European Journal of Operational Research 204 (2010) 139–152

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Stochastics and Statistics

Triple seasonal methods for short-term electricity demand forecasting

James W. Taylor *

Saïd Business School, University of Oxford, Park End Street, Oxford OX1 1HP, UK

ARTICLE INFO

Article history:
Received 1 September 2008
Accepted 6 October 2009
Available online 17 October 2009

Keywords:
Forecasting
Electricity demand
Seasonality
Exponential smoothing
ARMA

ABSTRACT

Online short-term load forecasting is needed for the real-time scheduling of electricity generation. Univariate methods have been developed that model the intraweek and intraday seasonal cycles in intraday load data. Three such methods, shown to be competitive in recent empirical studies, are double seasonal ARMA, an adaptation of Holt–Winters exponential smoothing for double seasonality, and another, recently proposed, exponential smoothing method. In multiple years of load data, in addition to intraday and intraweek cycles, an intrayear seasonal cycle is also apparent. We extend the three double seasonal methods in order to accommodate the intrayear seasonal cycle. Using six years of British and French data, we show that for prediction up to a day-ahead the triple seasonal methods outperform the double seasonal methods, and also a univariate neural network approach. Further improvement in accuracy is produced by using a combination of the forecasts from two of the triple seasonal methods.

1. Introduction

The efficient management of power systems requires accurate forecasts of electricity demand (load) for a range of lead times. To enable real-time scheduling, system operators produce forecasts for short lead times using online automated methods. In deregulated energy markets, such predictions are also of importance to market participants to support energy transactions (Bunn, 2000).

Electricity demand is often modelled in terms of weather variables (e.g. Hor et al., 2005; Cancelo et al., 2008). However, univariate methods are frequently considered to be sufficient for short lead times because the weather variables tend to change in a smooth fashion over short time frames, and this will be captured in the demand series itself. Furthermore, weather-based online systems require default procedures in order to ensure robustness (Bunn, 1982). In a recent empirical study (Taylor, 2008a), a univariate method was shown to outperform a multivariate method up to about four hours ahead, and a combination of forecasts from the two methods was found to be the most accurate approach up to a day ahead. This shows that univariate methods have a valuable role to play in short-term load forecasting. Of course, such methods are the only option when forecasting load in locations where weather forecasts are either unavailable or too costly (Soares and Medeiros, 2008). In this paper, we focus on the use of univariate methods for prediction up to one day ahead.

A variety of univariate methods have been used for short-term load forecasting. Seasonal ARIMA models of various types have been considered by, among others, Hagan and Behr (1987), Moghram and Rahman (1989), Mbamalu and El-Hawary (1993), Taylor and McSharry (2007) and Soares and Medeiros (2008). The approach of Park et al. (1991) mixes autoregressive modelling with exponential smoothing and exponentially weighted least squares. The latter is also known as general exponential smoothing, and this was the focus of the early study of Christiaanse (1971). Given its widespread use in many business forecasting applications (see Hyndman et al., 2008), it is surprising that exponential smoothing had not received more attention from load forecasting researchers prior to the study by Taylor (2003). Although the nonlinear and nonparametric features of neural networks are attractive for modelling load in terms of weather variables (see Hippert et al., 2001), their appeal for univariate modelling is less obvious and rests on the possible existence of nonlinearities in the structure of the load time series. Pai and Hong (2005) implement an ARIMA model and a neural network as benchmarks in their study of support vector machines, which also enable nonlinear regression modelling. A univariate method based on principal component analysis is considered by Taylor and McSharry (2007). The approach aims to simplify the forecasting task by extracting the important underlying seasonal components in the time series.

A feature of time series of intraday electricity demand is the presence of intraday and intraweek seasonal cycles. In two recent empirical studies, several univariate statistical methods were considered that aim to accommodate both of these seasonal cycles (Taylor and

* Tel.: +44 (0)1865 288927; fax: +44 (0)1865 288805.
E-mail address: james.taylor@sbs.ox.ac.uk

0377-2217/$ - see front matter
doi:10.1016/j.ejor.2009.10.003

McSharry, 2007; Taylor, 2008a). Three of the more successful methods were a double seasonal ARMA model, an adaptation of Holt–Winters exponential smoothing for double seasonality, and an exponential smoothing method recently proposed by Gould et al. (2008). Those empirical studies used intraday time series each of length 30 weeks. With multiple years of observations, in addition to intraday and intraweek cycles, a further seasonal cycle is evident, namely the annual cycle. In this paper, our primary contribution is the extension of the three double seasonal methods to include the intrayear seasonal cycle. Although these extensions for triple seasonality are relatively simple, they have not previously been considered in the literature.

Of course, there is no guarantee that allowing for the intrayear cycle in a statistical model will lead to improved accuracy for prediction up to just one day ahead. Therefore, an important contribution of this paper is the empirical study in which we use two time series of half-hourly demand to compare post-sample forecasting results for single, double and triple seasonal formulations of the methods. The empirical results provide insight for the number of seasonal cycles to include in a method, for the relative importance of the different types of seasonal cycles, and for the choice between the two exponential smoothing methods and ARMA modelling. The study also evaluates the benefit in incorporating a residual autocorrelation term in the exponential smoothing method of Gould et al.

Research in this area would seem to be timely, given the increasing availability of intraday data in a variety of other applications, such as traffic management and call centre staff scheduling (see, for example, Lam et al., 2006; Taylor, 2008b). Another potential application area for the methods in this paper is the modelling of intraday electricity prices, which is of primary importance for trading electricity. The methods would have to be carefully adapted for prices because they typically possess rather complex structure with substantial volatility and mean-reverting spikes.

In Section 2, we describe the two load series. Section 3 presents ARMA model formulations for data with single, double and triple seasonality. The focus of Section 4 is exponential smoothing formulations for single, double and triple seasonality based on the Holt–Winters method. In Section 5, the method of Gould et al. (2008) is presented, along with its extension for triple seasonality. In Section 6, we describe the neural network that we use as a sophisticated benchmark in our empirical comparison of methods. Section 7 presents our empirical study, which evaluates forecast accuracy for lead times from one half-hour ahead up to one day ahead. In Section 8, we summarize and provide concluding comments.

2. The British and French load series

In this study, we consider two load series; one is for Great Britain and the other is for France. Each consists of six full years of half-hourly observations for electricity demand from 2001 to 2006, inclusive. Our empirical analysis used the first five years of data to estimate forecasting method parameters and the remaining year to evaluate post-sample forecast accuracy. This gave 87,648 half-hourly observations for estimation, and 17,520 for evaluation. The two series are shown in Fig. 1. Both show a strong intrayear seasonal cycle with demand much lower during the warmer months of each year. A closer look at these warmer months in the French series reveals a period of four weeks of particularly low demand, which corresponds approximately to August. Although this is not apparent from a visual inspection of the British series in Fig. 1, the British transmission company does try to model the effect of the three summer weeks when a large amount of industry closes.

Fig. 2 shows a winter fortnight and a summer fortnight from both series. Within each fortnight, the weekdays show similar patterns of demand, but the Saturdays and Sundays are somewhat different. These features are typical of series of intraday electricity demand. Interestingly, for both the British and French data, the seasonality is considerably different for the winter and summer. This indicates that, when modelling the seasonality, the pattern should be allowed to change across the year.

In each series, we identified a number of unusual days, termed 'special' days, on which demand differed considerably from the regular seasonal pattern. These were either public holidays or days adjacent or close to public holidays. In short-term online prediction systems, these days are typically replaced by forecasts prepared offline. As our study is concerned with online prediction, prior to fitting and evaluating the forecasting methods, we chose to smooth out the special days leaving the natural periodicities of the data intact. In our study, this smoothing was performed by simple averaging procedures, which usually involved replacing demand on each special day period by

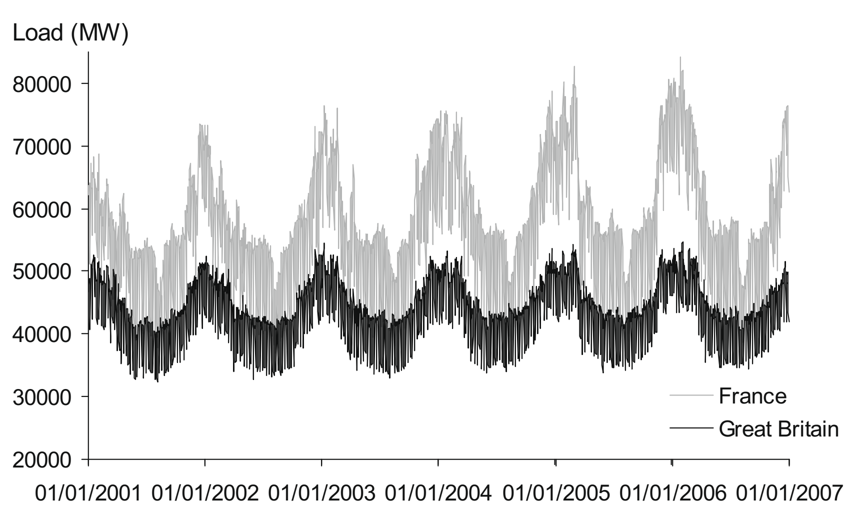

Fig. 1. Electricity load in France and Great Britain for 2001–2006.