Forecasting Frequency-Corrected Electricity Demand to Support Frequency Control

James W. Taylor and Matthew B. Roberts

IEEE Transactions on Power Systems, 2016, Vol. 31, pp. 1924-1932.

***Abstract*—Electricity demand forecasts are needed for decisions regarding generation dispatch for lead times as short as just a few minutes. Imbalance between generation and demand causes deviation of the system frequency from its target, which in Great Britain is 50Hz. This, in turn, causes a change in demand, due largely to motor loads. For Great Britain, the change is estimated to be 2.5% of demand per 1Hz of frequency deviation from its target. This can be used to calculate the demand that would have occurred if frequency had been at 50Hz. Modeling and forecasting the resulting frequency-corrected demand provides a better basis for dispatching generation. This paper evaluates methods for forecasting frequency-corrected demand up to 10 minutes ahead. We introduce an exponential smoothing model that, like the system operator's proposed Kalman Filter approach, jointly models frequency and demand. We also evaluate a set of univariate methods applied directly to the series of frequency-corrected demand. These methods have not previously been considered for lead times less than 10 minutes. In our empirical analysis, the best results were produced by a seasonal exponential smoothing method applied directly to the series of frequency-corrected demand.**

***Index Terms*— Electricity demand forecasting, frequency-load control, system frequency, time series models.**

I. INTRODUCTION

A slight imbalance between electricity demand and generation causes the system frequency to fluctuate [1]. The transmission system operator (TSO) in Great Britain is required to ensure that the system frequency does not deviate greatly from the target of 50Hz. With this aim, throughout the day, the TSO instructs generators as to how much electricity to provide. These instructions can be reactionary, in response to sudden changes in load and deviations of frequency from the required level, but they should also be based on demand forecasts [2], as generators require several minutes notice if adjustments to generation are required. If demand turns out to be greater than the amount generated, the system frequency will fall, as the energy has to come from slowing down the various generators. The lower frequency will result in a lower demand, and self-regulating effect helps to stabilize the system frequency [3,§4.2]. It is also the case that, if generation exceeds demand, the frequency rises, leading to a rise in demand. The sensitivity of load to frequency is due largely to motor loads, as well as power devices controlled by semiconductors that allow variation in the power supplied, such as switched-mode power supplies in personal computers [4]. An example of the self-regulating effect is provided in [4], in which a fault led to an unexpected generation loss of 1,050MW in Scotland. Only 64% of the loss was recovered using reserves, with the remainder accounted for by the natural sensitivity of demand to frequency. In view of this, the sensitivity of load to frequency should be considered in setting the optimal level of reserve [4].

The sensitivity of demand to deviations in frequency varies across different electricity systems (see, for example, [5,6,7]). For the Great Britain electricity grid, the TSO estimates the impact as about 2.5% of demand per 1Hz of frequency deviation from the target of 50Hz. In view of this, we can derive the level of demand that would have occurred if frequency had been at 50Hz. We refer to this as the frequency-corrected demand C_t, and it is calculated as in expression (1), where D_t is the demand, F_t is the frequency, and c is the load-frequency sensitivity, which we term the *correction factor* and which we assume to be 2.5% for Great Britain.

$$C_t = D_t + c\left(50 - F_t\right)D_t \qquad (1)$$

In considering future periods, the aim is to have the system operating at its target frequency of 50Hz. Therefore, generation dispatchers need to predict the demand under the assumption that the frequency will be at this target. If frequency turns out not to be 50Hz, the forecasts should not be evaluated against observed demand, but should instead be evaluated against the demand that would have been observed if the system had been operating at 50Hz. This is the frequency-corrected demand. Indeed, to try to maintain a stable frequency at 50Hz, instead of delivering instructions to generators based on forecasts of demand, it is appropriate to use forecasts of frequency-corrected demand C_t. In this paper, our focus is to evaluate the point forecast accuracy of different methods for predicting minute-by-minute observations of C_t. In view of the limitations on generator response times, for Great Britain, the main focus of the TSO is a lead time of 3 minutes. However, to provide broader insight, we investigate accuracy for lead times from 1 to 10 minutes.

Load is often modeled in terms of weather variables. However, this has tended not to be the case when forecasting at lead times of less than about an hour, which can be termed *very short-term* load forecasting. For such lead times, weather predictions are of questionable worth, as the recent load

J. W. Taylor is with the Saïd Business School, University of Oxford, Park End Street, Oxford OX1 1HP, UK (e-mail: james.taylor@sbs.ox.ac.uk).

M.B. Roberts is with National Grid, St Catherine's Lodge, Bearwood Rd, Sindlesham, RG41 5BN, UK (e-mail: matthew.b.roberts@nationalgrid.com).

observations can adequately capture the very-short term evolution in the load time series. For very short-term forecasting, artificial neural networks (ANNs) are considered in [2], along with models based on fuzzy logic and autoregressive models. The analysis uses only the historical load observations. ANNs are also used in [8], and although temperature is considered as a possible input, the models presented include inputs that are functions only of past load. In [9], hourly load forecasts are employed as the basis for predicting load from 5 minutes up to 2 hours ahead. A cubic spline is used to convert hourly forecasts to predictions at a 5-minute sampling rate. A Kalman Filter is then applied to these predictions and load data recorded at a 5-minute to produce forecasts from 5 minutes to 2 hours ahead. Abductive networks are used in [10], and it is found that, for one hour-ahead prediction, temperature forecasts are not useful, with load from various lags being sufficient inputs. A forecast horizon of 15 minutes is considered in [11], where an ANN is implemented, based on fuzzy logic and chaotic dynamics reconstruction techniques. The approach is applied to a dataset consisting of only the historical load observations. This is also the case with the wavelet ANNs used in [12,13], which incorporate a pre-processing filter in order to smooth spikes occurring in the data due to malfunctioning of data recording devices.

In this paper, we compare a set of forecasting methods that are applied directly to the time series of frequency-corrected demand. Several of these were included in the empirical study in [14], which looked at the forecasting of demand for lead times from 10 to 30 minutes. In this paper, our focus differs because we apply these methods to frequency-corrected demand, and we consider shorter lead times. Also, we evaluate other methods from the literature, including an ANN, and an exponential smoothing method based on singular value decomposition. In addition, we implement two methods that are not applied to the time series of frequency-corrected demand, but instead jointly model frequency and demand. One of these is the Kalman Filtering approach proposed by the TSO.

Section II describes the data used in our empirical study. Section III presents the two forecasting methods that jointly model frequency and demand. Section IV describes the methods that directly model the frequency-corrected demand series. Section V presents the empirical results. Section VI provides concluding comments.

II. The Frequency and Demand Data

The data used in this study was supplied by the TSO in Great Britain. It consists of the 29 weeks of minute-by-minute observations for electricity demand and frequency in Great Britain from Sunday 7 April 2013 to Saturday 26 October 2013. This constitutes 292,320 observations for each series. Following the advice of the TSO, if a value of demand was more than 1,000MW lower than the previous demand value, it was assumed that the recording meter had malfunctioned, and so the value was replaced by the average of demand from the two adjacent periods. This was needed for only 12 periods. An alternative to this simple smoothing approach is to use the filtering approach in [12]. We used the first 20 weeks of data to estimate forecasting method parameters, and the remaining 9 weeks to evaluate post-sample forecast accuracy.

The methods considered in this paper are not satisfactorily able to model special days, such as public holidays. Prediction for such days is typically performed separately offline. In our 29-week sample, there were just three special days; two public holidays in the estimation sample and one in the post-sample period. For the estimation sample, we smoothed over each period of the special days by using the average of the corresponding periods in the two adjacent weeks. For the special day in the post-sample period, we smoothed by averaging the observations from the corresponding periods in the two previous weeks. The resulting smoothed values were not used for model estimation or evaluation. An alternative to forecasting special days offline is to adapt the models presented in [15], which were designed for half-hourly data.

Fig. 1 shows the first fortnight of the electricity demand time series. The plot shows, at least for the weekdays, a repeating cycle of length m_1=24×60=1,440 minutes. There is also a repeating cycle across each week, which is of length m_2=7×24×60=10,080 minutes. The time series of system frequency, for the same two-week period, is shown in Fig. 2. This series shows high volatility around the target of 50Hz.

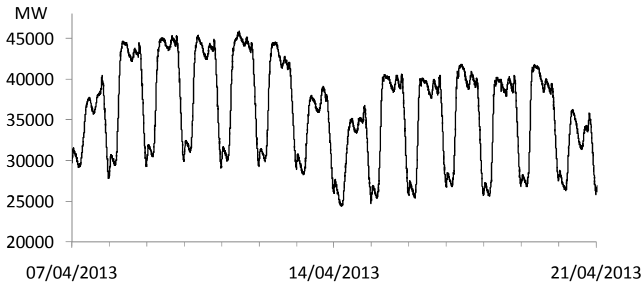

Fig. 1. Minute-by-minute electricity demand for a two-week period.

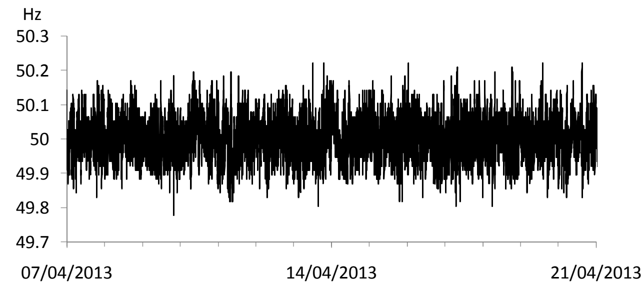

Fig. 2. Minute-by-minute frequency for a two-week period.

Fig. 3 presents the autocorrelation functions (ACF) for demand and frequency. The demand ACF confirms the seasonality apparent in Fig. 1, with strong autocorrelation at lags that are multiples of m_1, and particularly high autocorrelation at multiples of m_2. The frequency ACF shows strong autocorrelation at the early lags, and also increased values of the autocorrelation at lags that are multiples of m_1.

These values, which are significant at the 5% level, indicate some degree of seasonality in the frequency series.

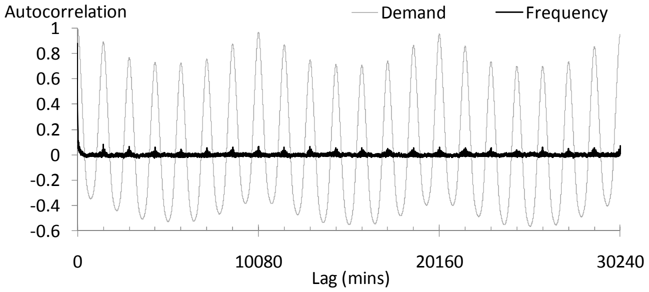

Fig. 3. Autocorrelation function for demand and frequency.

A plot of a fortnight of the frequency-corrected demand series shows the same repeating cyclical patterns evident in the demand plot of Fig. 1. As we are interested in lead times of just a few minutes, it is informative to plot frequency-corrected demand for a much shorter period, along with demand and frequency. We do this in Fig. 4 for one hour on the first Monday of our dataset. The figure shows the frequency-corrected demand for three values of the correction factor, c. Although the TSO assumes c=2.5%, this value is an estimate [4], and so it is interesting to consider alternative values. The plot of frequency in Fig. 4 confirms the finding from Fig. 3 that it is autocorrelated. As indicated by expression (1), when frequency is below 50Hz, Fig. 4 shows frequency-corrected demand above demand, and when frequency is above 50Hz, frequency-corrected demand is below demand. Given the form of expression (1), it is understandable that frequency-corrected demand shows greater volatility than demand in Fig. 4.

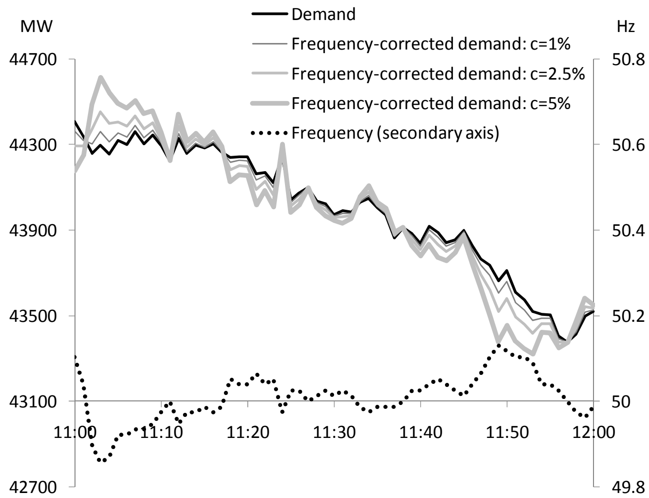

Fig. 4. For a one hour period on Monday 8 April 2013, minute-by-minute observations for frequency (plotted against the secondary axis), demand and frequency-corrected demand for three values of the correction factor c.

III. Jointly Modeling Frequency and Demand

In this section, we present two methods that jointly model demand and system frequency. In each method, demand and frequency forecasts are generated, and then fed into expression (1) to give predictions for frequency-corrected demand.

A. Kalman Filter Approach

We implemented a Kalman Filter approach proposed by the TSO that involves modeling state variables for demand d_t, growth in demand Δ_t^d, frequency f_t, and growth in frequency Δ_t^f. The state equations for the model are the following:

$$d_t = d_{t-1} + \Delta_{t-1}^d + c\,d_{t-1}\Delta_{t-1}^f + u_{1,t-1} \quad (2)$$

$$\Delta_t^d = \Delta_{t-1}^d + u_{2,t-1} \quad (3)$$

$$f_t = f_{t-1} + \Delta_{t-1}^f + u_{3,t-1} \quad (4)$$

$$\Delta_t^f = \Delta_{t-1}^f + u_{4,t-1} \quad (5)$$

where the u_{it} are error terms. Expression (2) views the change in demand as consisting of the sum of three terms: growth, the impact of the frequency growth on demand, and error.

The observed demand D_t and frequency F_t are related to the state variables through the following observation equations:

$$D_t = d_t + v_{1,t}$$

$$F_t = f_t + v_{2,t}$$

where the v_{it} are error terms. The TSO assumed the following: v_{it} and u_{it} are Gaussian, serially uncorrelated, and uncorrelated with each other; v_{1t} and v_{2t} are uncorrelated; and u_{1t} and u_{2t} are not correlated with u_{3t} or u_{4t}. The variances and correlations were treated as parameters. To assist prediction for a given lead time, parameters λ_1 and λ_2 were incorporated in the forecasting of the demand and frequency using the following:

$$D_{t+k|t} = d_{t|t} + \lambda_1\,\Delta_{t|t}^d$$

$$F_{t+k|t} = f_{t|t} + \lambda_2\,\Delta_{t|t}^f$$

The parameters were optimized separately for each lead time by minimizing the sum of squared errors for the forecasts of frequency corrected-demand. The TSO proposed constraints on the parameters to assist the optimization. However, we found improved accuracy by removing these constraints. For a lead time of 3 minutes, which is the horizon of greatest interest to the TSO, we obtained λ_1=2.73, λ_2=0, var(v_{1t})=3.36, var(v_{2t})=0, and the covariance matrix of u_{it} was estimated as:

$$\begin{pmatrix} 8.89 & 0.655 & 0 & 0 \\ 0.655 & 0.046 & 0 & 0 \\ 0 & 0 & 9.38 & 0.0760 \\ 0 & 0 & 0.0760 & 0 \end{pmatrix}$$

B. Exponential Smoothing for Frequency and Demand

A weakness of the Kalman Filter approach is that it does not model seasonality. Standard state-space formulations for seasonality would imply a model with more than m_2 states (see, for example, [16,§3.2.1]), which would be impractical. This motivates an alternative joint model of demand and frequency. Our proposal is to adapt HWT exponential smoothing, which aims to model intraday and intrayear seasonality. The HWT model is presented in [17] as follows:

$$y_t = l_{t-1} + d_{t-m_1} + w_{t-m_2} + \phi e_{t-1} + \varepsilon_t \quad (2)$$

$$e_t = y_t - \left(l_{t-1} + d_{t-m_1} + w_{t-m_2}\right) \quad (3)$$

$$l_t = l_{t-1} + \alpha\, e_t \quad (4)$$

$$d_t = d_{t-m_1} + \delta\, e_t \quad (5)$$

$$w_t = w_{t-m_2} + \omega e_t \tag{6}$$

where y_t is the target variable, l_t is the smoothed level; d_t is the seasonal index for the intraday cycle; w_t is the index for the intraweek cycle remaining after the intraday cycle is removed; ε_t~N(0,σ^2); σ^2 is a constant variance; α, δ and ω are smoothing parameters; and ϕ is a residual autocorrelation parameter.

We used the HWT model for frequency F_t. Our model for demand was adapted from the HWT model, and is presented in expressions (7)-(11). For this model, the underlying demand is viewed as being best represented by frequency-corrected demand, and so we model the level and seasonality of this variable. Consequently, the indices l'_t, d'_t and w'_t, represent the level and seasonality in frequency-corrected demand. In expression (7), we express demand in terms of the sum of these indices minus the frequency correction.

$$D_t = l'_{t-1} + d'_{t-m_1} + w'_{t-m_2} - c(50 - F_{t-1})D_{t-1} + \phi' e'_{t-1} + \varepsilon'_t \tag{7}$$

$$e'_t = D_t - \left(l'_{t-1} + d'_{t-m_1} + w'_{t-m_2} - c(50 - F_{t-1})D_{t-1}\right) \tag{8}$$

$$l'_t = l'_{t-1} + \alpha' e'_t \tag{9}$$

$$d'_t = d'_{t-m_1} + \delta' e'_t \tag{10}$$

$$w'_t = w'_{t-m_2} + \omega' e'_t \tag{11}$$

We treated the modeling of frequency and demand as a joint model, and estimated all the parameters in one step by minimizing the sum of squared errors for the forecasts of frequency corrected-demand. We initially based estimation on one step-ahead forecast error, which is the standard approach for exponential smoothing. However, we obtained better results by following the TSO's Kalman Filter approach in estimating parameters separately for each lead time. In Section V, we report results for the model with this approach to estimation. For 3 minutes ahead, we obtained the following estimates: α=0.000, δ=0.301, ω=0.117, ϕ=0.300, α'=0.000, δ'=0.046, ω'=0.137 and ϕ'=0.997. For all exponential smoothing methods considered in this paper, we initialized the smoothed components using averages of the early observations.

IV. Modeling Frequency-Corrected Demand

In this section, we present a variety of univariate methods that we applied directly to frequency-corrected demand. These methods did not involve the modeling of frequency or demand. In addition to seasonal methods, we felt it was useful to evaluate non-seasonal methods, as our interest is in very short lead times, and because the Kalman Filter approach of Section III.A does not model seasonality. Prior to implementing the methods in this section, we applied the natural log transform in order to stabilize the variance and produce additive structure.

A. Random Walk

As a naïve benchmark, we used the value of the series at the forecast origin as the forecast for all lead times.

B. Simple AR Models

As additional simple benchmarks, we implemented an autoregressive (AR) model including only lags 1 to 5, and a second AR model including lags 1 to 5, as well as lags m_1 and m_2. In both models, all terms were significant at the 5% level.

C. SARMA and ARMA Models

We used the same form of double seasonal autoregressive moving average (SARMA) model considered in [14] for modeling a minute-by-minute series of demand. The model incorporates the product of three lag polynomials for both the AR and the MA parts of the model. The first polynomial is written in terms of the lag operator L, with the aim of capturing the variation in the level of the series. The second polynomial is written in terms of L^{m_1}, in order to model the intraday seasonality. The third polynomial is written in terms of L^{m_2}, with the aim of modeling the intraweek cycle. We considered lag polynomials up to order five, and found all terms to be significant at the 5% level. Estimation involved the standard approach of maximizing a Gaussian likelihood function.

As a benchmark, we also implemented a non-seasonal ARMA model for frequency-corrected demand. We included lags of up to order five for both the AR and MA terms.

D. ANN

We implemented a similar type of ANN to that used in [17] for half-hourly demand. This model is a univariate single hidden layer feedforward ANN with single output. We implemented a separate ANN for each lead time. To be consistent with our SARMA modeling, we considered a larger set of candidate inputs than used in [17]. For lead time h, the potential inputs consisted of the value of the output variable at the forecast origin and at lags: 1, 2, 3, 4, m_1-h, 2m_1-h, 3m_1-h, 4m_1-h, 5m_1-h, m_2-h, 2m_2-h, 3m_2-h, 4m_2-h and 5m_2-h. From these variables, we selected inputs using cross-validation with a hold-out sample consisting of the final 4 weeks of the estimation sample. This led to the following inputs: the output variable at the forecast origin and at lags: 1, 2, 3, 4, m_1-h, and m_2-h.

E. Non-seasonal Exponential Smoothing

We used damped additive trend exponential smoothing, which smoothes the level and trend of a series, and dampens the trend when forecasting (see, for example, [18]). The method does not model seasonality. We optimized the parameter values by the standard approach of minimizing the sum of squared in-sample 1 step-ahead forecast errors.

F. HWT Exponential Smoothing

We implemented the HWT exponential smoothing method of expressions (2)-(6) for frequency-corrected demand. We first optimized parameters by minimizing the sum of squared 1 step-ahead forecast error. However, we also optimized the model separately for each lead time, as this was the approach taken with the two methods of Section III. For 3 minute-ahead prediction, this led to α=0.001, δ=0.047, ω=0.140 and ϕ=0.997. In applying the HWT model to minute-by-minute demand [14], accuracy for 10-30 minutes-ahead improved when parameters were estimated based on 30 minute-ahead error. In this paper, we extend this idea by optimizing the parameters separately for each lead time.

G. SVD-Based Exponential Smoothing

For modeling half-hourly load, an exponential smoothing method has been introduced in [17] that involves singular value decomposition (SVD), which is a dimension reduction technique. For minute-by-minute data, the method proceeds by applying SVD to the data arranged as a ($w \times m_2$) matrix $\boldsymbol{Y}$, where w is the number of weeks of data. Each column of $\boldsymbol{Y}$ contains the observations for a particular minute of the week. Predicting the next row of $\boldsymbol{Y}$ amounts to forecasting the next week of the time series. SVD transforms $\boldsymbol{Y}$ to a new matrix $\boldsymbol{P}$, with columns ordered according to the extent to which they capture the variation in the columns of $\boldsymbol{Y}$. The approach is then to reduce the matrix $\boldsymbol{P}$ to fewer columns, in order to capture the main variation in $\boldsymbol{Y}$, whilst simplifying the forecasting task to one of predicting the next row of a reduced matrix.

If SVD is applied to a column-centered matrix, it is equivalent to principal component analysis. The columns of matrix $\boldsymbol{P}$ are then the principal components. The empirical study in [17] focused on British and French half-hourly load data, and used accuracy for a lead time of 1 half-hour to select the dimension of the reduced matrix. This led to dimensions of 31 and 38, respectively. In our study of minute-by minute data, we selected the dimension using 1 to 10 minute-ahead accuracy for a cross-validation sample consisting of the final 4 weeks of the estimation sample. For this, the optimal dimension of the reduced matrix was just 1 column. This suggests that, for very short lead times, the seasonality is adequately described by the average intraweek cycle, with the variation around this cycle captured by non-seasonal terms within the SVD-based exponential smoothing model. We optimized the parameter values by minimizing the sum of squared in-sample 1 step-ahead forecast errors.

V. Evaluation of Post-Sample Forecast Accuracy

For lead times from 1 to 10 minutes, we evaluated point forecast accuracy for each minute of the 9-week post-sample period, using mean absolute error (MAE), mean absolute percentage error, mean squared error, and mean squared percentage error. The rankings of the methods at each lead time were similar for all four measures, and so we present results for just the MAE. These results are given in Fig. 5.

Fig. 5 does not include the results for the exponential smoothing model of Section III.B, because the results for this joint model of frequency and demand were very similar to those of HWT exponential smoothing applied directly to frequency-corrected demand, which we described in Section IV.F. The joint modeling is probably not superior because it is difficult to predict the frequency time series beyond one step-ahead. Indeed, in separate analysis of frequency, we found that sophisticated models were only very slightly able to outperform the naïve random walk forecasting approach.

Fig. 5 shows the random walk, simple non-seasonal AR and simple seasonal AR methods performing relatively poorly beyond 1 minute-ahead. The results for the latter two methods were very similar, indicating that the seasonal terms did not benefit a simplistic AR modeling.

The ANN can be seen to have performed poorly for the early lead times, with accuracy improving with the lead time. This was also the case for a similar ANN implemented in [17] for half-hourly data, although the poor performance for the shortest lead time was less apparent in that study because it considered prediction up to 24 hours ahead, and the ANN was reasonably competitive beyond about 12 hours ahead. The ANN results suggest that there are not significant nonlinearities in the time series structure of the load data. We experimented with no differencing, with the use of the output at the forecast origin as the only input, and with the inclusion of all 15 potential inputs. However, this did not deliver improved post-sample forecast accuracy. It seems that the strong seasonality in the data leads to a relatively large network in terms of hidden units, but that this is overly complex for very short lead times. Of course, there are many other forms of ANN, and it could be that one of these would be more successful than our ANN for minute-by-minute frequency-corrected demand.

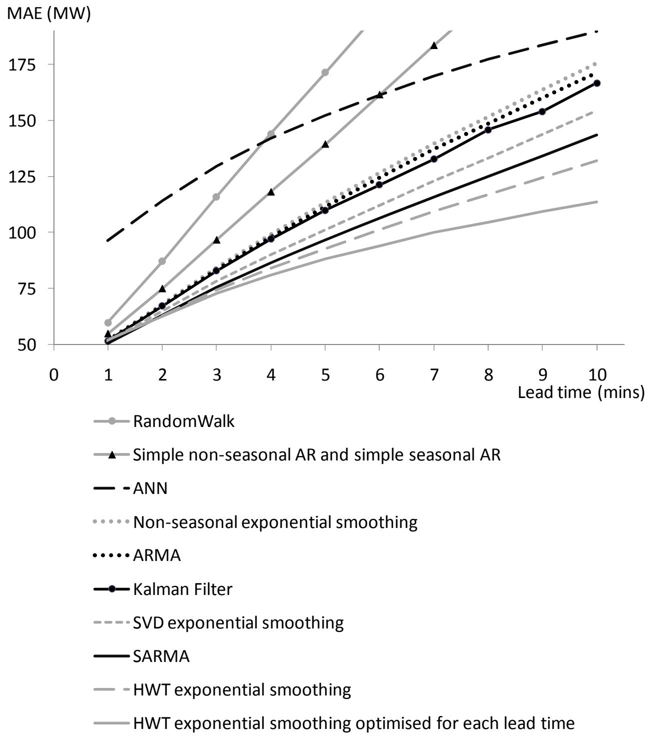

Fig. 5. Post-sample MAE for 1 to 10 minute-ahead prediction.

Fig. 5 shows fairly similar accuracy for the Kalman Filter approach, the non-seasonal exponential smoothing method, and the ARMA model. Unlike these three methods, the four best performing methods in Fig. 5 all model the seasonal cycles. The relative complexity of the SVD-based exponential smoothing approach would not seem to be justified, because the method is comfortably outperformed by the simpler HWT exponential smoothing method. This method was also more accurate than the SARMA model. Fig. 5 shows that the accuracy of the HWT method was improved by optimizing the method's parameters separately for each lead time.

Although we considered lead times from 1 to 10 minutes, the TSO's main interest is in 3 minutes. In Figs. 6 and 7, for three of the methods, we show how the MAE for this lead time varies across the minutes of the day. The figures are consistent with Fig. 5 in showing that the Kalman Filter was less accurate than the SARMA model, which in turn was less accurate than the HWT method. However, it is interesting to note from Fig. 7 that the HWT method is notably better than the SARMA model only around 6am and 8pm. For all three methods, the MAE is highest around these periods of the day, reflecting the difficulty of accurately modeling the series during the morning pick-up and the evening fall in demand.

As it is difficult to estimate the frequency correction factor c [4], we repeated our analysis for two alternative values, c=1% and 5%. With each of these values, we obtained the same rankings of methods that we have reported above for c=2.5%.

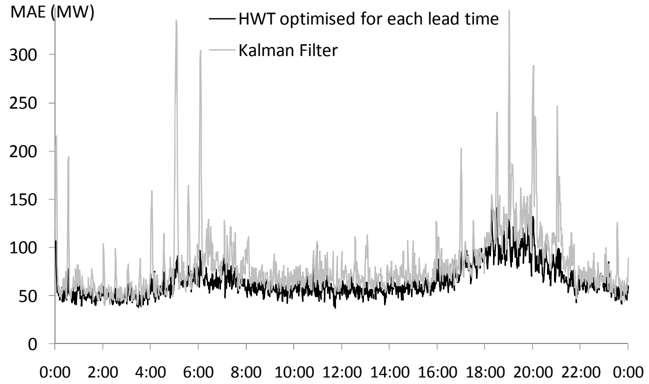

Fig. 6. MAE for the Kalman Filter approach and HWT exponential smoothing with parameters optimized for each lead time.

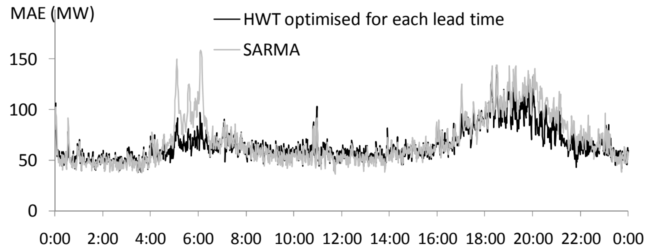

Fig. 7. MAE for the SARMA model and HWT exponential smoothing with parameters optimized for each lead time.

VI. Summary and Concluding Comments

In this paper, we first described how very short-term generation dispatch decisions should be based on predictions of frequency-corrected demand. We then compared the accuracy of methods for forecasting frequency-corrected demand. We considered methods that jointly modeled frequency and demand, but the best performing methods directly modeled frequency-corrected demand, and included terms for intraday and intraweek seasonal cyclicality. This is perhaps surprising, given that the lengths of these cycles are 1,440 and 10,080 minutes, respectively, while our focus was on lead times of 1 to 10 minutes. The most accurate method across all lead times was HWT exponential smoothing with parameters optimized separately for each lead time.

The focus in this paper has been on point forecasting, which reflects the main interest of the TSO. However, prediction intervals are also of potential interest. With regard to the best performing method in our study, HWT exponential smoothing, we note that the formulation of expressions (2) to (6) can be written as a single source of error state space model. This can be used as the basis for producing theoretical prediction intervals [19,Ch.6]. However, our simpler proposal is to use empirical prediction intervals, constructed from the distribution of historical forecast errors for the lead time of interest.

VII. References

[1] S.K. Pandey, S.R. Mohanty, and N. Kishor, "A literature survey on load-frequency control for conventional and distribution generation power systems," *Renewable and Sustainable Energy Reviews*, vol. 25, pp. 318-334, 2013.

[2] K. Liu, S. Subbarayan, R.R. Shoults, M.T. Manry, C. Kwan, F.L. Lewis, and J. Naccarino, "Comparison of very short-term load forecasting techniques," *IEEE Trans. Power Syst.*, vol. 11, pp. 877-882, 1996.

[3] H. Omara, "A methodology for determining the load frequency sensitivity," PhD Thesis, University of Manchester, 2012. https://www.escholar.manchester.ac.uk/uk-ac-man-scw:165577

[4] R. Pearmine, Y.H. Song, T.G. Williams, and A Chebbo, "Identification of a load-frequency characteristic for allocation of spinning reserves on the British electricity grid," *IEE Proceedings - Generation, Transmission and Distribution*, vol. 153, pp. 633-638, 2006.

[5] G.J. Berg, "System and load behaviour following loss of generation. Experimental results and evaluation," *Proceedings of the Institution of Electrical Engineers*, vol. 119, pp. 1483-1486, 1972.

[6] C. Concordia, and S. Ihara, "Load representation in power system stability studies," *IEEE Trans. on Power Apparatus and Systems*, vol. 4, pp. 969-977, 1982.

[7] E. Welfonder, H. Weber, and B. Hall, "Investigations of the frequency and voltage dependence of load part systems using a digital self-acting measuring and identification system," *IEEE Trans. on Power Systems*, vol. 4, pp. 19-25, 1989.

[8] W. Charytoniuk, and M.-S. Chen, "Very short-term load forecasting using artificial neural networks," *IEEE Trans. Power Syst.*, vol. 15, pp. 263-268, 2000.

[9] D.J. Trudnowski, W.L. McReynolds, and J.M. Johnson, "Real-time very short-term load prediction for power-system automatic generation control," *IEEE Trans. on Control Sys. Technol.*, vol. 9, pp. 254-260, 2001.

[10] R.E. Abdel-Aal, "Short-term hourly load forecasting using abductive networks," *IEEE Trans. Power Syst.*, vol. 19, pp. 164-173, 2004.

[11] H.Y. Yang, H. Ye, G. Wang, J. Khan, and T. Hu, "Fuzzy neural very short-term load forecasting based on chaotic dynamics reconstruction," *Chaos, Solitons and Fractals*, vol. 29, pp. 462-469, 2006.

[12] C. Guan, P.B. Luh, L.D. Michel, Y. Wang, and P.B. Friedland, "Very short-term load forecasting: wavelet neural networks with data pre-filtering," *IEEE Trans. Power Syst.*, vol. 28, pp. 30-41, 2014.

[13] C. Guan, P.B. Luh, L.D. Michel, and Z. Chi, "Hybrid Kalman Filters for very short-term load forecasting and prediction interval estimation," *IEEE Trans. Power Syst.*, vol. 28, pp. 3806-3817, 2013.

[14] J.W. Taylor, "An evaluation of methods for very short-term load forecasting using minute-by-minute British data," *Int. J. Forecast.*, vol. 24, pp. 645-658, 2008.

[15] S. Arora, and J.W. Taylor, "Short-term forecasting of anomalous load using rule-based triple seasonal methods," *IEEE Trans. Power Syst.*, vol. 28, pp. 3235-3242, 2013.

[16] J. Durbin, and S.J. Koopman, *Time Series Analysis by State Space Methods*. Oxford, UK: Oxford University Press, 2001.

[17] J.W. Taylor, "Short-term load forecasting with exponentially weighted methods," *IEEE Trans. Power Syst.*, vol. 27, pp. 458-464, 2012.

[18] E.S. Gardner Jr., "Exponential smoothing: the state of the art - Part II," *International Journal of Forecasting*, vol. 22, pp. 637-666, 2006.

[19] R.J. Hyndman, A.B. Koehler, J.K. Ord, and R.D. Snyder, *Forecasting with Exponential Smoothing: the State Space Approach*. Berlin, Heidelberg, Germany: Springer-Verlag, 2008.

James W. Taylor is Professor of Decision Science at the Saïd Business School, University of Oxford. His research interests include energy forecasting, exponential smoothing and probabilistic forecasting.