



BIROn - Birkbeck Institutional Research Online

Athanasopoulou, Maria Eleni and Deveikyte, Justina and Mosca, Alan and Peri, Ilaria and Proveti, Alessandro (2021) A hybrid model for forecasting short-term electricity demand. In: 2nd ACM International Conference on AI in Finance, 03-05 Nov 2021, London, UK. (In Press)

Downloaded from: <https://eprints.bbk.ac.uk/id/eprint/47349/>

Usage Guidelines:

Please refer to usage guidelines at <https://eprints.bbk.ac.uk/policies.html>
contact lib-eprints@bbk.ac.uk.

or alternatively

A Hybrid Model for Forecasting Short-Term Electricity Demand

Maria Eleni Athanasopoulou, Justina Deveikyte, Alan Mosca, Ilaria Peri and Alessandro Proveti

Birkbeck Institute for Data Analytics

Birkbeck, University of London

London WC1E 7HX, UK

ABSTRACT

Currently the UK Electric market is guided by load (demand) forecasts published every thirty minutes by the regulator. A key factor in predicting demand is weather conditions, with forecasts published every hour. We present HYENA: a hybrid predictive model that combines feature engineering (selection of the candidate predictor features), mobile-window predictors and finally LSTM encoder-decoders to achieve higher accuracy with respect to mainstream models from the literature. HYENA decreased MAPE loss by 16% and RMSE loss by 10% over the best available benchmark model, thus establishing a new state of the art for the UK electric load (and price) forecasting.

KEYWORDS

Hybrid models, Neural Networks, Regression, Feature Engineering.

ACM Reference Format:

Maria Eleni Athanasopoulou, Justina Deveikyte, Alan Mosca, Ilaria Peri and Alessandro Proveti. 2021. A Hybrid Model for Forecasting Short-Term Electricity Demand. In *2nd ACM International Conference on AI in Finance (ICAIF'21)*, November 3–5, 2021, Virtual Event, USA. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3490354.3494371>

1 INTRODUCTION

Electricity cannot be generated in advance and stored in large quantities to be used at a later time, it needs to be generated in real time. Thus, having accurate demand forecasts is crucial to ensuring security of supply. Furthermore, generators, suppliers and non-physical traders that participate in the electricity market, all need to have an accurate view on future demand.

For this reason, the academic literature examined several methods to improve the accuracy of the demand forecasting, but no model has been proven consistently superior regardless of the dataset and the market. Furthermore although the wider topic of load forecasting has been widely researched, there has been significantly less research on electricity systems that have a lot of embedded renewable generation, such as the UK.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICAIF'21, November 3–5, 2021, Virtual Event, USA

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-9148-1/21/11...\$15.00

<https://doi.org/10.1145/3490354.3494371>

The existence of a significant amount of embedded renewable generation in the system, suppresses load as it covers some of the demand for electricity and as a result less is requested from the grid. As the renewable generation is intermittent it introduces more uncertainty to the system, making accurate forecasting more challenging and of higher importance. In the UK, there are approximately 30 GWs of embedded generation connected to the distribution system [1], making predicting UK national demand extremely challenging.

Today, the key features that would intuitively be the basis for an electric demand prediction are basically public. In particular, past load levels and weather conditions, arguably two key features, are available for extraction from public Web sites. Other auxiliary features such as holiday and renewable installed capacity data are also available.

Nonetheless, predicting load on the time-scale of industrial operations, i.e., every thirty minutes, remains a challenge; even well-established methods, including neural networks, incur in huge losses, as shown in our preliminary study on forecasting the current UK electric demand.

In this article we respond to the challenge by defining a new hybrid model, called Hybrid Energy Analyser (HYENA), which we have designed and validated specifically for load forecasting (demand prediction) in the UK electric market. Our comparative assessment over a long period of historical demand shows that the HYENA model is on average more accurate than any of the established models over which it is based, thus confirming that hybrid models are the best option for this type of forecasting problem.

We approached the challenge by leveraging public data to train a multi-level hybrid system with memory-based machine learning at its core; using traditional forecasting methods to capture seasonal effects and feature engineering to filter the input data; and deploying Machine Learning to the ‘internal’ problem of learning the deviation in demand from the seasonally-adjusted data.

Our resulting architecture combines Linear regression, to capture seasonality, and a Long-Short Term Memory (LSTM) auto-encoder, to capture the “residual” variations in demand. To validate the proposed model, we used Tao Hong’s vanilla model as a benchmark. We also compared the forecasting ability of the hybrid model with that of a wide range of methods such as SARIMA, SARIMA with exogenous variables, Random Forests and simpler LSTM models, using common load forecast accuracy metrics. The proposed model outperformed all other tested methods, achieving a significant error reduction with respect to the benchmark model.

1.1 Literature on Electric load forecasting

Academic research has focused on demand forecasting over the years and different forecasting methods have been suggested depending on the forecasting horizon. For instance, when forecasting

very short-term horizons up to 6 hours, using univariate models such as ARIMA and Holt-Winters exponential smoothing [2], [3] is sufficient [4], while for longer horizons multivariate models are more appropriate. For forecasting up to two weeks ahead, demand is driven by exogenous factors such as weather, calendar and special events. Temperature is the main weather-related driver of demand. Its relationship with load is non-linear and it is driven by heating and cooling demand [5], [6]. Other factors such as humidity [7], cloud and wind speed [8] subtly affect demand and research has shown that using them as predictors improves forecasting accuracy. For systems such as California and the UK, where there is a lot of embedded solar generation in the systems suppressing demand, using solar radiation and installed solar capacity as predictors of demand further improves model accuracy [9].

One of the most widely used forecasting methods both in academia [10],[11] and industry [6], is Multiple Linear Regression (MLR). In fact, MLR is able to capture seasonality well with the help of seasonal variables; and can describe the more complex patterns, such as the relationship between load and temperature, using polynomials.

MLR is highly interpretable and faster to implement than other more complex methods such as Neural Networks. Random Forests [12] using weather and calendar variables have also been proven to be highly accurate, surpassing the forecasting abilities of other flexible approaches such as Neural Networks. Random Forests have the advantage that they can accept large volume of features without majorly impacting their performance when these inputs are not relevant [13]. Deep learning methods, such as various architectures of LSTM [14] and LSTM hybrid models [15], [16], have also proven to be promising, as LSTM models are able to learn long-term temporal connections and do not suffer from the vanishing gradient problem.

Combining LSTM ability to learn long-term dependencies with other models, such as CNN, to capture local trend in data has been shown to outperform simpler LSTM implementations [15]. Moreover, [17] has shown the benefits of using a Stacked Denoising Auto-Encoder, a type of multilayer perceptron (MLP), to learn efficient data coding and use it to accurately forecast demand.

2 DESIGN AND IMPLEMENTATION

The project followed a multi-step process. Firstly, we selected the weather stations we would use in our models and gathered the load, weather, holiday and embedded generation installed capacity data. In the next stage, we carefully analysed the historic load to uncover the main drivers of the UK electricity demand and engineered features that the analysis showed that could be useful in predicting it. We then created our baseline model and other candidate models, applying dimensionality reduction techniques where relevant. We fine-tuned the models and, finally, we evaluated them using error metrics commonly used in academia.

2.1 Input data and sources

Load: the load data used to build our models was the national demand consumption for the UK and spans between January 2018 and February 2020. However, to get a better understanding of the data, its trends and its drivers, we examined the load demand from 2014

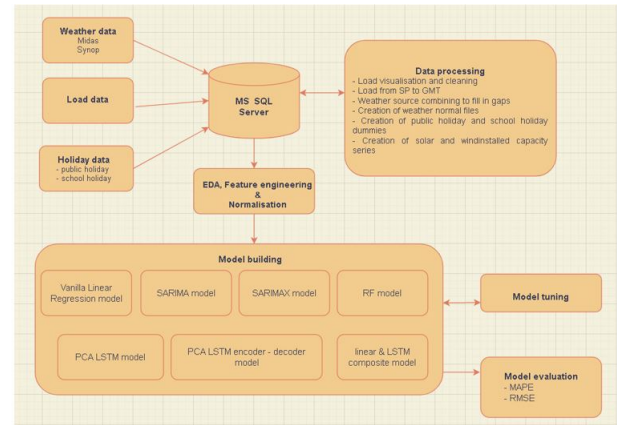


Figure 1: Methodology flowchart

until February 2020. The data was obtained from the BM Reports API which publishes real time data with half hourly granularity.

Weather: the weather data was obtained by the CEDA website using as a primary source the Land Synop data [18]; to complete any missing data we used the Midas open datasets.

Other inputs. For the period covered by this study we collected the daily estimation of the embedded solar and wind installed capacity reported by National Grid¹. Also, got the exact dates of holidays from the UK government website.

2.2 Data exploration

Looking at the historic electricity demand, it is easy to notice a slow year-on-year decline, shown on Figure 2, which displays, month by month, the average half-hourly demand. The drift is considered a consequence of technological advances, economic factors and the increasing installed capacity of renewable generation. To account for this long-term tendency, we used a trend variable in our models.

Examining the historic load at a finer level of detail one can see that the solar and wind installed capacity increase over last few years has in fact changed the shape of demand. Between the hours that solar panels can produce substantial amounts of energy (11:00-17:00) we see a dip in demand which is especially pronounced between April and October. We also see a dip in the first morning hours due to the embedded wind generating enough power to cover a significant part of the load needed in these low demand hours.

Given the big impact that embedded renewable generation has had on demand, it is evident that modelling demand requires including some proxy for the embedded generation factor. To account for this aspect in the models, at each half-hourly period, we used solar radiation and wind as proxies, while to account for the year on year difference in embedded generation, we used installed capacity as an input.

To account for the long-, medium- and short-term trends seen in the data, we had to include features representing both daily, weekly and yearly seasonality.

¹Please see the National Grid's Data Explorer.

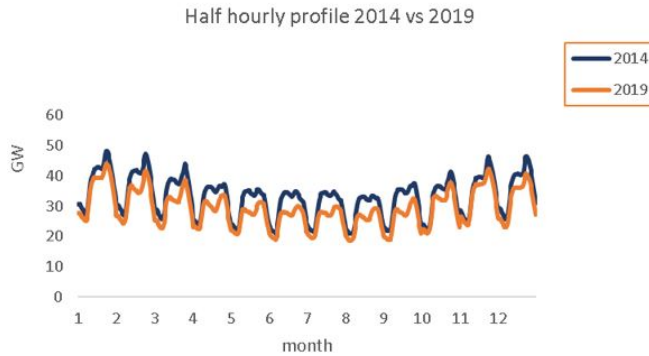


Figure 2: For each month, we plot the average half-hourly load for 2019 (in orange) and 2014 (in blue), the year when renewable installed capacity started its growth. Demand now has a more pronounced dip between 12:00 and 16:00, as a bigger share of demand is covered by solar radiation and less energy is requested from the grid.

A further aspect of the load series is its high variance. In order to “stabilise” it over time, we have applied a log transformation on the load, as suggested by [5] before running the benchmark models described below. However, the input load to the HYENA model was not transformed. Vice-versa, as LSTM is sensitive to data scale, the input for the LSTM(-AE) benchmarks described below was normalised to the $[0,1]$ range.

2.3 Weather data

Weather is one of the main drivers of demand, especially when forecasting more than 6 hours ahead. As we set out to model national demand, we included a small set of weather stations that would nonetheless be considered representative of the country’s weather. One important point is that the stations need to be geographically dispersed so they are not too collinear between them, as collinearities could create unstable load forecasts (please see, e.g., [19] for a discussion). Keeping this condition in mind, we used population density as proxy for both economic activity in the region and residential energy usage, when we selected the weather stations used in the model. Given the significant effect of renewable generation on demand, we also selected weather stations on areas that have high solar and wind installed capacity to act as proxy for the renewable energy being injected to the system.

2.4 Feature engineering and variable selection

From the exploratory analysis of the load, we have already discovered a few variables that would be important predictors of demand. Demand decreases slightly year on year; thus we included a trend in our models. To account to the weekly seasonality, we included a variable for the days of the week. Monthly seasonality is captured with monthly dummies.

To account for the difference in levels between working days and holidays, we included a dummy variable that has the value 1 when the day is a public holiday. To capture the daily shape of demand, we included a variable for time. Next, we consider the

school calendar: as during school holidays we see a later ramp up in demand in the mornings, we added a dummy variable denoting school holidays to our models.

One of the biggest drivers of demand, beyond seasonality, is temperature. Load and temperature have a non-linear relationship. As shown in Figure 3, when temperature is low, demand increases, while when it is high, demand decreases until it reaches around 291 degrees Kelvin, when the cooling demand starts. This can be captured by taking the second-degree polynomial function of temperature.

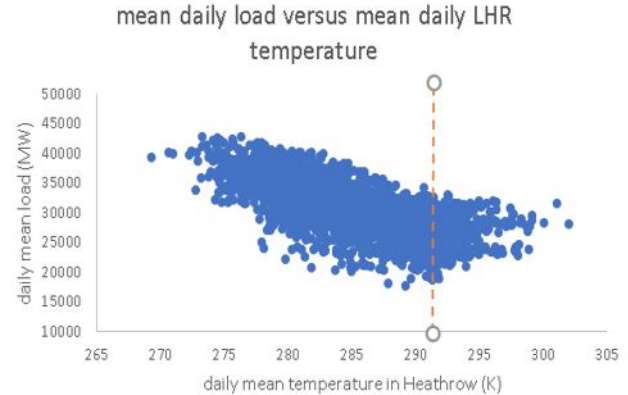


Figure 3: Daily load plotted against mean daily temperature in Heathrow; the non-linear relationship between load and temperature can be observed.

As we saw in the load exploratory analysis, renewable embedded generation has a big effect on the shape of demand. To capture the long-term changes in demand caused by the high renewable penetration, we included variables capturing the installed embedded solar and wind capacity; while to capture the half hourly effect of renewables on load, we included variables showing the solar radiation and wind speed per half hour for the selected 5 weather stations. Weather does not have an instantaneous effect on load only.

After we identified the variables that would be good predictors of demand, we used feature importance from Extremely Randomized Regression Trees to validate the variable selection we made based on the exploratory analysis. The Extremely Randomized Regression Trees is an ensemble method that fits a number of randomised trees on various sub-samples of the dataset without bootstrapping and uses averaging to improve the predictive accuracy and control overfitting. As expected, the variables selected explained enough of the variance in the load to justify their inclusion in the model; with the most important variables being the time related variables. As the UK is not large enough to have multiple climates, we expected some collinearities between weather stations despite purposely selecting them so that they are geographically dispersed. This was verified by looking at the correlations between stations. However, as we are aiming to predict national demand, we chose to include all selected weather stations as we needed to capture weather differences between regions. We did however, apply Principal Component Analysis(PCA) in an effort to reduce the dimensions of the dataset and

eliminate collinearities in the data. PCA identifies the hyperplane that lies closest to the data, projects the data on it, and selects the axis that preserves the maximum amount of variance, so as to lose as little information as possible. While it removes any excessive multi-collinearity and reduces data dimensionality, PCA is also known to remove some explainability from data. Since our focus is on prediction, we chose to apply it nonetheless.

3 THE HYENA MODEL

HYENA attempts to capture seasonality and trend with a simple linear model, creating a forecast of what the load would be with average weather; next it models the difference between such forecast and the actual observed load. This is a novel solution where, instead of trying to predict the signal per se, we first removed the known factors of seasonality and focused on predicting the residual.

We have implemented HYENA by training the linear model with data between 2014 and 2016 using calendar, installed capacity and holiday data. For the second step we used an LSTM auto-encoder to forecast the difference between the linear model forecast and the actuals for the period between 2017 and 2019:

$$\varepsilon_t = \hat{y}_t - y_t \quad (1)$$

with

$$\begin{aligned} \hat{y}_t = & \beta_0 + \beta_1 M_{2t} + \beta_2 M_{3t} + \beta_3 M_{4t} + \beta_4 M_{5t} \\ & + \beta_5 S M_{6t} + \beta_6 M_{7t} + \beta_7 M_{8t} + \beta_8 M_{9t} + \beta_9 M_{10t} \\ & + \beta_{10} M_{11t} + \beta_{11} M_{12t} + \beta_{13} t_t + \beta_{14} wd_t \\ & + h_{public} + Trend_t + C_{solar} + C_{wind} \end{aligned}$$

where $M_{2t} \dots M_{12t}$ are monthly dummies, t_t is time, wd_t is week-day, h_{public} is a public holiday dummy, C_{solar} is the solar installed capacity and C_{wind} is the wind installed capacity.

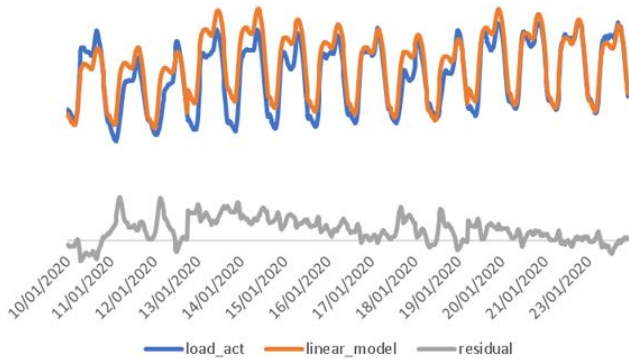


Figure 4: Linear forecast (in orange) to create predictions with average weather and residual demand (in grey) to predict the remaining effects not captured by seasonality and event variables.

Once the forecast is produced by the linear model, we take its difference from the actual load to create the residual series (Figure 4), which represents the deviation from normal demand and is dependent mostly on weather. To validate this and to select which

of the available variables to use in the LSTM auto-encoder model, we used feature importance from Extremely Randomised Trees. Using current and lagged variables for temperature and solar radiation, as well as time to capture some remaining seasonality, we built an encoder-decoder LSTM model with an LSTM layer for encoding and an LSTM layer for decoding, followed by a dense layer. To optimise the model, we deployed the well-known Adam optimiser with a small batch size and early stopping to avoid over training.

4 COMPARATIVE MODEL EVALUATION

To validate the HYENA approach, and in general to create a comprehensive assessment of predictive models for UK demand forecasting, we have implemented and tuned some of the best-known models from the literature, which are briefly introduced next. Our baseline model for this prediction task is Tao Hong’s “vanilla” model, which has been used as a benchmark for the GEFCom2012 energy competition and subsequently in many papers focusing on load forecasting. Next, we have set up and run instances of

- the Seasonal Autoregressive Integrated Moving Average (SARIMA): an extension of ARIMA that explicitly supports data with a seasonal component, which is the case here;
- the SARIMAX model, which extends the above by including exogenous variables [20] such as weather and calendar variables. SARIMAX thus addresses the potential shortcoming of SARIMA which takes into account past load values;
- Random Forest, the well-known ensemble method was deployed here by building a random-forest regressor with 200 trees, with maximum depth of 10 and at least 45 samples required to split to an internal node;
- LSTM, the well-known deep learning method which can handle non-stationary data (as it is the case with load series) but requires a careful manual tuning. Next, Principal component Analysis (PCA) was deployed to reduce the data dimensions down to 9. Finally, we implemented a many-to-many LSTM network with 48 timesteps, 9 features and 2 LSTM layers, where each cell per timestep emits a signal to the same timestep of the next layer, followed by a dense layer, and
- LSTM autoencoder, which improves LSTM by an encoding layer which suppresses the inputs and keeps only the most important signals for prediction. As we did for LSTM, the model was optimised by minimising the mean square error with the Adam optimiser; also Early stopping was employed to avoid over-training the network.

Each model has been tuned, when possible, to best fit historical data and thus make the comparison fair. For lack of space, we have put the description of this phase in the supplementary material to this article.

The candidate models were evaluated based on the day-ahead forecasts (i.e., predict the exact level of demand tomorrow). Our testing period ran from the 1st of January until the 14th of February 2020. Forecasts were evaluated against the retrospective actual values using the Mean Absolute Percentage Error (MAPE) and Root Mean Square Error (RMSE), which are standard for the electricity markets. While model training time could be a consideration in this

setting, we did not deem it a useful evaluation criterion as none of the models had a prohibitively long training time.

Metric	Baseline	SARIMA	SARIMAX	RF	LSTM	LSTM AE	Hybrid
MAPE	6.59	10.26	9.62	5.32	4.82	4.21	3.33
RMSE	2468	3787	3656	1965	1723	1611	1280

Figure 5: Average MAPE and RMSE during the testing period (Jan. 1-Feb. 14 2020).

The baseline model (Tao Hong’s vanilla model), for all its simplicity, is not easily outperformed, as all models struggled in predicting the ramp-ups in demand. However, the more flexible models managed a flatter error shape than the baseline for the rest of the day.

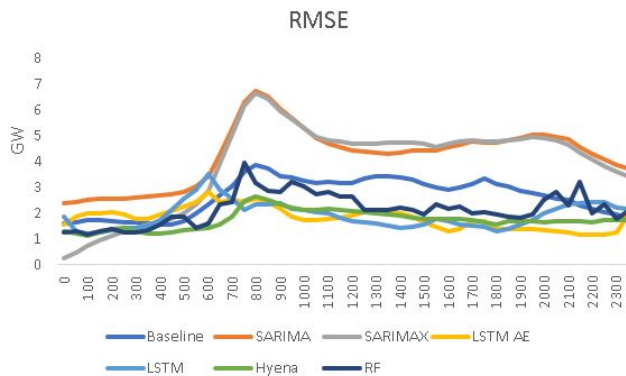


Figure 6: Average RMSE per hour during the testing period (1st Jan to 14th Feb 2020).

SARIMA and SARIMAX hardly outperformed Hong’s baseline model; and contrary to what most of the literature suggests, nor they outperformed LSTMs or Random Forest. A possible interpretation, is that our data presentation could have failed to capture the yearly seasonality of the load, as we had to restrict our training sample to rolling 30 days due to restricted computational resources.

SARIMAX, which takes into account temperature and calendar variables, outperformed SARIMA, but as it was lacking any proxies for solar radiation, it mis-assigned the solar effect leading to higher errors than SARIMA at peak hours for solar radiation (11 am to 3 pm). The Random Forest model outperformed the baseline but showed a slightly worse predictive ability than the LSTM models.

In comparison to the Random Forest, the LSTM model was more difficult to implement, as it required more pre-processing steps and a dimensionality reduction stage while the gains in predictive power over the former were minimal. Also the LSTM encoder-decoder implementation managed to outperform all previous models. As the LSTM-AE compresses its data, it managed to identify the most important signals in the data, which lead to significantly better performance and less time dependent error.

Finally, our HYENA models outperformed all benchmarks as it captured i) seasonality in data with a well-specified linear sub-model and ii) the residual demand effects with an LSTM encoder-decoder sub-model.

5 CONCLUSIONS

Our study of the UK electric market and the multiple factors that affect demand lead us to the design and implementation of HYENA, a hybrid model that shows excellent predictive power: on a long-term dataset containing half-hourly demand data for the first 45 days of 2020 it convincingly outperformed the accuracy of six benchmarks, including statistical (SARIMA, SARIMAX), Machine Learning (Random Forests) and Deep Learning methods (LSTM, LSTM autoencoder): HYENA decreased MAPE by 16% and RMSE by 10% over the best available model, thus establishing a new state of the art for the UK electric load forecasting.

HYENA combines existing methods to implement a novel approach to modelling demand, not necessarily electric demand: regression methods to capture seasonality and LSTM autoencoders to capture the “residual” variations in demand. However, several preliminary activities are required to obtain such an effective model.

A thorough data acquisition and analysis of the load and its drivers was conducted, which highlighted the challenges that are involved when real data is used for model design, such as automatic data collection from multiple inconsistent sources, gaps and temporal inconsistencies in the data, the difficulties of synthesising data from multiple sources to create as a complete a dataset as possible; and unavailability of all the required data. Next, a range of forecasting methods was investigated and several intermediate steps, e.g., normalisation and PCA, were executed and evaluated. An additional challenge was the limited computational power at our disposal, which was particularly felt for the SARIMAX implementation: more resources would mean a longer training set and a model extension which could account for yearly seasonality.

Finally, we plan further research on HYENA with a focus on predicting regional demand, which could then be aggregated to create a national demand forecast. We expect that this approach would lead to a sharp increase of accuracy, as regional load would be easier to forecast using the relevant weather station, rather than using a few stations to forecast the demand for the whole country.

This approach would also solve the problem of collinearity between stations being and would eliminate the need for dimensionality reduction as all of the data would be essential. Finally, we would be interested in including rain, humidity and cloud data to the proposed model as these would help explain a part of the error which we believe could be significant.

REFERENCES

- [1] UK The National Grid. Energy forecasting strategic project roadmap. <https://www.nationalgrideso.com/document/145941/download>, 2020.
- [2] J. W. Taylor. Short-term electricity demand forecasting using double seasonal exponential smoothing. *The Journal of the Operational Research Society*, 54(8):799–805, 2003.
- [3] James W. Taylor. An evaluation of methods for very short-term load forecasting using minute-by-minute british data. *International Journal of Forecasting*, 24(4):645–658, 2008.
- [4] James W. Taylor, Lilian de Menezes, and Patrick McSharry. A comparison of univariate methods for forecasting electricity demand up to a day ahead. *International Journal of Forecasting*, 22(1):1–16, 2006.
- [5] Souhaib Ben Taieb and Rob Hyndman. A gradient boosting approach to the kaggle load forecasting competition. *International Journal of Forecasting*, 30(2):382–394, 2014.
- [6] Tao Hong, Pierre Pinson, and Shu Fan. Global energy forecasting competition 2012. *International Journal of Forecasting*, 30(2):357–363, 2014.
- [7] J. Xie, Y. Chen, T. Hong, and T. D. Laing. Relative humidity for load forecasting models. *IEEE Transactions on Smart Grid*, 9(1):191–198, 2018.

- [8] Jingrui Xie and Tao Hong. Wind speed for load forecasting models. *Sustainability*, 9:795, 05 2017.
- [9] Zhiying Li, Xun Zou, and Steven Quiring. Comparison of three short-term load forecast models in southern california. *Energy*, 189:116358, 10 2019.
- [10] Tao Hong. *Short Term Electric Load Forecasting*. Springer, 2010.
- [11] A. Papalexopoulos and T. Hesterberg. A regression-based approach to short-term system load forecasting. In *Conference Papers Power Industry Computer Application Conference*, pages 414–423, 1989.
- [12] Ali Lahouar and J. Ben Hadj Slama. Day-ahead load forecast using random forest and expert input selection. *Energy Conversion and Management*, Volume 103:Pages 1040–1051, 07 2015.
- [13] Ali Lahouar and J. Ben Hadj Slama. Hour-ahead wind power forecast based on random forests. *Renewable Energy*, 109:–, 03 2017.
- [14] Ruiguo Yu, Jie Gao, Mei Yu, Wenhuan Lu, Tianyi Xu, Mankun Zhao, Jie Zhang, Ruixuan Zhang, and Zhuo Zhang. LSTM-EFG for wind power forecasting based on sequential correlation features. *Future Generation Computer Systems*, 93, 10 2018.
- [15] Chujie Tian, Jian Ma, Chunhong Zhang, and Panpan Zhan. A deep neural network model for short-term load forecast based on long short-term memory network and convolutional neural network. *Energies*, 11:3493, 12 2018.
- [16] K. Yan, W. Li, Z. Ji, M. Qi, and Y. Du. A hybrid LSTM neural network for energy consumption forecasting of individual households. *IEEE Access*, 7:157633–157642, 2019.
- [17] Peng Liu, Peijun Zheng, and Ziyu Chen. Deep learning with stacked denoising auto-encoder for short-term electric load forecasting. *Energies*, 12:2445, 06 2019.
- [18] Centre NCAS. LAND SYNOP reports from land stations collected by the met office MetDB system. Retrieved from <https://catalogue.ceda.ac.uk/uuid/9f80d42106ba708f92ada730ba321831>, 2020.
- [19] Carsten Dormann and et al. Collinearity: a review of methods to deal with it and a simulation study evaluating their performance. *Ecography*, 36(1):27–46, 2013.
- [20] George Papaioannou, Christos Dikaiakos, Anargyros Dramountanis, and Panagiotis Papaioannou. Analysis and modeling for short- to medium-term load forecasting using a hybrid manifold learning principal component model and comparison with classical statistical models (sarimax, exponential smoothing) and artificial intelligence models (ann, svm): The case of greek electricity market. *Energies*, 9(8):635, Aug 2016.