Artificial Intelligence in Medicine 73 (2016) 34–44

Contents lists available at ScienceDirect

Artificial Intelligence in Medicine

journal homepage: www.elsevier.com/locate/aiim

Out of hours workload management: Bayesian inference for decision support in secondary care

Iker Perez [a,*], Michael Brown [a], James Pinchin [a], Sarah Martindale [a], Sarah Sharples [b], Dominick Shaw [c], John Blakey [d]

[a] Horizon Digital Economy Research Institute, The University of Nottingham, Nottingham Geospatial Building, Jubilee Campus, Wollaton Road, Nottingham NG8 1BB, UK
[b] Faculty of Engineering, The University of Nottingham, Room B25, Coates Building, University Park, Nottingham NG7 2RD, UK
[c] Faculty of Medicine & Health Sciences, The University of Nottingham, Room B21, Nottingham City Hospital, Hucknall Road, Nottingham NG5 1PB, UK
[d] Department of Clinical Sciences, Liverpool School of Tropical Medicine, Pembroke Place, Liverpool L3 5QA, UK

ARTICLE INFO

Article history:
Received 26 April 2016
Received in revised form
29 September 2016
Accepted 29 September 2016

Keywords:
Healthcare management
Multivariate time series
Count data
Out of hours
Graphical model

ABSTRACT

Objective: In this paper, we aim to evaluate the use of electronic technologies in out of hours (OoH) task-management for assisting the design of effective support systems in health care; targeting local facilities, wards or specific working groups. In addition, we seek to draw and validate conclusions with relevance to a frequently revised service, subject to increasing pressures.
Methods and material: We have analysed 4 years of digitised demand-data extracted from a recently deployed electronic task-management system, within the *Hospital at Night* setting in two jointly coordinated hospitals in the United Kingdom. The methodology employed relies on Bayesian inference methods and parameter-driven state-space models for multivariate series of count data.
Results: Main results support claims relating to (i) the importance of data-driven staffing alternatives and (ii) demand forecasts serving as a basis to intelligent scheduling within working groups. We have displayed a split in workload patterns across groups of medical and surgical specialities, and sustained assertions regarding staff behaviour and work-need changes according to shifts or days of the week. Also, we have provided evidence regarding the relevance of day-to-day planning and prioritisation.
Conclusions: The work exhibits potential contributions of electronic tasking alternatives for the purpose of data-driven support systems design; for scheduling, prioritisation and management of care delivery. Electronic tasking technologies provide means to design intelligent systems specific to a ward, speciality or task-type; hence, the paper emphasizes the importance of replacing traditional pager-based approaches to management for modern alternatives.

1. Introduction

Secondary healthcare systems around the world are under increasing pressure [1,2]. Patient admissions are rising [3], and the number of available beds is falling [4]; simultaneously, the complexity of conditions and treatments is increasing [5]. Hence, healthcare systems must undergo major changes and optimise the use of limited resources.

This situation is especially acute during the *out of hours* (OoH) setting. For 75% of the working week hospitals are staffed by a skeletal team [6], and care must be delivered by a small and often junior group of clinicians over a wide range of medical specialities; frequently in large and complex sites [7]. Decreasing budgets [8], tighter controls of working hours [9] and the desire for separation of work and private life [10,11] have lead to a shrinkage of OoH working teams. Hence, in order to deliver safe healthcare of a consistently high quality the provision of this service is frequently revised [12–14], often without underlying comprehensive data or understanding of the demand placed on clinical teams.

Extensive research has been concerned with the study of expert and knowledge-based systems in healthcare management. This includes logistics, resource scheduling or estimation of service demands, and we refer the reader to [15–17] (and references therein) for some examples of this work. In particular, quantitative

* Corresponding author.
E-mail addresses: iker.perez@nottingham.ac.uk (I. Perez), michael.brown@nottingham.ac.uk (M. Brown), james.pinchin@nottingham.ac.uk (J. Pinchin), sarah.martindale@nottingham.ac.uk (S. Martindale), sarah.sharples@nottingham.ac.uk (S. Sharples), dominic.shaw@nottingham.ac.uk (D. Shaw), john.blakey@lstmed.ac.uk (J. Blakey).

http://dx.doi.org/10.1016/j.artmed.2016.09.005
0933-3657/

demand-forecasting studies focused on patient volumes (e.g. [18]) have confirmed that seasonal patterns and serial correlation structures play important roles in understanding demand loads. Additionally, the study of explanatory covariates in admission volumes has proved helpful in order to identify social pressures on workload (e.g. [19]). However, research restricted to global admission and consultation counts is insufficient in order to inform policy on local staff management; note that different patients receiving unrelated treatments over several medical disciplines put disparate workload pressures on specialist staff groups and grades; moreover, such pressures may vary drastically within distant geographical regions.

Hence, there is a need to employ modern embedded technologies for the design of effective management support systems that can target local facilities and specific working groups [20]; examples of such work include [21] and [22]. In this paper, we explore the use of electronic task-management for the study of OoH workload in secondary facilities. The purpose is twofold; on one hand, to exhibit the additional value that such tasking data provides in combination with modern machine learning methodology, for supporting intelligent scheduling, prioritisation and management of care delivery. On the other hand, to draw and validate conclusions with managerial relevance not restricted to the time-window and medical facilities covered in this study. We note that task-demand and completion numbers offer a better representation of workload, as opposed to admission or consultation counts; a lack of available sources has precluded previous quantitative studies of this kind.

To address these matters, we employ a state-space graphical model allowing for the extraction of patterns of cyclic variation in task-demand. This provides a useful framework to treat long series of observed multivariate counts, assuming independent observations conditionally on the values of a latent process. It allows to not underestimate true serial dependencies and control for discreteness and over-dispersion [23]; moreover, link functions may allow for intuitive interpretations of covariates' effects (e.g. [24]). Thus, by means of a latent parameter-driven model, we show it is possible to draw inference on contemporary and serial correlations on demand, over different medical and surgical specialities within a local facility. We also exhibit the ability to quantify future demand pressures, and we compare results with approaches relying on common methodology. Finally, we offer a summary of relevant conclusions as scrutinized by local medical staff, consultants and nation-wide healthcare organizations.

The data collected for this purpose was gathered from two major university hospitals, which combined provide secondary healthcare to over 2.5 million residents in the United Kingdom. In both hospitals involved, tasks for the team are requested, assigned and managed via web and mobile device interfaces; and the data is collected at each stage of this process, allowing work-demand to be monitored and analysed. The dataset used for this work contained 652,585 task requests and covers the period from January 2012 to December 2015.

Main results in the paper identify shared characteristics of OoH workload and display a significant split between medical and surgical specialities. Also, strong serial dependencies in demand series and a fast-decreasing predictive power over increasing time-windows emphasize the importance of short-term scheduling and prioritisation. Moreover, results support claims relating to both the importance of data-driven local staffing and work-demand forecasts serving as a basis to intelligent scheduling support. Patient and administrative needs vary significantly according to the day of the week or shift; notably, weekend planning must account for the variation between medical and surgical wards and bank holidays need to be treated as weekends; yet, workforce should (subject to few exceptions) be similar all year-round.

The rest of the paper is organised as follows; in Section 2 we review literature on state-space models with applications in time-series analysis, and explain the model employed in this study. Section 3 offers a discussion on alternative approaches and related work within a medical context. A description and summary of the data is offered in Section 4. In Section 5 we report results in the paper, and finally we offer conclusions and discuss policy recommendations in Section 6.

2. State-space models for multivariate count data

State-space models are a class of probabilistic graphical models (see [25]) that have found applications in time-series analysis for supply chain planning [26], text and music analysis [27], econometrics [28] or political science [29], to name a few. In particular, they allow us to describe the dependence between continuous latent state variables and discrete observed counts in terms of some probabilistic distribution; hence, they pose a useful mean to relate electronic tasking information with patterns of workload in healthcare facilities.

These models are either observation or parameter driven, and can extend generalized linear models by incorporating latent autoregressive processes within the conditional mean function (cf. [30,24]); thus introducing both auto-regression and over-dispersion. Less common multivariate extensions can handle both contemporary and serial correlations, and therefore tackle questions not addressed by marginal models (cf. [31,32]). In general, these models employ dynamic factors or vector auto-regressions at the latent level, and various technical examples of this work can be found in [28,31–33] and references therein.

Let $\boldsymbol{y}_t \in \mathbb{N}_0^n$ denote a vector of task counts as observed at day t; in our application, this reflects counts over n different medical and surgical specialities within the two hospitals. We now formulate a model such that $\boldsymbol{y}_t$ is drawn from a family of conditionally independent Poisson distributions, such that

$$y_{t,i}|h_t, \boldsymbol{\mu}_t \sim \mathcal{P}(h_t \cdot \mu_{t,i}), \quad \text{for } i \in \{1, \cdots, n\}, \tag{1}$$

and

$$\log(\boldsymbol{\mu}_t) = \sum_j \boldsymbol{f}_j(t) + \boldsymbol{\nu}_t, \quad \text{for } t \in \{1, \cdots, T\}.$$

Here, $\boldsymbol{\mu}_t \in \mathbb{R}_+^n$ denotes a latent rate vector of hourly-tasks, and $h_t \in \{16, 24\}$ indicates the working hours of the day. Thus, we aim to capture the relation within daily workload and temporal or autoregressive patterns by means of a log-link function. In the following, we discuss the components of the model along with a simplified example for the univariate series of global counts in Fig. 1.

2.1. Temporal trends in task-demand

The family $\{\boldsymbol{f}_j(t) \in \mathbb{R}^n : j \geq 1\}_{t \geq 0}$ are vector sequences defined to capture cyclic trends of task-demand; we note that these are log-multiplicative factors scaling hourly-task rates and hence the expected demand in (1). Particularly, $\boldsymbol{f}_1(t)$ includes $n \times 2$ coefficients $\boldsymbol{\beta}_1$ that represent independent linear trends of workload over different specialities; for instance, in Fig. 2 (left) we observe a credible interval for a linear growing trend in global task-demand specific to the 4 years covered in this data set, there we notice an approximate 48% increase on demand. Next,

$$\boldsymbol{f}_j(t) = \sum_{k=1}^{K_j} \boldsymbol{\beta}_{j,k} \cdot \left(\cos\left(\frac{2\pi kt}{P_j}\right), \sin\left(\frac{2\pi kt}{P_j}\right) \right)', \quad j \in \{2, 3\}$$

are Fourier terms with periodicities $P_2 = 7$ and $P_3 = 365$. These are frequently employed (cf. [28,33]) in both frequentist and Bayesian