Ocean Engineering 36 (2009) 348–356

Contents lists available at ScienceDirect

Ocean Engineering

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

Forecasting ocean wave energy: Tests of time-series models

Gordon Reikard *

Sprint-Nextel, 6180 Sprint Parkway, Overland Park, KS 66211, USA

ARTICLE INFO

Article history:
Received 14 June 2008
Accepted 3 January 2009
Available online 20 January 2009

Keywords:
Forecasting
Time-series models
Wave energy
Nonlinear variability

ABSTRACT

This paper evaluates the ability of time-series models to predict the energy from ocean waves. Data sets from four Pacific Ocean sites are analyzed. The energy flux is found to exhibit nonlinear variability. The probability distribution has heavy tails, while the fractal dimension is non-integer. This argues for using nonlinear models. The primary technique used here is a time-varying parameter regression in logs. The time-varying regression is estimated using both a Kalman filter and a sliding window, with various window widths. The sliding window method is found to be preferable. A second approach is to combine neural networks with time-varying regressions, in a hybrid model. Both of these methods are tested on the flux itself. Time-varying regressions are also used to forecast the wave height and wave period separately, and combine the forecasts to predict the flux. Forecasting experiments are run at an hourly frequency over horizons of 1–4 h, and at a daily frequency over 1–3 days. All the models are found to improve relative to a random walk. In the hourly data sets, forecasting the components separately achieves the best results in three out of four cases. In daily data sets, the hybrid and regression models yield similar outcomes. Because of the intrinsic variability of the data, the forecast error is fairly high, comparable to the errors found in other forms of alternative energy, such as wind and solar.

1. Introduction

Recent technological advances have made it possible to capture the energy of ocean waves for electricity generation at a competitive cost. Wave farms are being developed in the United States, Portugal, England, and Scotland, while wave-powered devices are being tested in other locations. Examples of this technology include the Pelamis Wave Energy Converter, which was successfully tested at the European Marine Energy Center in 2004, and the Energen Wave Power device. Utilities need to forecast power output for operational planning and short-term trading. The horizons involved in short-term management of power grids range from as little as a few hours to as long as several days. The existing work on wave forecasting falls into two broad categories, large-scale energy balance models and time-series methods. This paper uses the latter approach, evaluating the ability of several time-series models to predict the wave energy flux.

Let Y_t denote the power, in watts per meter (W/m) of crest length, g denote the acceleration caused by gravity (9.086 m/s^2), ρ denote the density of seawater (1025 kg/m^3), H denote the wave height in meters, and T denote the wave period in seconds. To distinguish between stochastic series and constants, H and T are written with time subscripts. The wave energy flux is measured as

$$Y_t = (g^2\rho/64\pi)H_t^2T_t \approx 0.492(H_t^2T_t) \tag{1}$$

Although the actual power derived from ocean waves is considerably less than the flux, this measure is a straightforward to calculate from the existing databases, and serves as a convenient unit for analysis.

The organization of this paper is as follows. Section 2 consists of an overview of the data and its statistical properties. Section 3 presents the forecasting models. The main emphasis is on time-varying parameter regressions, and hybrid models combining regressions with neural networks. Section 4 consists of a series of forecasting experiments. Two sets of tests are run, for the hourly and daily data. Section 5 concludes.

2. The data

The data sets, summarized in Table 1, are from four locations in the Pacific Ocean, and are available at a basic resolution of 1 h. Two of these, Point Loma and Half Moon Bay, CA, are coastal sites. The other two, Kauai, HI, and Aberdeen, WA are deep ocean sites several hundred miles from the coastlines. The time series consist of the wave height, the wave period, and in three cases wind speed, in m/s. The time span ranges from 20 to 36 months. The source is the National Oceanographic and Aeronautics Administration's National Data Buoy Center website

* Tel.: +1 913 315 6553.
E-mail address: Reikarsen@msn.com

0029-8018/$ - see front matter © 2009 Elsevier Ltd. All rights reserved.
doi:10.1016/j.oceaneng.2009.01.003

Table 1
Statistical properties of the data (all data are available at a 1-h resolution).

Database and series	Time span	Observations	Order of integration	Characteristic exponent (α)	Fractal codimension
Kauai, HI (170 NM west)	January 1, 2005–December 31, 2007	26,280			
Energy flux			0.43	1.18 ($\pm$0.008)	0.059 ($\pm$0.016)
Wave height			0.41	1.44 ($\pm$0.022)	0.048 ($\pm$0.045)
Average period			0.39	1.96 ($\pm$0.034)	0.012 ($\pm$0.069)
Wind speed			0.18	1.73 ($\pm$0.067)	0.051 ($\pm$0.141)
Point Loma, CA	March 26, 2007–December 7, 2007	14,934			
Energy flux			0.25	1.48 ($\pm$0.035)	0.061 ($\pm$0.071)
Wave height			0.22	1.75 ($\pm$0.063)	0.052 ($\pm$0.127)
Dominant period			0.32	1.92 ($\pm$0.074)	0.008 ($\pm$0.156)
Aberdeen, WA (315 NM west)	April 7, 2006–December 31, 2007	15,149			
Energy flux			0.39	1.16 ($\pm$0.008)	0.055 ($\pm$0.016)
Wave height			0.42	1.31 ($\pm$0.017)	0.049 ($\pm$0.054)
Dominant period			0.26	1.92 ($\pm$0.037)	0.000 ($\pm$0.074)
Wind speed			0.17	1.39 ($\pm$0.012)	0.089 ($\pm$0.025)
Half Moon Bay, CA	January 1, 2006–December 31, 2007	17,520			
Energy flux			0.29	1.44 ($\pm$0.028)	0.121 ($\pm$0.056)
Wave height			0.28	1.76 ($\pm$0.027)	0.078 ($\pm$0.054)
Dominant period			0.27	1.97 ($\pm$0.052)	0.000 ($\pm$0.141)
Wind speed			0.17	1.35 ($\pm$0.010)	0.081 ($\pm$0.021)

Source: National Data Buoy Center, National Oceanographic and Aeronautics Administration. Website: http://seaboard.ndbc.noaa.gov

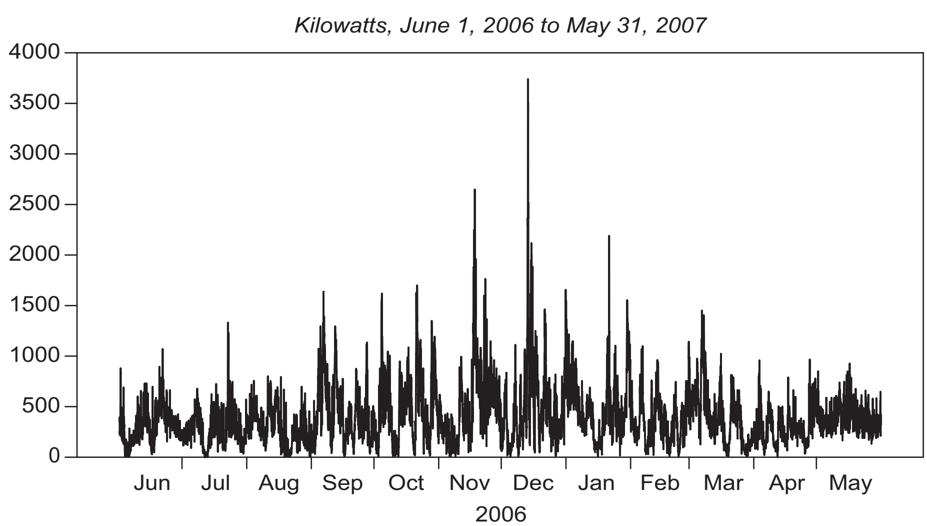

Fig. 1. The flux, Kauai.

(http://seaboard.ndbc.noaa.gov). In all the databases, there were missing dates and values. In some instances, dates were available with no values; in other cases, the dates themselves were missing. The missing values were interpolated in two stages. First, the date fields were used to create a continuous calendar. Second, the hourly wave height and period values were interpolated using the same regressions as in the forecasting models. Other interpolation methods yielded similar results (see on this issue Londhe, 2008). Once the continuous hourly data sets were created, the values were averaged to a one-day frequency in order to be able to run forecasting tests for daily values.

Figs. 1–4 show the flux, in kilowatts, at the hourly resolution, for selected periods. Table 1 presents three measures of its statistical properties, the order of integration, the characteristic exponent, and fractal codimension. These were estimated using multiple scaling, a method derived from fractal theory (Schertzer et al., 1997). All the tests were run on logs. The order of integration, denoted $I(\xi)$, is a measure of the dependence between time points at distant intervals. The flux is fractionally integrated, with coefficients lying in the range 0.25–0.43. Fractionally integrated series in the range $0<\xi<0.5$ do not trend over long horizons, but are characterized by long memory: the dependence between time points persists over long separation intervals (Beran, 1994; Baillie, 1996). The flux shows greater persistence in deep water locations; there is less serial dependence at the coastal sites. The wave height and period both also exhibit considerable persistence. Wind speed, however, shows significantly lower coefficients of integration.

The characteristic exponent, denoted α, measures the degree of tail thickness in the probability distribution. The special cases of $\alpha=2$ and 1 correspond to the Gaussian and Cauchy distributions. When $1<\alpha<2$, the distribution has heavier tails than in the

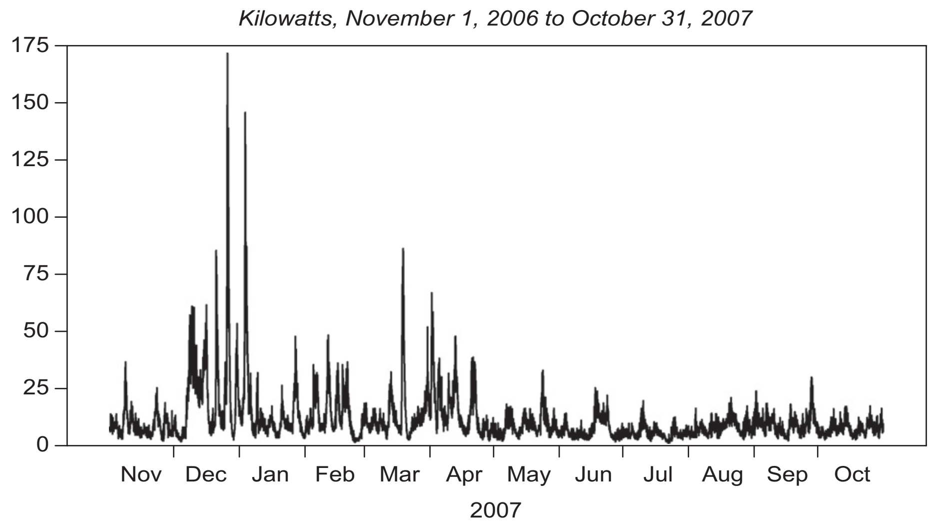

Fig. 2. The flux, Point Loma.

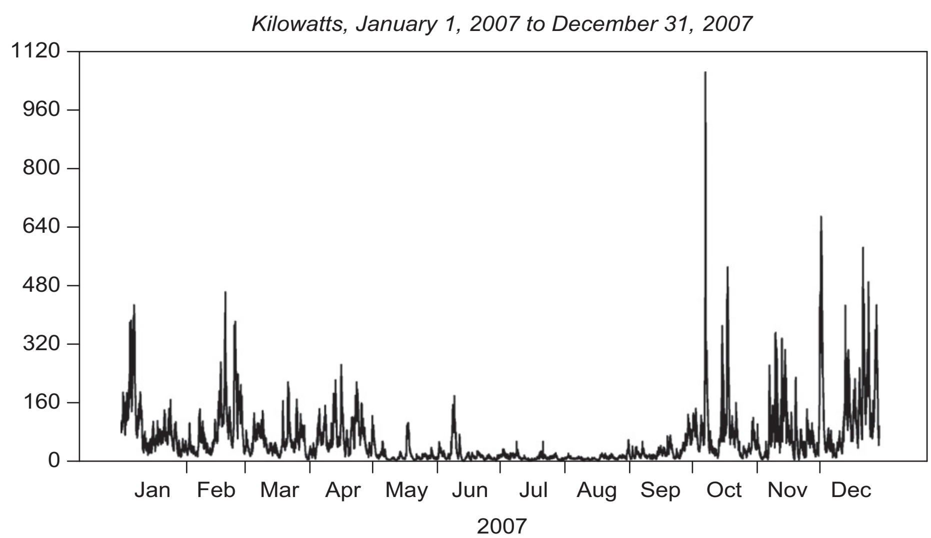

Fig. 3. The flux, Aberdeen.

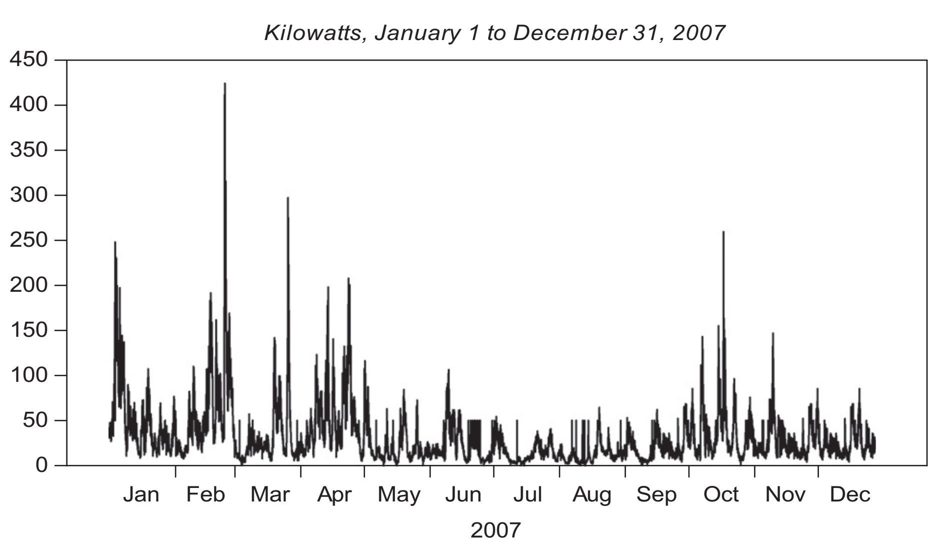

Fig. 4. The flux, Half Moon Bay.

standard normal. The flux exhibits heavy tails, with α estimated at 1.16–1.18 for the deep water sites, and 1.44–1.48 for the coastal sites. This is consistent with the results for wave height in Song (2006). The wind speed also exhibits heavy tails, with α estimated in the range 1.35–1.74.

The fractal codimension is a measure of the probability of events outside of a given threshold. Let Y_t be a time series, and let Υ_t denote the absolute log-difference of Y_t, i.e., $\Upsilon_t = |\ln Y_t - \ln Y_{t-1}|$. Let L be a characteristic length, such as a rate of change and let N be the number of observations. A probabilistic measure of dimensionality can be written as a power-law relationship:

$$\breve{D} \approx \ln[N(\Upsilon_t > L)/N(\Upsilon_t)]/\ln L; \quad [N(\Upsilon_t > L)/N(\Upsilon_t)] \approx L^{\breve{D}} \tag{2}$$

The fractal codimension is the difference between the embedding dimension of the space (D) and the fractal dimension ($\breve{D}$): $C_1 = D - \breve{D}$. For all $C_1 \neq 0$, the process is said to be fractal. When C_1 tends toward zero, there are fewer extreme events. As C_1 increases, there is a higher probability of fluctuations outside the threshold. There is strong evidence of fractality in the flux. In three cases, the codimension lies in the range of 0.055–0.061, with the estimates for the wave height marginally lower. The Point Loma data shows extremely high values, with the codimension estimated as 0.081 for wind speed, 0.077 for wave height, and 0.121 for the flux. The finding of fractality does not indicate the presence of deterministic attractors, but rather is accounted for by the multiplicative interactions in the flux, and by wave–wave and wind–wave interactions.

3. The forecasting models

The literature on wave forecasting by statistical methods originates with Sverdrup and Munk (1947). Large-scale frequency-domain models have been in use since the 1960s. In fully developed sea states, the Pierson and Moskowitz (1964) wave spectrum has often been applied; for partially developed sea states, the JONSWAP (Joint North Sea Wave Project) spectrum has been used (Hasselmann et al., 1976, 1980). These models relate the propagation and change in the wave spectrum to the energy input by wind, nonlinear energy transfer by interactions among waves, and dissipation. The wind input term includes the resonant interaction between waves and turbulent pressure patterns in the atmosphere, and the feedback between waves and induced pressure patterns (Phillips, 1957, 1958; Miles, 1957; see also Snyder et al., 1981). Nonlinear interactions in the wave field involve the transfer of energy from the dominant region, near the spectral peak, to higher and lower frequencies (Hasselmann et al., 1985). Recent advances in this field include third-generation simulation models, such as the WAMDI Group (1988) model, which represent state-of-the art knowledge of the physics of wave evolution (Janssen, 1991; Komen et al., 1994).

A second branch of the literature, originating for the most part in the 1990s, has used time domain methods. The most popular method has been neural networks. Representative works include Deo and Naidu (1998), Deo et al. (2001), Tsai et al. (2002), Deo and Jagdale (2003), Makarynskyy (2004), Londhe and Panchang (2006), Jain and Deo (2007), Tseng et al. (2007) and Zamani et al. (2008). Other methods have also been essayed. Ho and Yim (2006) propose transfer functions. Roulston et al. (2005) adopt a probabilistic approach. Malmberg et al. (2005) use Kalman filters, while Gaur and Deo (2008) adapt genetic programming for forecasting waves. In the last few years, there have been several important papers on forecasting time series with fractal properties, mainly in financial economics (Calvet and Fisher, 2001, 2002, 2004; Lux, 2008). These argue that fractal series can be modeled as multiplicative combinations of states, and recommend state transition or Markov-switching models. More recently, however, Granger (2008) has demonstrated that any nonlinear model, including state transition paradigms, can be approximated using regressions with time-varying coefficients.

The basic time-varying parameter model is specified as follows. Let Y_t = the wave energy flux, let ln denote natural logs, let ω denote a coefficient, let the t-subscript denote time variation, and let ε_t denote the residual:

$$\ln Y_t = \omega_{0t} + \omega_{1t}\ln Y_{t-1} + \omega_{2t}\ln Y_{t-2} + \omega_{3t}\ln Y_{t-3} + \omega_{4t}\ln Y_{t-4} + \varepsilon_t, \quad \varepsilon_t \sim P(0, \sigma_t^2) \tag{3}$$

where P is the probability distribution, and σ_t^2 is the residual variance. Experiments were run with various lag lengths. For the hourly data, lags 1 through 4 were statistically significant. At the daily frequency, the number of lags is lower. A simple enhancement is to include wind speed (W_t) on the right-hand side:

$$\ln Y_t = \omega_{0t} + \omega_{1t}\ln Y_{t-1} + \omega_{2t}\ln Y_{t-2} + \omega_{3t}\ln Y_{t-3} + \omega_{4t}\ln Y_{t-4} + \omega_{5t}\ln W_t + \varepsilon_t \tag{4}$$

Since wind can have zero values, it is transformed to logs by adding a small positive constant. Forecasting models with exogenous forcing factors can be specified in two ways. Either the flux can be forecasted using lags of wind speed, so that the term on the RHS of Eq. (4) becomes $\ln W_{t-i}$. Or the regressions can be estimated using the contemporaneous value of the forcing factor, in which case $\ln Y_{t+i}$ is predicted using $\ln W_{t+i}$. In this case, wind speed must also be forecasted. The latter method was used here. The model for wind speed is from Reikard (2008):

$$\ln W_t = \omega_{0t} + \omega_{1t}\ln W_{t-1} + \omega_{2t}\ln W_{t-2} + \omega_{3t}\ln W_{t-3} + \omega_{24t}\ln W_{t-24} + \omega_{Kt}\ln K_t + \varepsilon_t \tag{5}$$

where $\ln K_t$ is ambient temperature, in log Kelvins.

Wave energy data can be non-stationary over particular horizons. Because of seasonal cycles, wave height may trend upward or downward over periods of several months. When time series are non-stationary, a common procedure is to first-difference, as in the well-known Autoregressive, Integrated, Moving Average (ARIMA) class of models (Box and Jenkins, 1970). While the notation associated with this class of models is different, they are essentially regressions on differences, with additional terms for moving averages or causal inputs. In essence, the issue here is whether to estimate regressions on levels or differences. Both methods are tested.

A further possibility is to forecast the components of wave energy separately and combine the forecasts to predict the flux. The wave height can depend on the period, so this term is included in the equation:

$$\ln H_t^2 = \omega_{0t} + \omega_{1t}\ln H_{t-1}^2 + \omega_{2t}\ln H_{t-2}^2 + \omega_{3t}\ln H_{t-3}^2 + \omega_{4t}\ln H_{t-4}^2 + \omega_{5t}\ln T_t + \omega_{6t}\ln W_t + \varepsilon_t \tag{6}$$

The period itself is predicted by a regression on lags:

$$\ln T_t = \omega_{0t} + \omega_{1t}\ln T_{t-1} + \omega_{2t}\ln T_{t-2} + \omega_{3t}\ln T_{t-3} + \omega_{4t}\ln T_{t-4} + \varepsilon_t \tag{7}$$

There is as yet little evidence as to how well fractal series can be approximated by non-parametric models, such as neural networks. However, neural nets have been widely used to model nonlinear data sets. Consequently, while the main emphasis here is on the time-varying coefficient models, neural nets are tested as a possible alternative. The system architecture used here consists of a multilayer perceptron, trained using a backpropagation algorithm. The net can be trained by epoch (a forward and backward pass over all the observations in the sample) or by example (a pass over individual observations). In iterative forecasting exercises, the input and bias weights can be retained from the previous period or can be restarted at each time point.

A final idea is to combine the neural network and time-varying regressions, in a hybrid model. Hybrid models have recently attracted some interest because they can combine the best features of different techniques. References include Nelson et al. (1999), Hibbert et al. (2000), Zhang (2003), Zhang and Qi (2005), Hibon and Evgeniou (2005) and Aburto and Weber (2007). Typically, in these hybrid models, an initial regression is estimated, and the residuals are then processed using a neural net. The procedure here is actually the opposite. First, the net was trained, then the memory vector from the net was used as an input in a time-varying regression, and this equation was forecasted. In effect, the output of the neural net is used as a causal factor. The regression then adjusts the time-varying weights associated with the memory vector and other causal factors at each observation.

4. Forecasting tests

Two sets of forecasting tests were run, for the hourly data, over horizons from 1–4 h, and for the daily data, at horizons of 1–3 days. The forecasts were iterative. In other words, the models were initially estimated over an initial set of starting values or training sample, then models were forecasted, and then at the next step, the models were re-estimated, continuing through the entire data set. All the predicted values are true out-of-sample forecasts, in that they use only data prior to the start of the forecast horizon. The reported error for horizons beyond one period is the forecast for that horizon only; intervening values are omitted. In other words, the error at t+2 omits the error at t+1. The criterion used to evaluate the forecasts is the mean absolute percent error, defined as the average value of [Abs (log of forecast−log of actual)]. As recommended in Theil (1971), the forecast errors are evaluated relative to a random walk, i.e., setting the forecast equal to the most recently observed value. The tables report both the actual error and the ratio of the error to the error from the random walk.

The first specification issue is estimating the time-varying coefficients. When an unrestricted Kalman filter (Kalman, 1960) was used, the coefficients were observed to fluctuate excessively, causing forecast accuracy to deteriorate. The forecast errors from the Kalman filter were only slightly lower than in the random walk. Given this, the regressions were estimated using a sliding window. Experiments were run with different window widths, and the span that yielded the smallest forecast errors was used. At short window widths, the results were similar to the Kalman filter. In the hourly tests, as the window width extended to more than 960 h, the forecasts became too inertial. The smallest errors were found in a range of 400–800 h. The tests are reported for a window width of 600 h for the Kauai and Aberdeen data, and 480 h for the other two data sets. In the daily tests, the optimal window width was in the range of 60 days. When shorter window widths were used, the coefficients again became too volatile.

The second specification issue is the neural net. Experiments were run with different numbers of hidden layers, and learning gradients. In the daily data, one direct connection and one hidden layer produced the best forecasts. In the hourly data the preferred configuration was one direct connection and three hidden layers. The inputs were lags of the flux, and wind speed. The bias and input weights were allowed to change at each point in the iterative forecasts. When the weights were allowed to change at each interval, training by example caused the output to be excessively variable. Consequently, the nets were trained by epoch. The optimal training sample used for the neural nets was significantly shorter than for the regressions. When the training samples were set to the same lengths as the window widths, the net forecasts were too inertial. Training samples, on the order of 144 h and 2 weeks were used in the hourly and daily tests. The initial neural net forecasts generally yielded somewhat smaller errors than in the Kalman filter tests, but larger errors than in the sliding window regressions. By comparison, in the hybrid models, the errors were always lower than in the initial pass with the neural net. Consequently, the results from the hybrid rather than the net are reported.

The code for all the models was written using the Regression Analysis of Time Series program owned by Estima, Inc., a software company in Illinois (www.estima.com). Additional experiments with neural nets were run using Neural Solutions software (www.nd.com).

In the hourly tests, Model 1 is the time-varying regression (Eq. (3)). Model 2 includes wind speed (Eq. (4)). This is omitted in Point Loma data, where wind speed is unavailable. Model 3 is the regression in differences. Model 4 is the hybrid. Model 5 combines the forecasts for the components (Eqs. (6) and (7)). For reference, the errors from the best forecasts for the components are also reported. Table 2 reports the results from the hourly tests. In the Kauai data set, all the models reduce the error by about 5 percent, at all horizons. Forecasting the two components separately and the regression in differences are marginally better than the alternatives. In Point Loma, forecasting the two components separately does slightly better than the other models, with an improvement of about 8 percent relative to the random walk. The hybrid is slightly better than the regressions. In Aberdeen, the models on average do about 6 percent better than the random walk at the 1-h horizon, although as the horizon extends, the improvement diminishes. Forecasting the components separately achieves the best results, with an improvement of 7 percent. The hybrid and the regression with wind speed are very close. In Half Moon Bay, the regression in differences does better than the other models, probably because this model is better able to deal with localized trending in the data.

Figs. 5–8 present the absolute percent error from the best model at the 1-h horizon. In each case, the average error is raised by the presence of frequent outliers. In some data sets these occur intermittently. In other instances, they are pervasive.

Table 3 presents the results for the tests at a daily resolution. In the Kauai data set, forecasting the components separately does better than fitting models to the flux, but only at the 1-day horizon. However, at the 2–3-day horizons, the hybrid achieves the best results. In Point Loma, the hybrid does better at 1- and 3-day horizons, while at the 2-day horizon three models do about equally well. In Aberdeen, the hybrid achieves the consistently best results, although at the 1-day horizon the regression with wind speed is similar. Forecasting the components separately also does better at the 1-day horizons, but there is no improvement thereafter. In the Half Moon Bay data set, three models achieve comparable results, the regression, the separate forecast of the components and the hybrid. The regression in differences did not predict well with any of the daily data sets.

5. Conclusions

Because the flux is the product of two stochastic series, the forecast error at high frequencies is necessarily high. At the hourly frequency, it is often preferable to forecast the wave height and period separately. This procedure yielded the best results in three-fourths of the data sets. Among the paradigms that were estimated for the flux directly, the experiments at the hourly frequency favored the time-varying parameter regression and

Table 2
Comparison of the mean absolute percent forecast error, hourly data.

Model	Forecast horizon							
	1 h		2 h		3 h		4 h	
	Error	Ratio	Error	Ratio	Error	Ratio	Error	Ratio
Kauai								
Random walk	0.2679		0.3211		0.3622		0.3963	
Regression: logs, stochastic coefficients	0.2556	0.9541	0.3060	0.9530	0.3420	0.9442	0.3755	0.9475
Regression with wind speed	0.2574	0.9608	0.3064	0.9542	0.3453	0.9533	0.3767	0.9505
Regression in differences	0.2552	0.9526	0.3058	0.9524	0.3412	0.9420	0.3749	0.9460
Hybrid (regression, neural network)	0.2563	0.9567	0.3059	0.9527	0.3442	0.9503	0.3766	0.9503
Components forecasted separately	0.2548	0.9511	0.3049	0.9495	0.3429	0.9467	0.3743	0.9445
Error for the wave height squared	0.2269		0.2796		0.3193		0.3519	
Error for the wave period	0.0889		0.0938		0.0991		0.1046	
Point Loma								
Random walk	0.1739		0.1896		0.2039		0.2185	
Regression: logs, stochastic coefficients	0.1609	0.9252	0.1758	0.9272	0.1903	0.9333	0.2049	0.9378
Regression in differences	0.1604	0.9224	0.1751	0.9235	0.1891	0.9274	0.2040	0.9336
Hybrid (regression, neural network)	0.1605	0.9229	0.1752	0.9241	0.1884	0.9240	0.2024	0.9263
Components forecasted separately	0.1590	0.9143	0.1741	0.9182	0.1887	0.9255	0.2039	0.9332
Error for the wave height squared	0.1063		0.1262		0.1461		0.1646	
Error for the wave period	0.1045		0.1098		0.1211		0.1211	
Aberdeen								
Random walk	0.1819		0.2071		0.2345		0.2631	
Regression: logs, stochastic coefficients	0.1722	0.9467	0.1995	0.9633	0.2273	0.9693	0.2566	0.9753
Regression with wind speed	0.1709	0.9395	0.1965	0.9488	0.2235	0.9531	0.2509	0.9536
Regression in differences	0.1722	0.9467	0.1998	0.9648	0.2282	0.9731	0.2573	0.9780
Hybrid (regression, neural network)	0.1716	0.9434	0.1976	0.9541	0.2247	0.9582	0.2513	0.9552
Components forecasted separately	0.1689	0.9285	0.1938	0.9358	0.2198	0.9373	0.2482	0.9434
Error for the wave height squared	0.1149		0.1413		0.1696		0.1993	
Error for the wave period	0.1025		0.1114		0.1198		0.1291	
Half Moon Bay								
Random walk	0.1989		0.2256		0.2517		0.2776	
Regression: logs, stochastic coefficients	0.1955	0.9829	0.2224	0.9858	0.2491	0.9897	0.2754	0.9921
Regression with wind speed	0.1958	0.9844	0.2221	0.9845	0.2490	0.9893	0.2755	0.9924
Regression in differences	0.1894	0.9522	0.2158	0.9566	0.2419	0.9611	0.2681	0.9658
Hybrid (regression, neural network)	0.1947	0.9789	0.2197	0.9738	0.2475	0.9833	0.2743	0.9881
Components forecasted separately	0.1911	0.9608	0.2182	0.9672	0.2449	0.9730	0.2719	0.9795
Error for the wave height squared	0.1551		0.1841		0.2121		0.2408	
Error for the wave period	0.0889		0.0877		0.1058		0.1135	

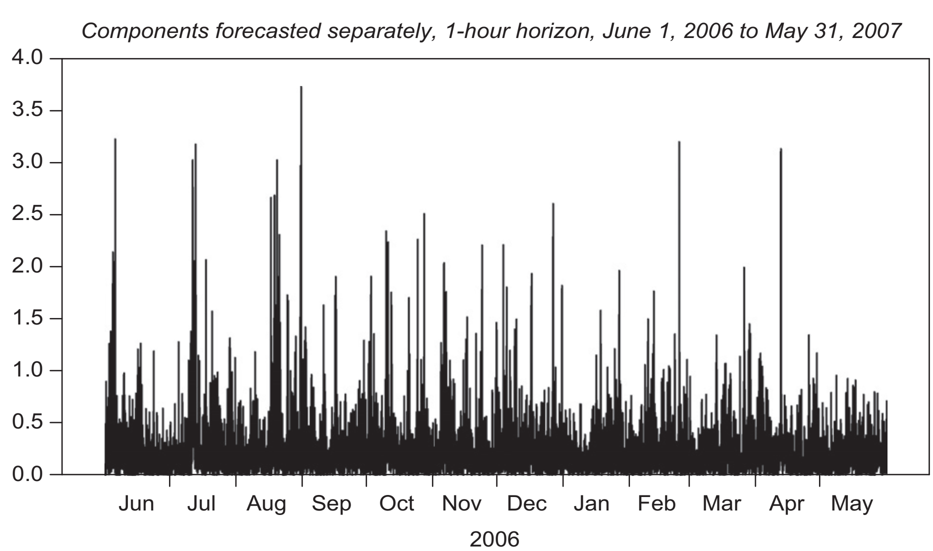

Fig. 5. The absolute forecast error, Kauai.

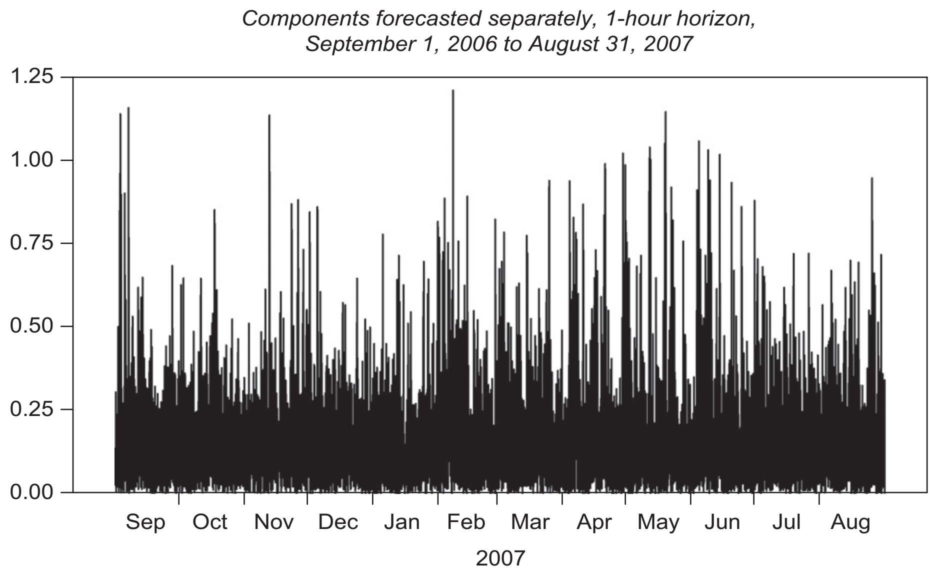

Fig. 6. The absolute percent error, Point Loma.

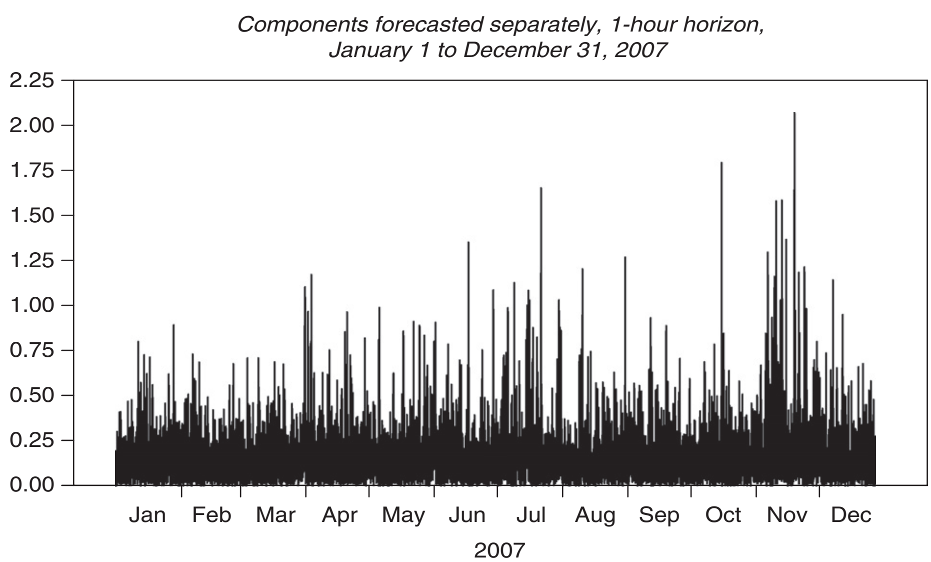

Fig. 7. The absolute percent error, Aberdeen.

hybrid models. Including the wind speed generally did not improve predictive accuracy. This outcome is at first sight difficult to reconcile with the finding that wind speed is statistically significant in the regressions. However, much of the information in wind speed that is relevant to prediction is already captured by the lags of the flux. Further, there is no improvement from including wind speed at the longer horizons, in part because the accuracy of the wind speed forecasts degrades. Whether to run the regressions in levels or differences depends largely on whether the data exhibits localized trending. In data sets where the data trends over short-term horizons, due to seasonal variations, differencing can do slightly better than regression on levels.

The tests generally support the time-varying parameter methodology. In this respect, the most important specification issue in the time-varying regressions was the width of the window. There is no hard-and-fast rule for the optimal window width. Instead, this must be determined empirically.

The favorable results from the regressions should not be read as rejecting neural nets as a forecasting methodology. The findings for the net are subject to the caveat that only one type of system architecture was used. It is possible that with other architectures such as genetic optimization, the results in the preliminary experiments would have been preferable. The ability of neural nets to predict fractal series is as yet a largely unexplored area, and should be investigated. In this respect, an interesting finding is that the length of the training sample had a significant impact on the accuracy of the neural net forecasts. Again, the optimal length of the training sample needs to be determined by empirical testing. The hybrid models, however, generally did well, particularly at the daily resolution. In essence, the memory vector from the neural net can contribute additional information that

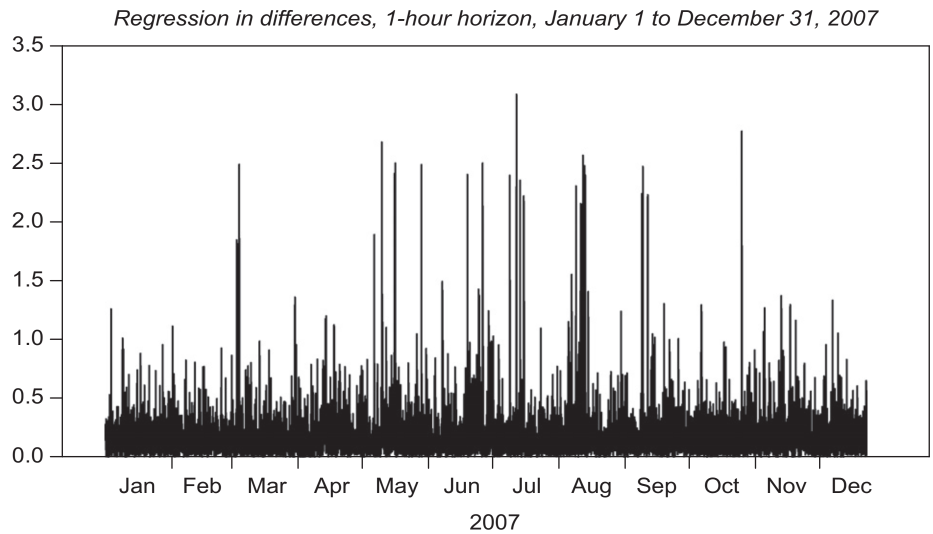

Fig. 8. The absolute percent error, Half Moon Bay.

Table 3
Comparison of the mean absolute percent forecast error, daily data.

Model	Forecast horizon					
	1 day		2 days		3 days	
	Error	Ratio	Error	Ratio	Error	Ratio
Kauai						
Random walk	0.3251		0.4602		0.5134	
Regression: logs, stochastic coefficients	0.3097	0.9526	0.4218	0.9166	0.4586	0.8933
Regression with wind speed	0.3091	0.9508	0.4221	0.9172	0.4589	0.8938
Regression in differences	0.3132	0.9634	0.4325	0.9398	0.4703	0.9160
Hybrid (regression, neural network)	0.3112	0.9572	0.4215	0.9159	0.4551	0.8864
Components forecasted separately	0.3054	0.9394	0.4234	0.9200	0.4654	0.9065
Error for the wave height squared	0.2583		0.3688		0.4018	
Error for the wave period	0.1153		0.1516		0.1636	
Point Loma						
Random walk	0.3164		0.4752		0.5278	
Regression: logs, stochastic coefficients	0.2857	0.9030	0.4104	0.8636	0.4349	0.8240
Regression in differences	0.2955	0.9339	0.4311	0.9072	0.4743	0.8986
Hybrid (regression, neural network)	0.2815	0.8897	0.4101	0.8630	0.4289	0.8126
Components forecasted separately	0.2887	0.9125	0.4108	0.8645	0.4351	0.8244
Error for the wave height squared	0.2638		0.3615		0.3911	
Error for the wave period	0.1105		0.1481		0.1571	
Aberdeen						
Random walk	0.4436		0.6134		0.6553	
Regression: logs, stochastic coefficients	0.4192	0.9450	0.5612	0.9149	0.5974	0.9116
Regression with wind speed	0.4052	0.9134	0.5758	0.9387	0.6013	0.9176
Regression in differences	0.4097	0.9236	0.5730	0.9341	0.6057	0.9243
Hybrid (regression, neural network)	0.4041	0.9110	0.5592	0.9116	0.5925	0.9042
Components forecasted separately	0.4192	0.9450	0.5643	0.9200	0.6025	0.9194
Error for the wave height squared	0.3839		0.5066		0.5539	
Error for the wave period	0.1347		0.1661		0.1697	
Half Moon Bay						
Random walk	0.4408		0.6241		0.6801	
Regression: logs, stochastic coefficients	0.4198	0.9524	0.5968	0.9563	0.5972	0.8781
Regression with wind speed	0.4231	0.9598	0.5972	0.9569	0.6259	0.9203
ARIMA: logs, stochastic coefficients	0.4282	0.9714	0.6012	0.9633	0.6502	0.9560
Hybrid (regression, neural network)	0.4124	0.9356	0.5916	0.9479	0.5971	0.8780
Components forecasted separately	0.4163	0.9444	0.5918	0.9482	0.6038	0.8878
Error for the wave height squared	0.4076		0.5435		0.5753	
Error for the wave period	0.1221		0.1534		0.1595	

enhances the predictive ability of the time-varying parameter regressions.

The major caveat associated with these findings is that the errors from the various forecasting models all lie within a fairly narrow range. Evidently, there are limits to the predictability of the flux, even with the most advanced models currently available. The range of error in these data sets is 17–25 percent at an hourly horizon, 28–41 percent at a daily resolution. This is similar to the range of errors found in other forms of renewable energy, such as solar and wind. The high error implies that utilities will need backup sources of power in order to compensate for the inherent variability of waves.

References

Aburto, L., Weber, R., 2007. Improved supply chain management based on hybrid demand forecasts. Applied Soft Computing 7, 136–144.

Baillie, R.T., 1996. Long memory processes and fractional integration in econometrics. Journal of Econometrics 73, 5–59.

Beran, J., 1994. Statistics for Long-Memory Processes. Chapman & Hall, New York.

Box, G.E.P., Jenkins, G.M., 1970. Time Series Analysis: Forecasting and Control. Holden-Day, San Francisco.

Calvet, L.E., Fisher, A.J., 2001. Forecasting multifractal volatility. Journal of Econometrics 105, 27–58.

Calvet, L.E., Fisher, A.J., 2002. Multifractality in asset returns: theory and evidence. Review of Economics and Statistics 84, 381–406.

Calvet, L.E., Fisher, A.J., 2004. How to forecast long-run volatility: regime-switching and the estimation of multifractal processes. Journal of Financial Econometrics 2, 49–83.

Deo, M.C., Jagdale, S.S., 2003. Prediction of breaking waves with neural networks. Ocean Engineering 30, 1163–1178.

Deo, M.C., Naidu, C.S., 1998. Real time wave forecasting using neural networks. Ocean Engineering 26, 191–203.

Deo, M.C., Jha, A., Chaphekar, A.S., Ravikant, K., 2001. Neural networks for wave forecasting. Ocean Engineering 28, 889–898.

Gaur, S., Deo, M.C., 2008. Real-time wave forecasting using genetic programming. Ocean Engineering 35, 1166–1172.

Granger, C.W.J., 2008. Non-linear models: where do we go next–time varying parameter models? Studies in Nonlinear Dynamics and Econometrics 12 (3) (Article 1. 〈http://www.bepress.cxom/snde/vol12/iss3/art1〉).

Hasselmann, D.E., Dunckel, M., Ewing, J.A., 1980. Directional wave spectra observed during JONSWAP 1973. Journal of Physical Oceanography 10, 1264–1280.

Hasselmann, K., Ross, D.B., Müller, P., Sell, W., 1976. A parametric wave prediction model. Journal of Physical Oceanography 6, 200–228.

Hasselmann, S., Hasselmann, K., Allender, J.H., Barnett, T.P., 1985. Computations and parameterizations of the non-linear energy transfer in a gravity wave spectrum. Part II: parameterizations of the non-linear energy transfer for application in wave models. Journal of Physical Oceanography 15, 1378–1391.

Hibbert, H., Pedreira, C., Souza, R., 2000. Combining neural networks and ARIMA models for hourly temperature forecasts. In: Proceedings of the International Conference on Neural Networks IJCNN 2000, pp. 414–419.

Hibon, M., Evgeniou, T., 2005. To combine or not to combine: selecting among forecasts and their combinations. International Journal of Forecasting 21, 15–24.

Ho, P.C., Yim, J.Z., 2006. Wave height forecasting by the transfer function model. Ocean Engineering 33, 1230–1248.

Jain, P., Deo, M.C., 2007. Real-time wave forecasts off the western Indian coast. Applied Ocean Research 29, 72–79.

Janssen, P., 1991. Quasi-linear theory of wind-wave generation applied to wave forecasting. Journal of Physical Oceanography 21, 1631–1642.

Kalman, R.E., 1960. A new approach to linear filtering and prediction problems. Transactions of the American Society of Mechanical Engineers. Series D, Journal of Basic Engineering 83, 35–45.

Komen, G.J., Cavaleri, L., Donelan, M., Hasselmann, K., Hasselmann, S., Janssen, P.A.E.M., 1994. Dynamics and Modelling of Ocean Waves. Cambridge University Press, Cambridge.

Londhe, S.N., 2008. Soft computing approach for real-time estimation of missing wave heights. Ocean Engineering 35, 1080–1089.

Londhe, S.N., Panchang, V., 2006. One-day wave forecasts based on artificial neural networks. Journal of Atmospheric and Oceanic Technology 23, 1593–1603.

Lux, T., 2008. The Markov-switching multifractal model of asset returns: GMM estimation and linear forecasting of volatility. Journal of Business and Economic Statistics 26, 194–210.

Makarynskyy, O., 2004. Improving wave predictions with artificial neural networks. Ocean Engineering 31, 709–724.

Malmberg, A., Holst, U., Holst, J., 2005. Forecasting near-surface ocean winds with Kalman filter techniques. Ocean Engineering 32, 273–291.

Miles, J.W., 1957. On the generation of surface waves by shear flows. Journal of Fluid Mechanics 3, 185–204.

Nelson, M., Hill, T., Remus, W., O'Connor, M., 1999. Time series forecasting using neural networks: should the data be deseasonalized first? Journal of Forecasting 18, 359–367.

Phillips, O.M., 1957. On the generation of waves by a turbulent wind. Journal of Fluid Mechanics 2, 417–445.

Phillips, O.M., 1958. The equilibrium range in the spectrum of wind-generated waves. Journal of Fluid Mechanics 4, 426–434.

Pierson, W.J., Moskowitz, L., 1964. A proposed spectral form for fully developed wind-seas based on the similarity theory of S. A. Kitaigorodskii. Journal of Geophysical Research 69, 5181–5190.

Roulston, M.S., Ellepola, J., von Hardenberg, J., Smith, L.A., 2005. Forecasting wave height probabilities with numerical weather prediction models. Ocean Engineering 32, 1841–1863.

Reikard, G., 2008. Using temperature and state transitions to forecast wind speed. Wind Energy 11, 431–443.

Schertzer, D., Lovejoy, S., Schmitt, F., Chigirinskaya, Y., Marsan, D., 1997. Multifractal cascade dynamics and turbulent intermittency. Fractals 5, 427–471.

Snyder, R.L., Dobson, F.W., Elliott, J.A., Long, R.B., 1981. Array measurements of atmospheric pressure fluctuations above surface gravity waves. Journal of Fluid Mechanics 102, 1–59.

Song, J.B., 2006. Probability distribution of random wave-current forces. Ocean Engineering 33, 2435–2453.

Sverdrup, H.U., Munk, W.H., 1947. Wind, Sea and Swell: Theory of Relations for Forecasting. US Navy Hydrographic Office, Washington, DC.

Theil, H., 1971. Principles of Econometrics. Wiley, New York.

Tsai, C.P., Lin, C., Shen, J.N., 2002. Neural network for wave forecasting among multi-stations. Ocean Engineering 29, 1683–1695.

Tseng, C.M., Jan, C.D., Wang, J.S., Wang, C.M., 2007. Application of artificial neural networks in typhoon surge forecasting. Ocean Engineering 34, 1757–1768.

Hasselmann, S., Hasselmann, K., Bauer, E., Janssen, P.A.E.M., Komen, G., Bertotti, L., Lionello, P., Guillaume, A., Cardone, V.C., Greenwood, J.A., Reistad, M., Zambresky, L., Ewing, J.A., WAMDI Group, 1988. The WAM model: a third generation wave prediction model. Journal of Physical Oceanography 18, 1775–1810.

Zamani, A., Solomatine, D., Azimian, A., Heemink, A., 2008. Learning from data for wind-wave forecasting. Ocean Engineering 35, 953–962.

Zhang, G.P., 2003. Time series forecasting using a hybrid ARIMA and neural network model. Neurocomputing 50, 159–175.

Zhang, G.P., Qi, M., 2005. Neural network forecasting for seasonal and trend time series. European Journal of Operational Research 160, 501–514.