

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository: <https://orca.cardiff.ac.uk/id/eprint/33198/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Gillard, Jonathan William and Knight, Vincent Anthony 2014. Using singular spectrum analysis to obtain staffing level requirements in emergency units. *Journal of the Operational Research Society* 65 (5) , pp. 735-746. 10.1057/jors.2013.41

Publishers page: <http://dx.doi.org/10.1057/jors.2013.41>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies. See <http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.



Using Singular Spectrum Analysis to Obtain Staffing Level Requirements in Emergency Units

Jonathan Gillard and Vincent Knight
Cardiff School of Mathematics
Cardiff University
{GillardJW, KnightVA}@Cardiff.ac.uk

Abstract

Many Operational Research (OR) techniques use historical data to populate model input parameters. Although the majority of these models take into account stochastic variation of the inputs, they do not necessarily take into account seasonal variations and other stochastic effects that might arise. One of the major applications of OR lies within healthcare, where ever increasing pressure on healthcare systems is having major implications on those who plan the provision of such services. Coping with growing demand for healthcare, as well as the volatile nature of the number of arrivals at a healthcare facility makes modelling healthcare provision one of the most challenging fields of OR. This paper proposes the use of a relatively modern time series technique, Singular Spectrum Analysis (SSA), to improve existing algorithms that give required staffing levels. The methodology is demonstrated using data from a large teaching hospital's emergency unit. Using time dependent queueing theory, as well as SSA, staffing levels are obtained. The performance of our technique is analysed using a weighted mean square error measure, introduced in this paper.

1 Introduction

Many traditional OR methods have been used to ensure the efficient running of healthcare systems. Examples include: the optimisation of surgery schedules [9], the optimal location of healthcare clinics [29], and the modelling of capacity requirements in a critical care directorate [20]; amongst many others. Unscheduled care is one of many aspects that a hospital has little control over. For example, much media attention has focused on Emergency Units (EU), often reporting long waits for arriving patients until they are seen by a healthcare professional (see for example [4]). Consequently, EU's have been the focus of much research [2, 6, 10, 28].

Time dependent queueing theory has been used to schedule the number of clinical decision makers needed in an EU to allow only for a small probability of a patient waiting longer than a specified duration [19]. Many of the methods for modelling EU's require a robust prediction of the amount of demand likely to be exerted upon the EU. Classical approaches use historical data to obtain distributions of the interarrival time of patients. This paper aims to showcase the potential of a relatively novel technique known as Singular Spectrum Analysis (SSA) to forecast the number of patients arriving in the future, and consequently use these forecasts to ensure that an appropriate number of clinical decision makers are scheduled to be at work. A schematic summary of the proposed methodology applied to staffing levels is given in Figure 1. We begin by analysing historical data, and then after the processing of the algorithms described in the paper, as well as recording further data, the process may repeat in a cycle.




Figure 1: High-level description of proposed process

The paper is organised as follows. Section 2 will give a brief overview of time dependent queueing theoretical techniques used in this paper, as well as giving a detailed description of the particular algorithm used to obtain the appropriate number of staff. Section 3 will introduce SSA. In section 4 we will illustrate our approach with a case study. Conclusions and suggestions for further work are given in section 5.

2 Time Dependent Queueing Theory and Staffing Levels

Healthcare systems often operate in extremely stochastic conditions. Being able to cope with (and plan) for these conditions is of uttermost importance. One stochastic aspect is in the way demand for certain healthcare services varies over time. For example an EU will be relatively quiet in the middle of the week at night, but can expect an influx of patients over the weekend, particularly in the evening. As such it is reasonable to approximate an EU by a $M(t)/M/s(t)$ queue [19]. For readers that are not familiar with Kendall's notation, $M(t)/M/s(t)$ denotes a queueing system with negative exponential interarrival times (for this work, of patients) with time varying rate (denoted by $\lambda(t)$), a negative exponential service rate (denoted by μ) and time varying number of servers (denoted by $s(t)$). Various approaches have been used to give the smallest number of servers needed (for each time period t) to ensure a given performance standard. For a review the readers are encouraged to see [22].

Many existing methods use an amalgamation of steady state models [17, 18], whilst other methods are based on the 'offered load' of a system (the number of patients present in the equivalent $M(t)/M/\infty$ queue [23]). In [22] various approximations are tested and compared to what is referred to as the 'exact' method. In the 'exact' method the so-called Chapman-Kolmogorov equations are solved numerically in an iterative fashion, incrementing the staffing levels for each time period until the required service level is reached in order to meet some pre-specified target with high probability. The Chapman-Kolmogorov equations are given below (1).

$$\begin{aligned}
\frac{dp_0}{dt} &= -\lambda(t)p_0(t) + \mu p_1(t) \\
\frac{dp_n}{dt} &= \lambda(t)p_{n-1}(t) + (n+1)\mu p_{n+1}(t) - (\lambda(t) + n\mu)p_n(t), \quad 1 < n < s(t) \\
\frac{dp_n}{dt} &= \lambda(t)p_{n-1}(t) + s(t)\mu p_{n+1}(t) - (\lambda(t) + s(t)\mu)p_n(t), \quad n \geq s(t)
\end{aligned} \tag{1}$$

($p_n(t)$ denotes the probability of having n patients in the system at time t , $n \geq 1$.)

In [22] the role of this ‘exact’ method is to serve as a ‘gold-standard’ against which various approximation methods may be compared. The purpose of this paper is not to compare queueing theoretical models but the forecasts that feed them. Thus, we also use an exact method. The algorithm used is implemented in VBA for Excel and solves (1) using Euler’s algorithm. A copy of this programme is available at [25].

Most EUs work to targets based on the duration of wait. Given the nature of the injuries likely to arrive at an EU, major patients should ideally have minimal wait before receiving treatment. Therefore we aim to find minimum staffing numbers to ensure that a particular probability of waiting before being served is achieved. This probability is given by:

$$P_{\text{wait}} = 1 - \sum_{i=0}^{s(t)-1} p_i(t)$$

In the previously mentioned work, other considerations are given such as ensuring that the probability of a particular length of wait and/or abandonment levels are low. The emphasis for major patients on P_{wait} is justified by the severity of the injuries and the implications that a major patient suffering from a non immediate service has on the Ambulance service [3].

Figure 2 shows the staffing levels for 8 hour shifts over a week period obtained using this exact method that ensures $P_{\text{wait}} < .05$. The arrival rates used are taken from a case study described in section 4.




Figure 2: Staffing requirements over a week period ensuring $P_{\text{wait}} < .05$

As previously stated, the purpose of this paper is to investigate the forecasts that feed staffing level approximation algorithms. SSA will be used to generate such forecasts. In the bulk of the literature hourly averages are used (see for example [17, 18, 22]) and the references therein. Importantly however, this method fails to take into account the natural seasonality of the demand and thus can be improved using various techniques from time series (such as SSA, which is described in the next section).

The seasonality of EU data is typically complex, with the data exhibiting many seasonalities of different periodicities (see for example [32]). Examples include:

1. Yearly seasonality. Demand upon an EU typically increases year on year, with higher demand during the colder months. People with breathing and chest complaints typically have more problems during the warmer months.
2. Weekly seasonality. Demand upon an EU typically is higher at the weekend than the weekday.
3. Daily seasonality. Demand is often higher in the evenings (particularly at the weekends) than during working hours.

3 Singular Spectrum Analysis

3.1 Introduction

The SSA technique is a novel and powerful technique of time series analysis incorporating elements of classical time series analysis, multivariate statistics, multivariate geometry, dynamical systems and signal processing.

The birth of SSA is usually associated with the publication of papers by Broomhead and King (e.g. [7, 8]). A thorough description of the theoretical and practical foundations of SSA technique (with many examples) can be found in [16] and [21]. An elementary introduction to the subject can be found in [12]. Below, only a brief description of the methodology of SSA is given; for further details the reader is referred to [16] and [21] (and the references therein).

The main purpose of SSA is to decompose the originally observed time series into a sum of series, so that each component in this sum can be identified as either a trend, periodic or quasi-periodic component, or noise. This is followed by a reconstruction of the original series using selected components. The advantage of SSA is that we do not need to fit a parametric model to data. Whilst traditional time series models require restrictive distributional and structural data assumptions, these assumptions are not required by SSA. SSA has also been shown to have superior performance when compared to more traditional methods of time series modelling (such as ARIMA models); see for example [21]. SSA has proved to be very successful, and has already become a standard tool in the analysis of climatic, meteorological and geophysical time series; see, for example [11, 13, 30, 31]. Recently, SSA has been used to model demand upon the Welsh ambulance service [14, 32].

3.2 Basic algorithm

Let $F_N = (f_0, \dots, f_{N-1})$ be a one-dimensional series of length N . The main stages in performing SSA fall under the following four headings:

1. Embedding:

Choose a *window length* L ($1 < L < N$) and construct $K = N - L + 1$ lagged vectors

$$\mathbf{X}_i = (f_{i-1}, \dots, f_{i+L-2})^T, \quad 1 \leq i \leq K,$$

and put them into the $L \times K$ matrix

$$\mathbf{X} = [\mathbf{X}_1 : \dots : \mathbf{X}_K].$$

This is called the trajectory matrix. \mathbf{X} is a Hankel matrix, that is, all elements along the anti-diagonals are equal. This Hankelisation of the one-dimensional time series F_N onto the multidimensional series of lagged vectors can be viewed as a mapping $H: \mathbb{R}^N \rightarrow \mathbb{R}^{L \times K}$.

2. Decomposition:

The singular value decomposition (SVD) of \mathbf{X} yields

$$\mathbf{X} = \sum_{i=1}^d \sqrt{\lambda_i} U_i V_i^T$$

where $\{\lambda_i, i = 1, \dots, d\}$ are the set of eigenvalues of the matrix $\mathbf{X}\mathbf{X}^T$, and ordered such that $\lambda_1 \geq \dots \geq \lambda_d > 0$. $\sqrt{\lambda_i}$ is the i th singular value of \mathbf{X} . $(\sqrt{\lambda_i}, U_i, V_i)$ is known as the i th eigentriple. U_1, \dots, U_d and V_1, \dots, V_d are the orthonormal left and right singular vectors of \mathbf{X} respectively. U_1, \dots, U_d form a orthonormal basis of the column space L_d of \mathbf{X} .

3. Grouping:

By selecting m disjoint subsets I_1, \dots, I_m from the set of indices $\{1, \dots, d\}$ one obtains the decomposition

$$\mathbf{X} = \mathbf{X}_{I_1} + \dots + \mathbf{X}_{I_m},$$

where $\mathbf{X}_I = \sum_{i \in I} \sqrt{\lambda_i} U_i V_i^T$. This grouping step may be viewed as the decomposition of the trajectory space into an orthogonal sum of subspaces $L_d = \bigoplus_{j=1}^m L^{(j)}$, where $L^{(j)} = \text{span}(U_i, i \in I_j)$.

4. Reconstruction:

By averaging each of $\mathbf{X}_{I_1}, \dots, \mathbf{X}_{I_m}$ across their anti-diagonals yields the Hankelisation of $\mathbf{X}_{I_1}, \dots, \mathbf{X}_{I_m}$. Then $F_N = \tilde{F}^{(1)} + \dots + \tilde{F}^{(m)}$ where $\tilde{F}^{(j)} = (\tilde{f}_0^{(j)}, \dots, \tilde{f}_{N-1}^{(j)})$ and $\tilde{f}_i^{(k)}$ is the average value along the i th anti-diagonal of \mathbf{X}_{I_k} .

Hence in SSA we have to select two parameters: L , and the m disjoint subsets. A plot of the singular values (the square root of the eigenvalues $\sqrt{\lambda_i}$) identifies the number of eigenvectors to be taken to reconstruct the time series (in a similar manner to principal component analysis, see [24]). Explicit plateaux in the singular value spectra indicates pairs of eigenvectors that are likely to be important. Pairwise scatterplots of eigenvectors allow the visual identification of the eigenvectors corresponding to harmonic elements of the original series. Analysis of the periodograms from the original series, and of its eigenvectors, will inform of the frequencies that need to be considered when reconstructing the time series. Additional guidance for the selection of L and the m disjoint subsets is included in [16] and [21].

3.3 Forecasting using SSA: linear recurrent formulae (LRF)

SSA may be used to forecast any time series that (at least approximately) follow a linear recurrent formula (LRF). This follows due to the fact that if the number of terms in the SVD of \mathbf{X} is smaller

than the window length L , then the series satisfies some LRF. A key result to permit the forecasting of a series by SSA was stated in the book [16]: if the dimension of the linear space spanned by the columns of the trajectory matrix X is less than the window length L , then the series satisfies a LRF of dimension $L - 1$.

However, this assumption is not restrictive. For example, if the series F_N satisfies the LRF

$$f_{j+d} = \sum_{k=1}^d a_k f_{j+d-k}, \quad 0 \leq j \leq N - d - 1$$

for some real valued a_1, \dots, a_d , then the series F_N may also be written as a sum of products of exponentials, polynomials and harmonic components. That is,

$$f_j = \sum_{k=1}^q \alpha_k(j) \exp(\mu_k j) \sin(2\pi\omega_k j + \phi_k)$$

where $\{\alpha_k\}$ are polynomials, and $\{\mu_k\}$ $\{\omega_k\}$ and $\{\phi_k\}$ are arbitrary parameters.

The eigenvectors of \mathbf{X} provided in the SVD step described above yield the coefficients a_1, \dots, a_d . Confidence intervals for the resulting forecasts may be obtained by bootstrapping. For technical details the reader is referred to [16].

3.4 Implementing SSA

There are a number of ways people may implement SSA for their own use:

1. Specially developed software.

A number of software packages (some freely available) have been developed to perform SSA. Many offer additional extensions on the basic SSA algorithm described earlier. A webpage of the implementations currently available is maintained [1]. Facilities are also available to implement SSA through Visual Basic for Excel [15].

2. Using MATLAB.

An M-file which contains an implementation of the SSA algorithm as outlined in this paper is available from the authors.

3. Using statistical software such as R or SAS.

Recently the package `Rssa` for the open-source software R has been developed [26], additionally; the newer versions of SAS contains an implementation of the SSA algorithm [27].

In the next section we consider a case study using data from an EU of a large teaching hospital.

4 Case Study

In this work we use data obtained from the University Hospital of Wales (UHW). The UHW is a large teaching hospital based in Cardiff (Wales). EU data for the period between 01/02/2009 to 31/01/2010 concerning arrivals of patients categorised as ‘major’ (patients in urgent need of attention, not needing resuscitation) is used. The data are entered on a ‘real time’ basis by both administrative and clinical staff directly onto the hospital’s patient management system. The data

set contains a detailed description of the arrival times of patients which serves as input to the arrival rates. For the purpose of this paper, the data set is divided into two subsets. Firstly a ‘training set’, spanning the period 01/02/2009 to 01/01/2010. This subset is used to predict a month’s worth of arrivals and subsequently tested against the second data subset: the ‘test set’; spanning a period from 01/01/2010 to 31/01/2010. In practice, identifying service rates is very difficult in EUs. Patients are often blocked or waiting for large periods of time. The service rate ($\mu = 2$, per hour) chosen for this paper is chosen based on conversations held with members of the UHW management team and is in line with the service rates used in the literature [19].

4.1 Analysis of training set

Summary statistics of the hourly arrival rates are included in table 1.

Mean	3.93
Standard deviation	1.32
Median	4
Range	7.35
Min	0.875
Max	8.25

Table 1: Summary statistics of hourly arrival rates at UHW.

Large variation is observed in the arrival rates. This is to be expected; much media attention has focused on the volatility of the number of patients attending EU’s across the UK [5]. It is thus appropriate to carefully model this variation when forecasting future arrival rates. This volatility is also observed in the time series of the arrival rate (per eight-hour shift) given in Figure 3.




Figure 3: Graphs of arrival rates (per eight-hour shift)

As described in section 3, to perform the SSA method, we have to define two parameters: L and m . Often L is selected to be half of the number of recorded observations; see [16]. Using the notation of section 3 we have $N = 1002$ and hence $L = 501$. The plateaux in the plot of the logarithm of the singular values (recall, that these are the square root of the eigenvalues $\sqrt{\lambda_i}$) given in Figure 4 suggest that five components are needed to model the signal with a high accuracy; remaining components do not contribute significantly to the variation of the data. The techniques to select the number of components needed are identical to those employed in principal component analyses [24]. Table 2 gives the percentage of variation explained by each of the selected components.




Figure 4: Logarithm of first 10 singular values

Component	Percentage variation explained
First	89.962%
Second	3.117%
Third	3.105%
Fourth	0.118%
Fifth	0.116%

Table 2: Percentage of variation explained by each selected component (total 96.42%)

Plots of the first five components are included in Figure 5. The SSA reconstruction of the original series is obtained by simply summing each of these components. Pairwise scatterplots of these components give further insight of the nature of these components; by highlighting which are associated with periodic components of the originally observed data. This is already clear from investigations of the components given in Figure 5, but is even clearer upon inspection of the pairwise scatterplots. Figure 5 also contains pairwise plots of the second and third component; these components are associated with a daily periodicity corresponding to the division of a day into three shifts. The pairwise plot of the fourth and fifth component is also given; these components are associated with a weekly periodicity (corresponding to the division of a week into 21 shifts. Figure 6 consists of the periodogram of the second and third, and fourth and fifth components respectively. These periodograms demonstrate the seasonality of these paired components.

The first component is the trend, and explains the most variation in the time series. Two similar eigenvectors suggests the presence one sine-like component. As the second and third have a similar eigenvector, as do the fourth and fifth components, then the current reconstruction consists of a trend and two sine-like components. The decision is thus between having 3 components, and 5 components. 4 components would result in an incomplete sine being introduced into the reconstruction. Admittedly, the contribution of the fourth and fifth components together is rather small, and so the inclusion of them does not influence the reconstruction and forecasts that much. However, they do correspond to some seasonality inherent in the data, and it would be folly not to include this information in the reconstruction.

So the SSA reconstruction may be viewed as consisting of three disjoint subsets. The first subset consists of the first component, and is the overall trend. The second subset consists of the second and third component, corresponding to daily variation. The third subset consists of the fourth and fifth component, corresponding to weekly variation.

The SSA reconstruction for the original data is included in Figure 7.




Figure 5: Plot of first five principal components of the time series of the arrival rates (per eight-hour shift) and pairwise plots.




Figure 6: Plot of periodograms

4.2 Forecasting and comparing with test set

The forecast of the period from 01/01/2010 to 31/01/2010 is contained in Figure 7. The forecast compares extremely well to the actually observed arrival rates. The root mean square error between the SSA forecast and the actual arrival rates, was found to be: 0.8817 whilst the mean square error for predictions obtained using the the previously discussed 21 means was found to be: 2.0933. Note that the root mean square error between an estimate \tilde{y}_j and a true value y_j for n_{obs} observations is given by:

$$\sqrt{\frac{\sum_{j=1}^{n_{obs}} (\tilde{y}_j - y_j)^2}{n_{obs}}}$$




Figure 7: Original data (grey) with SSA fit and forecast based on the first five components (black)

4.3 Predicting staffing level requirements

Figure 2 shows the actual demand rate and staff requirements (obtained using the algorithm described in section 2) for the 1st week of the test set. The two approximations that will be tested against this method are found from:

1. Using the mean demand rate obtained from the training set for each eight hour shift of the week. Hence, this method uses $7 \times 3 = 21$ mean rates. Figure 8 shows the mean demand rates and the predicted staff requirements that ensure $P_{wait} < .05$. We refer to this approach as the ‘mean method’.
2. Using the predicted demand rate obtained using SSA. Figure 8 shows the mean demand rates and the predicted staff requirements that ensure $P_{wait} < .05$. We refer to this approach as the ‘SSA method’.

Although neither of these methods seem perfect, the SSA method appears to give a slightly better match to the staffing requirement based on actual demand. (Figure 2). We now set out two criteria




Figure 8: Staffing requirements over a week period ensuring $P_{wait} < .05$ using the mean and SSA method

for the evaluation of both of the mean method and SSA method which will allow for a more in-depth analysis.

One criteria for the evaluation and comparison of both methods would be the mean square error (MSE). For a set of staffing level estimations, \tilde{x} , and the set of staff requirements obtained observed using historical arrival rates output by the queueing model to ensure a particular level of P_{wait} , x (both indexed by shift), we have:

$$MSE(\tilde{x}) = \frac{1}{n_{obs}} \sum_{j=1}^{n_{obs}} (\tilde{x}_j - x_j)^2$$

The MSE penalises overestimation as much as underestimation. In terms of staffing for an EU, it is of utmost importance to have sufficient members of staff, particularly when treating major injuries. Thus we define the *weighted MSE* for a given $\alpha \in [0, 1]_{\mathbb{R}}$:


$$MSE_{\alpha}(\tilde{x}) = \frac{2}{n_{obs}} \left((1 - \alpha) \sum_{\tilde{x}_j > x_j} (\tilde{x}_j - x_j)^2 + \alpha \sum_{\tilde{x}_j < x_j} (\tilde{x}_j - x_j)^2 \right) \quad (2)$$

This measure is a modified MSE which weights the under-predictions by a factor of α , and the over-predictions by a factor of $(1 - \alpha)$. Note that (2) is normalised so that taking $\alpha = \frac{1}{2}$ recovers the standard MSE.


If the results obtained in this work were presented to hospital administrators, they would be interested in the actual cost of the staffing levels provided by each method. As such, the staffing of an EU can be interpreted as an inventory stock problem where there is a cost associated to having an insufficient level of supply (under-supply cost: C_u) and a cost associated to having a level of supply superior to that which is required (over-supply cost: C_o) (as such we have a situation similar to the so-called Newsboy Model inventory problem [33]). Given that the relative cost is of importance we take $C_o = 1$ and $C_u = k$ and use the following cost function:

$$cost_k(\tilde{x}) = \sum_{\tilde{x}_j > x_j} (\tilde{x}_j - x_j) + k \sum_{\tilde{x}_j < x_j} (x_j - \tilde{x}_j)$$


Initially we investigate the comparison of the mean method for predicting patient flow with the SSA method against the actual requirements for varying upper bounds of P_{wait} . Figure 9 presents the outcome measures for different values of α and k .




(a) $\alpha = .5$ and $k = 1$



(b) $\alpha = .65$ and $k = 1.5$



(c) $\alpha = .75$ and $k = 3$



(d) $\alpha = .85$ and $k = 10$

Figure 9: Weighted MSE and cost for different α and k respectively

Figure 9(a) uses the parameters $\alpha = .5$ and $k = 1$ which penalise overestimation and underestimation of staff levels equally. It is interesting to see that in this situation the mean method performs slightly better than the SSA method. In Figure 9(b) underestimation is penalised marginally more than overestimation and we begin to see that the SSA methodology produces better results. Figures 9(c) and 9(d) also demonstrate the superiority of the SSA method over the mean method; indeed, in Figure 9(d) we see that there is a potential saving (with cost parameter $k = 10$) of over 30% to be made if the SSA method is adopted.

We now pay particular attention to high levels of performance: that is an upper bound for P_{wait} of .05 or .1. Figure 10 gives the weighted MSE and $cost_k$ with $P_{wait} < .05$ and $P_{wait} < .1$ for varying parameters: α over $[0, 1]_{\mathbb{R}}$ and k over $[1, 19]_{\mathbb{Z}}$. Note that these two ranges give a weighted MSE that at first, penalises overestimation ($\alpha < .5$) and then penalises underestimation ($\alpha > .5$). For $cost_k$ however, we only assume that an underestimation is penalised.




Figure 10: Weighted MSE and cost for different α and k respectively, and different upper bounds for P_{wait}

The two pairs of graphs: Figures 10(a)-10(b) and Figures 10(c)-10(d) show that in both cases considered the SSA method performs better when seeking to minimise under-staffing of the EU. However, one immediately notices how the differences between the methods are much less for $P_{wait} < .1$ as opposed to when $P_{wait} < .05$. The explanation for this is based on the staffing algorithm. In essence the algorithm takes estimated arrival rates, and maps them onto the piecewise set consisting of particular staffing levels which meet some pre-described criteria. As the upper bound for P_{wait} decreases, the size of the segments of our piecewise set also decrease. This implies that the staffing algorithm slowly becomes a bijective mapping. As the SSA forecast follows the actual data much

closer (as seen in Figure 7) it is expected that for high levels of performance the SSA method does indeed outperform the mean method.

5 Conclusion

In this work we have investigated the use of singular spectrum analysis (SSA) to obtain improved staffing level rosters. The present paper builds on the various time dependent queueing theoretical approaches briefly described in section 2, by improving the forecasts which directly feed into the staffing algorithms discussed. SSA is a model-free technique which decomposes the original series into a sum of interpretable components such as a slowly varying trend and periodic components. It is also a non-parametric technique and so does not need the parametric assumptions needed to use other methods of time series analysis.

To assess the performance of our suggested approach, we used the classical mean square error (MSE) criterion and also introduced a weighted version of the MSE. This allows the underestimation of staff to be penalised more heavily than the overestimation of staff (or vice-versa). Our approach was shown to give favourable results when considering high levels of performance. This is due to the fact that the relationship between arrival rates and staffing levels becomes bijective as the required performance levels increase.

This work offers valuable insights that could be applied to other areas (and not only for healthcare settings) such as capacity models and traffic congestion models. Further work will include the use and development of metaheuristic approaches to obtain staffing rosters that match the obtained staffing requirements. For example, forecasts of the kind generated in this paper may be subsequently embedded into current operational research staffing methodologies to optimise resource allocation of staff to allow rapid response to potentially life-threatening emergencies so that, with high probability, response time targets are met. Faster patient intervention not only saves lives but reduces morbidity and potentially the need for longer-term healthcare needs and costs to national health programmes.

Acknowledgements

The authors would like to thank Professor Paul Harper (Cardiff School of Mathematics) for his constructive advice throughout the preparation and writing of this paper.

References

- [1] T. Alexandrov. SSAwiki: SSA knowledge base. <http://www.math.uni-bremen.de/theodore/ssawiki/>.
- [2] L. Au, G. B. Byrnes, C. A. Bain, M. Fackrell, C. Brand, D.A. Campbell, and P.G. Taylor. Predicting overflow in an emergency department. *IMA Journal of Management Mathematics*, 20(1):39, 2009.
- [3] BBC. Ambulances ‘facing regular delays at AE’. <http://www.bbc.co.uk/news/health-16052725>.
- [4] BBC. Patient handovers at hospitals ‘delaying ambulances’. <http://www.bbc.co.uk/news/uk-wales-politics-12190751>.

- [5] BBC. Welsh A&E waiting times miss target despite improvement. <http://www.bbc.co.uk/news/uk-wales-13563171>.
- [6] S. Brenner, Z. Zeng, Y. Liu, J. Wang, J. Li, and P. K. Howard. Modeling and analysis of the emergency department at University of Kentucky Chandler Hospital using simulations. *Journal of emergency nursing: JEN : official publication of the Emergency Department Nurses Association*, 36(4):303–10, July 2010.
- [7] D. S. Broomhead and G. P. King. Extracting qualitative dynamics from experimental data. *Phys. D*, 20(2-3):217–236, 1986.
- [8] D. S. Broomhead and G. P. King. On the qualitative analysis of experimental dynamical systems. In *Nonlinear phenomena and chaos (Malvern, 1985)*, Malvern Phys. Ser., pages 113–144. Hilger, Bristol, 1986.
- [9] B. Cardoen, E. Demeulemeester, and J. Belien. Operating room planning and scheduling: A literature review. *European Journal of Operational Research*, 201(3):921–932, 2010.
- [10] T.J. Coats and S. Michalis. Mathematical modelling of patient flow through an accident and emergency department. *Emergency Medicine Journal*, 18(3):190, 2001.
- [11] J. M. Colebrook. Continuous plankton records – zooplankton and environment, northeast atlantic and north sea, 1948–1975. *Oceanol. Acta*, 1:9 – 23, 1978.
- [12] J. B. Elsner and A. A. Tsonis. *Singular Spectrum Analysis: A New Tool in Time Series Analysis*. Plenum, 1996.
- [13] K. Fraedrich. Estimating the dimension of weather and climate attractors. *J. Atmos. Sci.*, 43:419–432, 1986.
- [14] J.W. Gillard, P.R. Harper, V.A. Knight, and J.L. Williams. Forecasting Welsh ambulance demand. In *2nd Student Conference in Operational Research*, 2010.
- [15] N. Golyandina, V. Nekrutkin, and K. Braulov. Caterpillar-SSA. <http://www.gistatgroup.com/cat/>.
- [16] N. Golyandina, V. Nekrutkin, and A. Zhigljavsky. *Analysis of time series structure*, volume 90 of *Monographs on Statistics and Applied Probability*. Chapman & Hall/CRC, Boca Raton, FL, 2001. SSA and related techniques.
- [17] L. Green, P. Kolesar, and J. Soares. Improving the SIPP approach for staffing service systems that have cyclic demands. *Operations Research*, 49(4):549–564, July 2001.
- [18] L. Green, P. Kolesar, and A. Svoronos. Some effects of nonstationarity on multiserver Markovian queueing systems. *Operations Research*, pages 502–511, 1991.
- [19] L. Green, J. Soares, J. Giglio, and R. Green. Using queueing theory to increase the effectiveness of emergency department provider staffing. *Academic emergency medicine : official journal of the Society for Academic Emergency Medicine*, 13(1):61–8, January 2006.
- [20] J.D. Griffiths, N. Price-Lloyd, M. Smithies, and J. Williams. A queueing model of activities in an intensive care unit. *IMA Journal of Management Mathematics*, 17(3):277, July 2006.
- [21] H. Hassani. Singular spectrum analysis: Methodology and comparison. *Journal of Data Science*, 5:239–257, 2007.

- [22] A. Ingolfsson, E. Akhmetshina, S. Budge, Y. Li, and X. Wu. A Survey and Experimental Comparison of Service-Level-Approximation Methods for Nonstationary $M(t)/M/s(t)$ Queueing Systems with Exhaustive Discipline. *INFORMS Journal on Computing*, 19(2):201–214, January 2007.
- [23] O.B. Jennings, A. Mandelbaum, W.A. Massey, and W. Whitt. Server staffing to meet time-varying demand. *Management Science*, pages 1383–1394, 1996.
- [24] I. T. Jolliffe. *Principal component analysis*. Springer Series in Statistics. Springer-Verlag, New York, second edition, 2002.
- [25] V. Knight. Personal Webpage. <http://cf.ac.uk/math/contactsandpeople/profiles/knightva.html>.
- [26] A. Korobeynikov. RSSA. <http://github.com/asl/rssa>.
- [27] M. Leonard, M. John, B. Elsheimer, and M. Kessler. An introduction to singular spectrum analysis with SAS/ets software. *SAS Global Forum 2010*, 2010.
- [28] S. A. Paul, M. C. Reddy, and C. J. DeFlicht. A Systematic Review of Simulation Studies Investigating Emergency Department Overcrowding. *SIMULATION*, 86(8-9):559–571, March 2010.
- [29] H. Smith, P. Harper, C. Potts, and A. Thyle. Planning sustainable community health schemes in rural areas of developing countries. *European Journal of Operational Research*, 193(3):768–777, March 2009.
- [30] R. Vautard and M. Ghil. Singular spectrum analysis in nonlinear dynamics, with applications to paleoclimatic time series. *Physica D*, 35:395 – 424, 1989.
- [31] B. C. Weare and J. S. Nasstrom. Examples of extended empirical orthogonal functions. *Monthly Weather Review*, 110:481 – 485, 1982.
- [32] J.L. Williams, J.W. Gillard, P.R. Harper, and V.A. Knight. Forecasting Welsh ambulance demand using singular spectrum analysis. In A. Testi, E. Tanfani, E. Ivaldi, G. Carello, R. Aringhieri, and V. Fragnelli, editors, *Operations research for patient - centered health care delivery. Proceedings of the XXXVI International ORAHS Conference*, 2010.
- [33] W. L. Winston. *Operations Research: Applications and Algorithms*. Brooks/Cole, 1998.