unrealistic. Currently, this is representing a case where all overflowing patients either have their treatment delayed or are sent to another facility. For example, stroke patients can be sent to a specialist centre instead of going to the general hospital. Hence, implementing a queuing system would improve the accuracy of the model, as it would differentiate the patients who have delayed care and those who are being treated at another facility. This system may be achieved by creating a second patient list of all those who are not admitted due to the maximum occupancy. Additionally, treatments could be processed hourly, instead of daily, to more accurately model the flow of patients. Queued patients may then be admitted depending on a specified priority.

One of the most influential factors on LOS in the model is patient age. For ease of data processing between different datasets, the patients are grouped into age bands of ten years. A distribution is found based on the probabilities calculated from the resulting data. While there are a total of ten age groups, there are still some significant assumptions about the behaviour of patients within each band. For example, patients assigned an age in the 0 to 9 group would all be assigned an LOS based on the same probabilities from the data. This is assuming that a newborn has the same recovery rate as a 9-year-old, which in reality would not be the case. Studies have shown that younger children have faster recovery times for conditions such as COVID-19 and bone fractures, meaning LOS generally increases with age [25, 26]. This indicates that there are LOS distributions within each of the age bands which are not considered. To improve the accuracy of patient age distribution and corresponding behaviour, smaller age bands should be considered, or further analysis of the NHS data with continuous age variability.

The model is based on data from a single year of observations. Both the provided data from NHS Ayrshire and Arran and the NHS UK data capture admissions between 2021 and 2022. The model assumes the distributions of each factor follow those given by these datasets. To improve the ability of the model to represent the population of Ayrshire, an analysis of the local residential area could be performed. The age, gender and admissions distributions would be adaptive, changing accordingly to population dynamics each year. This would require additional data to be collected for the surrounding area.

The discharge rate fluctuates throughout the forecast due to the factors implemented in the model. However, there are other contributing factors and patient attributes which are not included. Patients can be sent to other hospitals which might have better resources for their type of condition. This patient would then be registered as discharged from that hospital, reducing the amount of waiting patients. Patients that are medically fit to leave but still require attention could be directed to a home care scheme. This would help reduce bed occupancy but increases the demand for carers. The COVID-19 results have a low discharge rate, due to most resources being focused on treating a single condition. This impacts the flow of patients for other conditions as their admission rate is reduced [27]. Dynamic admission rates allow us to study the changes within a day-to-day system. It is evident that admission rates fall on the weekends, which results in an improved ratio of staff to patients.

The model is extended through the incorporation of dynamic admission rates. The results of this improved the correlation of the model to that of the Prophet algorithm, shown by an increase in the PCC and SRCC values. However, the DTW distance value increases when compared to the base case. This suggests the magnitude of weekday to weekend differences may have been set too large for the size of the hospital. It can be determined that the short-term trends between the two datasets have a higher correlation, becoming increasingly dissimilar over larger time periods. The result of this is due to Prophet's extraction of seasonal trends, in addition to weekly. The predictions from Prophet will therefore tend to increase or decrease due to the time of year, which is not taken into account in the agent-based model. An improved model would incorporate these changes to better represent long-term behaviour.

5 Conclusion

In conclusion, the agent-based model presents a simulation of the hospital with the output of a bed occupancy forecast. The model generates patients using NHS admissions data from 2021 - 2022 and simulates eight categories of emergency conditions. Four scenarios are simulated: the initial case, varying the LOS of all patients, a COVID-19 pandemic and incorporating dynamic admission rates. The initial model results show that the overall occupancy can be managed at around 80%, if the hospital solely admits emergency cases. Small changes in average LOS have a large impact on the discharge rate. The hospital can implement schemes such as DwD and OPAT, which target early discharge and provide care to patients in their homes. The COVID-19 simulation showed that the strain could be improved by referring patients to other hospitals for treatment. The model is improved by incorporating a dynamic admission rate to account for weekly trends in patient admissions. Statistical analysis showed that the model follows the expected trends closely in the short term, but moves away from this in the long term. The University Hospital Ayr can use the outcomes of this work to start evaluating and implementing new strategies to alleviate the current occupancy pressures they are facing.

References

[1] NHS UK. Bed Availability and Occupancy. https://www.england.nhs.uk/statistics/statistical-work-areas/bed-availability-and-occupancy/, 2022.

[2] Emma Reilly. NHS Ayrshire and Arran: Hospitals reach full capacity as crisis mounts. https://www.irvinetimes.com/news/23243003.nhs-ayrshire-arran-hospitals-reach-full-capacity-crisis-mounts/, 2023.

[3] Kevin Dyson. NHS Ayrshire's delayed discharge crisis driven by South Ayrshire's reliance on private care. https://www.dailyrecord.co.uk/ayrshire/nhs-ayrshires-delayed-discharge-crisis-29010026.

[4] Public Health Scotland. Acute hospital activity and NHS beds information (annual). https://rb.gy/ewvdiq, 2021.

[5] National Statistics. Hospital Admitted Patient Care Activity, 2021-22. https://digital.nhs.uk/data-and-information/publications/statistical/hospital-admitted-patient-care-activity/2021-22, 2022.

[6] Diabetes UK. Us, Diabetes and a lot of facts and stats. https://www.diabetes.org.uk/resources-s3/2019-11/facts-stats-update-oct-2019.pdf, 2019.

[7] GOV.UK. Cases in the United Kingdom. https://coronavirus.data.gov.uk/details/cases, 2023.

[8] NHS Digital. Weekly Admissions and Beds from 1 October 2021 up to 31 March 2022. https://www.england.nhs.uk/statistics/statistical-work-areas/covid-19-hospital-activity/, 2022.

[9] Tayebeh Baniasadi, Marjan Ghazisaeedi, Mehdi Hassaniazad, Sharareh Rostam Niakan Kalhori, and Mehraban Shahi. A study of factors affecting on patients' length of stay in a surgical ward: Toward optimization by technology-based solutions. 12 2019.

[10] David Westergaard, Pope Moseley, Freja Karuna Hemmingsen Sørup, Pierre Baldi, and Søren Brunak. Population-wide analysis of differences in disease progression patterns in men and women. *Nature communications*, 10(1):1–14, 2019.

[11] NHS UK. Hospital Admitted Patient Care Activity 2019-20. https://digital.nhs.uk/data-and-information/publications/statistical/hospital-admitted-patient-care-activity/2019-20/summary-reports---apc---patient, 2020.

[12] NHS. Mean and median length of stay of hospital inpatient episodes. https://digital.nhs.uk/data-and-information/find-data-and-publications/supplementary-information/2019-supplementary-information-files/mean-and-median-length-of-stay-of-hospital-inpatient-episodes, 2019.

[13] National Statistics. Summary Hospital-level Mortality Indicator (SHMI) - Deaths associated with hospitalisation, England, February 2021 - January 2022. https://digital.nhs.uk/data-and-information/publications/statistical/shmi/2022-06, 2022.

[14] Amazon Forecast: Developer Guide. Prophet algorithm. https://docs.aws.amazon.com/forecast/latest/dg/aws-forecast-recipe-prophet.html, 2023.

[15] NHS England and NHS Improvement. Advanced forecasting techniques. *NHS. London: NHS England*, 2020.

[16] Facebook. Prophet. https://facebook.github.io/prophet/, 2017.

[17] StatsWales. NHS beds by date and use. https://statswales.gov.wales/Catalogue/Health-and-Social-Care/NHS-Hospital-Activity/nhs-activity-and-capacity-during-the-coronavirus-pandemic/nhsbed-by-date-use, 2022.

[18] A. L. N. Chapman. Outpatient parenteral antimicrobial therapy. *BMJ*, 346(mar26 1):f1585–f1585, March 2013.

[19] NHS. NHS @home. https://www.england.nhs.uk/nhs-at-home/, 2023.

[20] Scottish Government. Discharge without delay. https://www.nhsgoldenjubilee.co.uk/application/files/6216/3489/4272/Discharge_without_Delay_Discussion_Paper_Final.pdf, 2021.

[21] ONS. Coronavirus (COVID-19) latest insights: Hospitals. https://www.ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/conditionsanddiseases/articles/coronaviruscovid19latestinsights/hospitals, 2023.

[22] NHS. Community copd service (chronic obstructive pulmonary disease). https://www.uhcw.nhs.uk/download/clientfiles/files/Patient%20Information%20Leaflets/Medicine/Respiratory/118510_Community_COPD_Service_(1547)_Jan19.pdf, 2022.

[23] William N. Evans and Beomsoo Kim. Patient outcomes when hospitals experience a surge in admissions. *Journal of Health Economics*, 25(2):365–388, March 2006.

[24] G George. Effect of population ageing on emergency department speed and efficiency: a historical perspective from a district general hospital in the UK. *Emergency Medicine Journal*, 23(5):379–383, May 2006.

[25] Jonas Ludvigsson. Systematic review of COVID-19 in children show milder cases and a better prognosis than adults. *Acta Paediatrica*, 109, 03 2020.

[26] Steven Frick. Skeletal Growth, Development, and Healing as Related to Pediatric Trauma. 12 2015.

[27] Syed Ahmar Shah. Impact of first UK COVID-19 lockdown on hospital admissions: Interrupted time series study of 32 million people. *eClinicalMedicine*, 49:101462, July 2022.

A Model Flowcharts

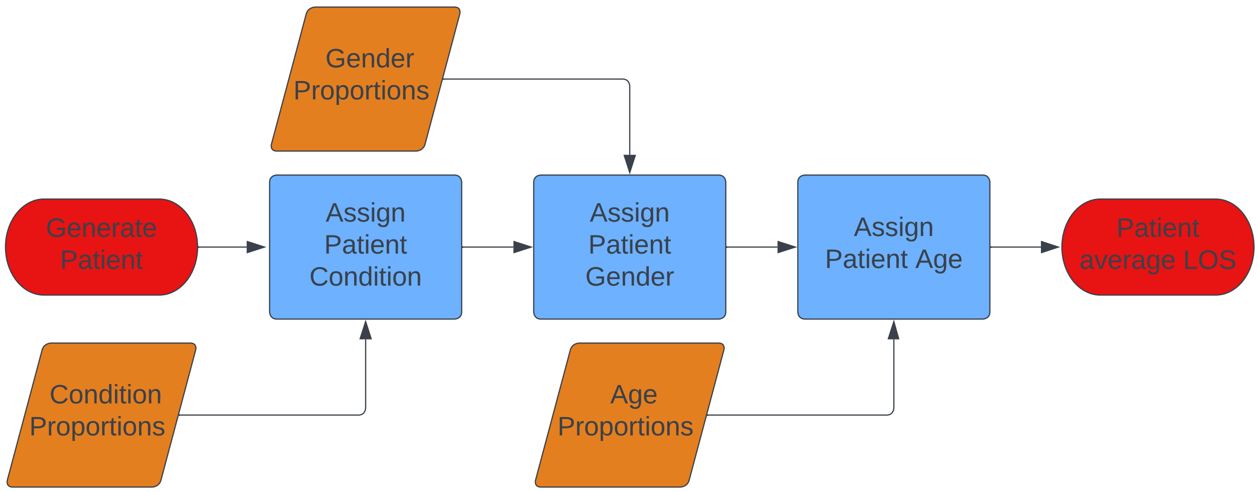

Figure 7: Flowchart showing patient generation within the model. The process applies to every patient before they are processed by the hospital. This will repeat n times, where n = admission rate.

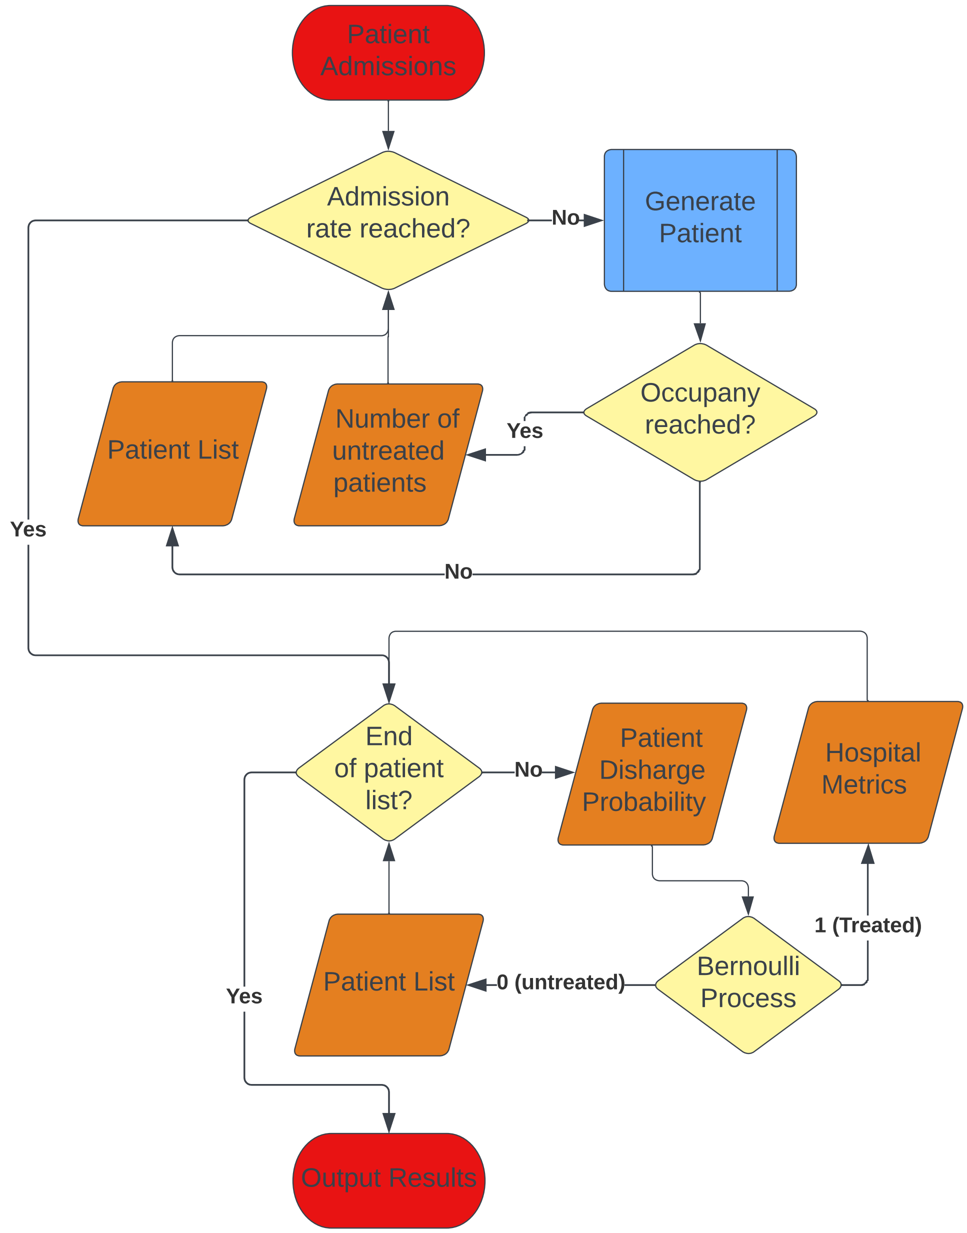

Figure 8: Flowchart showing process for one day of the simulation. Generate Patient is a pre-defined process outlined in Figure 7. The entire process repeats every day, until the simulation ends after m days, equal to the desired forecasting period.

B Dynamic Time Warping Cost Matrix

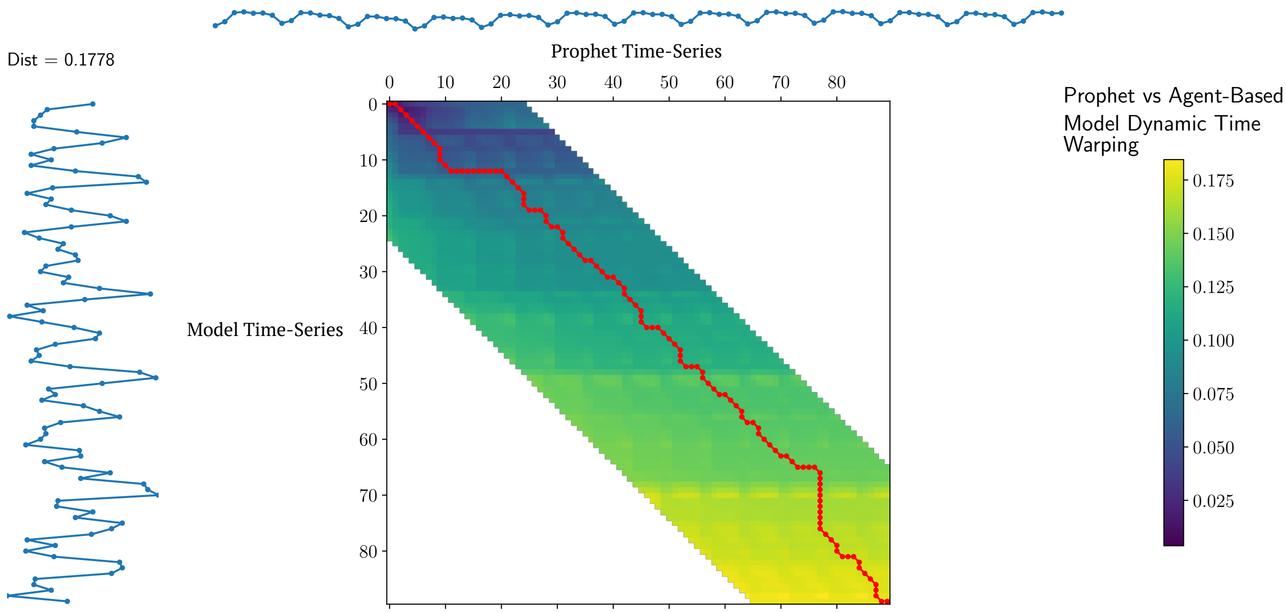

Figure 9: Visualisation of the cost matrix for all dynamic time warping paths, with optimal path shown in red for weekly trend results.