Operations Research for Health Care 5 (2015) 1–9

Contents lists available at ScienceDirect

Operations Research for Health Care

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

Forecasting demand for health services: Development of a publicly available toolbox

Mehdi Jalalpour [a,*], Yulia Gel [b], Scott Levin [c,d]

[a] Washkewicz College of Engineering, Cleveland State University, United States
[b] Department of Mathematical Sciences, University of Texas at Dallas, United States
[c] School of Medicine, Johns Hopkins University, United States
[d] Whiting School of Engineering, Johns Hopkins University, United States

ARTICLE INFO

Article history:
Received 21 February 2014
Accepted 10 March 2015
Available online 20 March 2015

Keywords:
GARMA
Time series of counts
Decision-making under uncertainty
Multiple-steps ahead forecasts
Maximum likelihood
MATLAB

ABSTRACT

Efficient health care delivery systems aim to match resources to demand for services over time. Resource allocation decisions must be made under stochastic uncertainty. This includes uncertainty in the number of individuals (i.e., counts) in need of services over discrete time intervals. Examples include counts of patients arriving to emergency departments and counts of prescription medications distributed by pharmacies. Accurately forecasting count data in health care systems allows decision-makers to anticipate the need for services and make informed decisions about how to manage resources and purchase supplies over time.

A publicly available toolbox to forecast count data is developed in this work. The toolbox is implemented in MATLAB environment with the newly developed generalized autoregressive moving average (GARMA) models with discrete-valued distributions. GARMA models treat count data in a mathematically coherent manner compared to Gaussian models, often inappropriately applied in health care applications. GARMA models can incorporate none to many exogenous variables hypothesized to influence the predicted responses (i.e., counts forecasted). The toolbox's primary purpose is to deliver one to multiple-steps ahead forecasts, but also gives information for model inference and validation. The toolbox uses the maximum likelihood method to estimate model parameters from the data. We demonstrate toolbox application and validity on two example health care count data sets and show how using integer-valued conditional distributions as offered by GARMA models can produce forecast models that outperform the traditional Gaussian models.

© 2015 Elsevier Ltd. All rights reserved.

1. Introduction

Health care delivery systems are often stressed by uncertain (i.e., variable) demand for services over time [1]. A large source of this variation stems from the natural pattern in which individuals within a population need specific health care resources [2,3]. Managing scarce resources often requires forecasts of discrete health care events in the form of count data. Forecasting applications for count data may be useful for both short-term (e.g., hourly to monthly) and long-term (e.g., years) management and planning.

Short-term forecasting applications are useful to decision-makers at both the organizational and public health system levels [4]. Provider organizations such as hospitals, clinics, and networks of facilities must anticipate demand for services and match capacity, staffing, and supplies to be most efficient [2,5–12]. A lack of resources may lead to crowded care areas, long waits, and inhibited access. Excess and idle resources incur unnecessary costs and become unviable. Specific examples include forecasting patients arriving to an intensive care unit to maintain safe staffing levels [13,14,6,15,8], predicting pharmacy medication distribution to determine efficient inventory levels [16], and predicting patient no-shows for clinic appointments to create optimal patient and staff scheduling [17–20]. Clinician managers and health care administrators are commonly forecasting in short-term to manage daily operations within their organizations and understand variability. However, these forecasts are often based on human intuition and experience with little scientific understanding of contributing factors [8]. The same need to forecast count data in the short-term exists at the environmental and public health level, albeit for a much different purpose. Government-based health care

* Corresponding author.
E-mail address: m.jalalpour@csuohio.edu (M. Jalalpour).

http://dx.doi.org/10.1016/j.orhc.2015.03.001
2211-6923/© 2015 Elsevier Ltd. All rights reserved.

agencies such as the Center for Disease Control and Prevention (CDC), World Health Organization (WHO), and public health researchers apply forecasting mostly within the context of disease surveillance (epidemiologic count data) and health care utilization [21,22]. This is done to establish epidemiologic patterns in disease spread, understand contributing factors, detect outbreaks early, and evaluate public health measures [23,24,4,25–31]. However, these models must predict a wide range of magnitudes, which may be challenging when low counts exist [32–35], and be mathematically coherent.

At present, data-driven methods to forecast counts include discrete event simulation [36,7,37], and time series models [38–43], which is the focus of this study. For these types of data, traditional time series models (Gaussian) may not be adequate, because continuous distributions, which may include negative values, are not coherent with count responses that must be positive integers. A more realistic approach is to use a generalized model which can specify the conditional distribution of the response variable (weekly counts for example), given past values of the response as well as observations of exogenous variables (or covariates). Such models were first proposed by Liang and Zeger [44] and Li [45] and subsequently extended by Benjamin et al. [46] to generalized mixed autoregressive moving average models (GARMA) with discrete-valued distributions suitable for count data. Another key advantage of GARMA models is their parsimonious parametrization that accommodates non-stationary dynamics of the observed time series.

The objective of this paper is to develop an open-source toolbox to forecast count data sets deploying GARMA models. This toolbox may be useful to health care decision-makers and managers studying health care operations from within provider organizations and across public health systems [36,7,37,39–43]. Despite the need, there is a clear lack of publicly available packages offering GARMA implementation. This is partly because GARMA models are nonlinear making their parameter estimation a non-trivial task. Benjamin et al. [46], who proposed these models initially, used `GLIM` (Generalized Linear Interactive Modeling) software [47] for model fitting, and they also mentioned the possibility of using a weighted least squares fitting procedure in `S-PLUS`. `GLIM`, however, is no longer actively supported or distributed, but has been advanced to GAMLSS within `R` environment [48]. The other `R` alternative is the VGAM package [49,50]. However, this package is limited to autoregressive structures and in-sample predictions, does not handle overdispersion, and is sensitive to initialization (likely due to the `R` optimizer quality). To the authors' knowledge, there are no other packages offering a GARMA option and no single implementation of GARMA in `MATLAB`, which is especially critical given that `MATLAB`, a high-level programming language, remains one of the most widely used analytical software systems in the field of operations research [51–54]. The newly developed **GARMA** toolbox also includes routines for multiple-steps ahead out-of-sample forecasting, model verification, and inference. In addition, we take advantage of `MATLAB` unconstrained optimization algorithm for delivering a more precise and stable maximum likelihood estimator. Previous research by our team studying the ability of Google influenza search query data (i.e., Google Flu Trends) [31] to predict weekly counts of patients with influenza presenting to an inner city emergency department (ED) motivated the development of this open-source toolbox.[1] The toolbox may be used to forecast counts at any pre-defined discrete time intervals. However, the time series approach deployed is most relevant for short-term forecasting in health care.

We illustrate application of the **GARMA** toolbox on a public health level data set of monthly counts of polio cases reported by the Center for Disease Control (CDC) and on weekly counts of patients with influenza at the Johns Hopkins Hospital emergency department at Baltimore, Maryland, US (health care facility level). Using these examples, we demonstrate the toolbox routines for model fitting, validation, inference, and out-of-sample forecasting. We also show how using integer-valued conditional distributions as offered by GARMA models can produce forecast models that outperform the traditional Gaussian models for these examples. We believe the improved performance for GARMA compared to Gaussian models may be generalized to other discrete-valued count data sets encountered in health care decision-making applications.

2. Methods

2.1. Notation and derivations

We assume responses (counts of patients for example) and exogenous variables data (average temperature and humidity for example) are collected at equal time intervals (weekly for example). There is a total of T time intervals, each denoted by t, and n exogenous variables. Therefore, the responses and exogenous variables can be stored in a vector $\mathbf{y}$ of length T and a matrix $\mathbf{X}$ with T rows and n columns respectively. Each observation for response and exogenous variables are denoted by y_t and $\mathbf{x}_t^\top$ respectively, where $\mathbf{x}_t^\top$ is $n \times 1$ vector. GARMA models assume that the *conditional* distribution of response given past information (responses, exogenous variables, and mean values) belongs to the family of exponential distributions (for example Poisson, Gaussian, and Negative-Binomial). Moreover, the conditional expected value of response is connected to past information by a one-to-one monotonic link function. The general GARMA(p, q) model for one-step ahead prediction with respect to the responses is expressed as [46]:

$$g(\mu_t) = \eta_t = \mathbf{x}_t^\top \boldsymbol{\beta} + \sum_{j=1}^{p} \phi_j \{g(y_{t-j}) - \mathbf{x}_{t-j}^\top \boldsymbol{\beta}\} + \sum_{j=1}^{q} \theta_j \{g(y_{t-j}) - \eta_{t-j}\}. \tag{1}$$

In Eq. (1), p and q are model autoregressive and moving average orders respectively, and $\boldsymbol{\beta}, \boldsymbol{\phi}, \boldsymbol{\theta}$ are vectors of model parameters with dimensions n, p, and q respectively. The link function $g(.)$ maps the expected value of response (μ_t) to the linear predictor η_t. The link functions are canonical, identical to generalized linear model (GLM) [55]. The second term in Eq. (1) is the autoregressive, and the third term is the moving average component of the model. A GARMA model without these terms reduces to a GLM.

The maximum likelihood method is used to estimate model parameters. The likelihood function for the data can be written as:

$$f(y_T, y_{T-1}, \cdots y_1 | \mathbf{D}_T; \boldsymbol{\gamma}) = f(y_T|\mathbf{D}_T; \boldsymbol{\gamma}) \times \cdots f(y_m|\mathbf{D}_m; \boldsymbol{\gamma}) \times f(y_{m-1}; \boldsymbol{\gamma}) \times \cdots f(y_1; \boldsymbol{\gamma}), \tag{2}$$

where $m = \max(p, q)$ and $\mathbf{D}_t = \left[\mathbf{x}_t, \ldots \mathbf{x}_1, y_{t-1}, \ldots, y_1, \mu_{t-1}, \ldots \mu_1\right]$ is the matrix of past information and $\boldsymbol{\gamma} = \boldsymbol{\beta}, \boldsymbol{\phi}, \boldsymbol{\theta}$ is the vector of model parameters. The left side of Eq. (2) is the joint probability density function of responses which may be expressed in terms of conditional and marginal distributions as seen in the right side of Eq. (2). Conditioned on the first $m-1$ observations and assuming that each response given past information follows the same exponential family distribution, the log-likelihood function for the data is expressed as [46]:

$$L = \sum_{t=m}^{T} \log f(y_t|\mathbf{D}_t; \boldsymbol{\gamma}), \tag{3}$$

where $f(.)$ is the probability density function (pdf) of the assumed exponential family. The set of exponential family distributions

[1] The new **GARMA** toolbox is currently available from https://sites.google.com/site/mehdijalalpour/research.

Download English Version:

https://daneshyari.com/en/article/7543691

Download Persian Version:

https://daneshyari.com/article/7543691

Daneshyari.com