

Short-Term Firm-Level Energy Consumption Forecasting

Subjects: Energy & Fuels

Contributor: Patricia Takako Endo, Andrea Maria Ribeiro

To minimise environmental impact, avoid regulatory penalties, and improve competitiveness, energy-intensive manufacturing firms require accurate forecasts of their energy consumption so that precautionary and mitigation measures can be taken. Deep learning is widely touted as a superior analytical technique to traditional artificial neural networks, machine learning, and other classical time series models due to its high dimensionality and problem solving capabilities. Despite this, research on its application in demand-side energy forecasting is limited. We compare two benchmarks (Autoregressive Integrated Moving Average (ARIMA), and an existing manual technique used at the case site) against three deep learning models (simple Recurrent Neural Networks (RNN), Long Short-Term Memory (LSTM), and Gated Recurrent Unit (GRU)) and two machine learning models (Support Vector Regression (SVR), and Random Forest) for short term load forecasting (STLF) using data from a Brazilian thermoplastic resin manufacturing plant. We use the grid search method to identify the best configurations for each model, and then use Diebold-Mariano testing to confirm the results. Results suggests that the legacy approach used at the case site is the worst performing, and that the GRU model outperformed all other models tested.

Keywords: Deep learning ; Energy forecasting

Introduction

The industrial sector is the largest consumer of delivered energy worldwide and energy-intensive manufacturing is the largest component in that sector ^[1]. Energy-intensity is driven by the mix of activity in these sectors including basic chemical feed-stocks, process (including heating and cooling) and assembly, steam and co-generation, and building-related energy consumption e.g. lighting, heating, and air conditioning ^[2]. World industrial energy consumption is forecast to grow from c. 242 quadrillion British thermal units (Btu) in 2018 to about 315 quadrillions Btu in 2050; the proportion of energy-intensive manufacturing is forecast to remain at approx. 50% during that period ^[1].

Extant literature has typically (i) focused on supply-side perspectives, (ii) aggregated energy costs, and (iii) failed to recognise the idiosyncrasies of the energy-intensive manufacturing sector and the associated centrality of energy management in production planning. There is a paucity of studies in demand-side process-related short-term load forecasting (STLF) using deep learning and machine learning for energy-intensive manufacturing. The limited studies that have been published do not compare deep learning performance against widely used machine learning models, classical time series models, or approaches used in practice. In addition to proposing prediction models, we also address this gap.

In this paper, we focus on performance analyses of deep learning and machine learning models for firm-level STLF considering energy-intensive manufacturing process of a thermoplastic resin manufacturer. We use energy consumption and production flow data from a real Brazilian industrial site. In addition to the energy consumption time series, we use data from two different stages of the thermoplastic resin production process – polymerisation and solid-state polymerisation. We propose three deep learning models - simple Recurrent Neural Networks (RNN), Long Short-Term Memory (LSTM) networks, and Gated Recurrent Unit (GRU) - and two machine learning models - Support Vector Regression (SVR), and Random Forest - for predicting daily energy consumption. These techniques were selected as they are referenced in the literature on STLF for demand-side energy consumption ^[3] ^[4]. We use the grid search method to identify the best model configurations. We compare the performance of the deep learning and machine learning models against (i) a classical time series model, Autoregressive Integrated Moving Average (ARIMA), and (ii) the current manual approach at the case site.

Finding models to predict energy consumption

The data used in this study was sourced from the Brazilian subsidiary of an international thermoplastic resin manufacturer, an energy-intensive manufacturing plant. Five years of data from 1 January 2015 to 31 December 2019 for daily total energy consumption at the plant (ENERGY dataset), as well as process-related data for two different stages of the manufacturing process - polymerisation (POLY_PRODUCTION dataset) and solid-state polymerisation (SSPOLY_PRODUCTION dataset) were provided. Each dataset comprised 1,826 values. These three series as input in to the models.

Due to the non-linear characteristics of the datasets used in this research, the need for both accuracy and fast run-times, and the promising results obtained in other works that used deep learning [5] [6] [7] [8], three deep learning techniques were selected for STLF in this study - simple RNN, LSTM, and GRU. In addition to these techniques, two different machine learning techniques were selected for the purpose of comparison - SVR, and Random Forest. These were selected as they feature in related works on STLF for demand side energy consumption [3] [4].

To determine the most suitable configuration for each deep learning model and machine learning model, we use the grid search method to determine the respective hyperparameters [9] [10] [11] [12] [13]. It is used widely as it is quick to implement, trivial to parallelise, and intuitively allows an entire search space to be explored [14].

Deep learning model setting

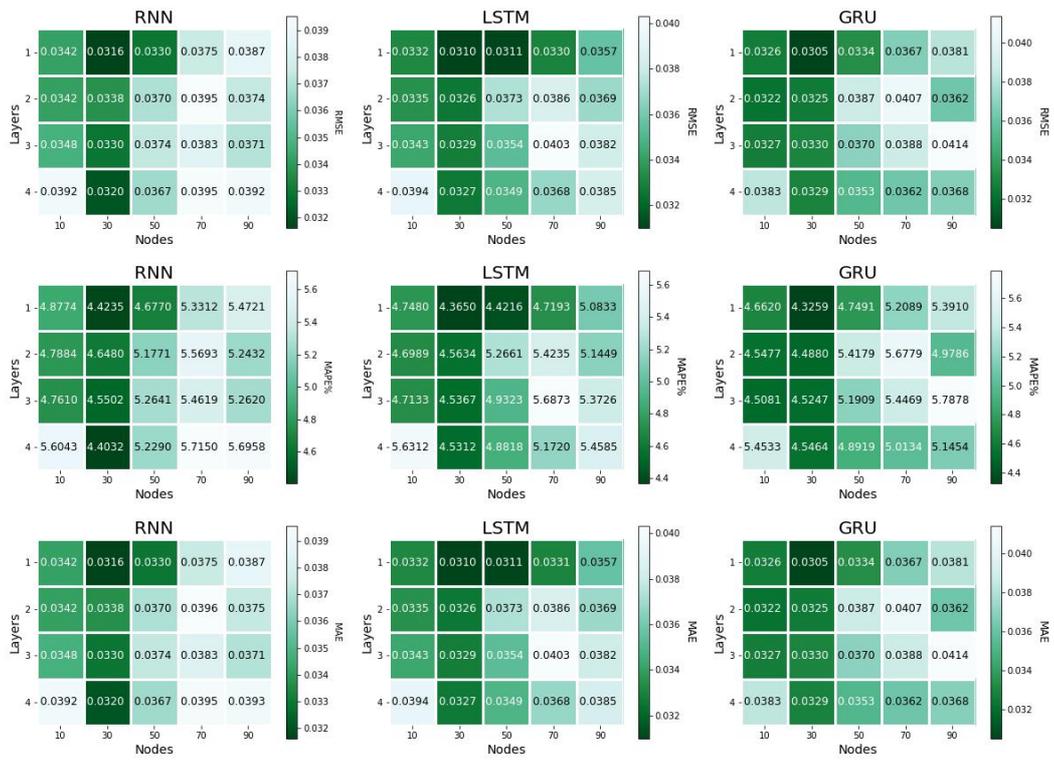
To perform the grid search, the dataset was separated into a training set consisting of 80% and 20% of the original dataset, respectively, using the percentage split. The hyperparameters evaluated by the grid search for deep learning techniques were (i) the number of layers, and (ii) the number of nodes in each layer (see Table 1).

Table 1. Machine learning parameters and levels.

Parameter	Levels
Number of nodes	From 10 to 90, step 20
Number of layers	From 1 to 4, step 1

Figures 1 present the grid search normalised results for RMSE, MAPE, and MAE, respectively. The configuration with one layer and 30 nodes generated the best models - RNN-1-30, LSTM-1-30, and GRU-1-30 - for all deep learning techniques, except in one instance where the average MAPE presented the best result for RNN with a configuration of four layers and 30 nodes i.e. RNN-4-30. These four model configurations will be used in our benchmark evaluation.

Figure 1. RMSE, MAPE and MAE grid search result for deep learning models.



Machine learning model settings

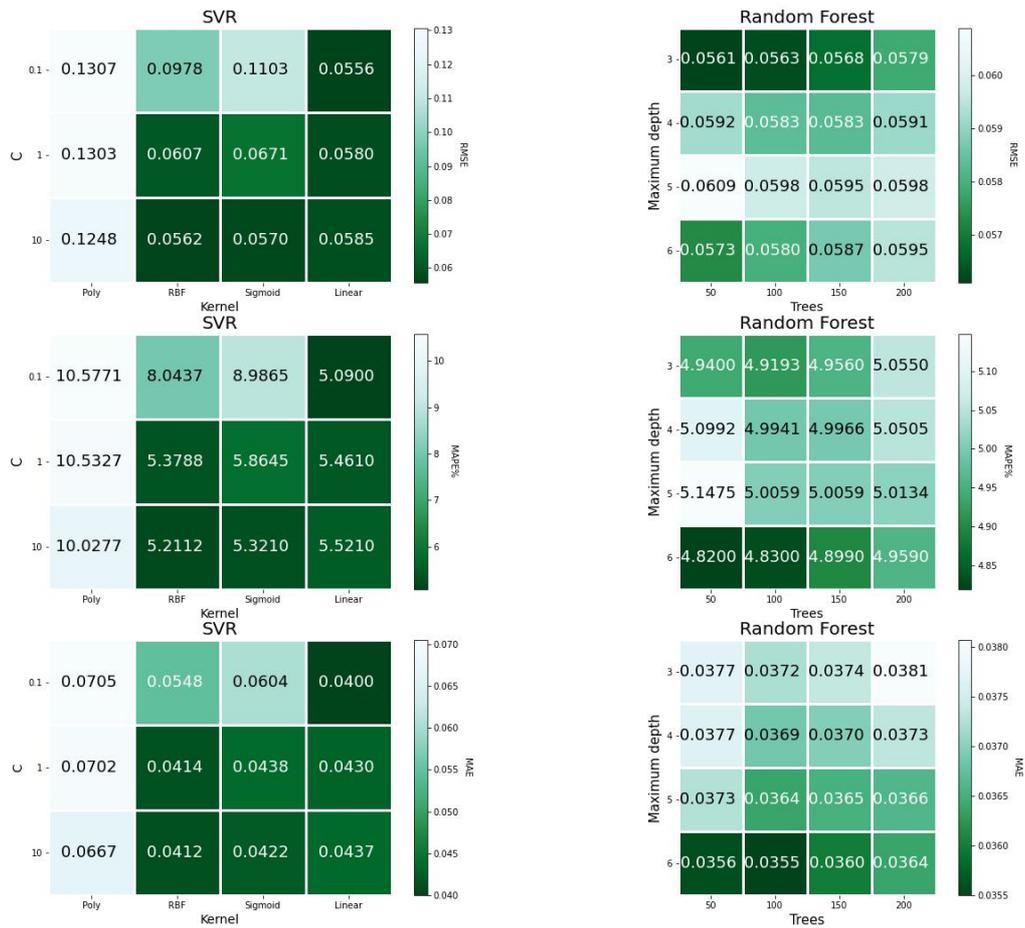
Similar to the deep learning models, we also perform a grid search to find the best hyperparameters of the machine learning models, using the same data set splitting procedure. The hyperparameters used vary according to the technique (see Table 2). For SVR, cost and the type of kernel were used whereas the maximum depth of trees and the number of trees were used for Random Forest.

Table 2. Machine learning parameters and levels.

Technique	Parameter	Levels
SVR	Number of C	0.1, 1 and 10
SVR	Type of kernel	Linear, polynomial and RBF
Random Forest	Number of max. depth	From 3 to 6, step 1
Random Forest	Number of trees	From 50 to 200, step 50

Figures 2 present the grid search results for RMSE, MAPE, and MAE, respectively. For SVR, the best configuration across the three metrics used - RMSE, MAPE, and MAE - is represented by SVR-0.1-linear, whose C value is 0.1 and used the linear kernel. For Random Forest, the models with configurations with (a) a maximum depth of three with 50 trees (Random Forest-3-50), (b) maximum depth of six with 50 trees (Random Forest-6-50), and (c) a maximum depth of six with 100 trees (Random Forest-6-100), generated the best results for RMSE, MAPE, and MAE, respectively. These four model configurations will be used in our benchmark evaluation.

Figure 2. RMSE, MAPE and MAE grid search result for machine learning models.



Results and discussions

Table 3 presents the RMSE, MAPE, MAE, and inference times results for the four deep learning models (RNN-1-30, RNN-4-30, LSTM-1-30, and GRU-1-30), and the four machine learning models (SVR-0.1-linear, Random Forest-3-50, Random Forest-6-50, and Random Forest-3-100) identified by the grid search method, as well as the manual and ARIMA benchmarks.

Table 3. RMSE, MAPE, MAE, and inference times results for selected analysed model.

Models	RMSE	MAPE (%)	MAE	Inference Time (s)
ARIMA	0.0471	3.52	0.0249	56.3565
RNN-1-30	0.0316	4.42	0.0316	0.3896
RNN-4-30	0.0320	4.40	0.0320	0.5939
LSTM-1-30	0.0310	4.37	0.0310	0.6751
GRU-1-30	0.0305	4.33	0.0305	0.7058
SVR-0.1-linear	0.0556	5.09	0.0400	0.0014

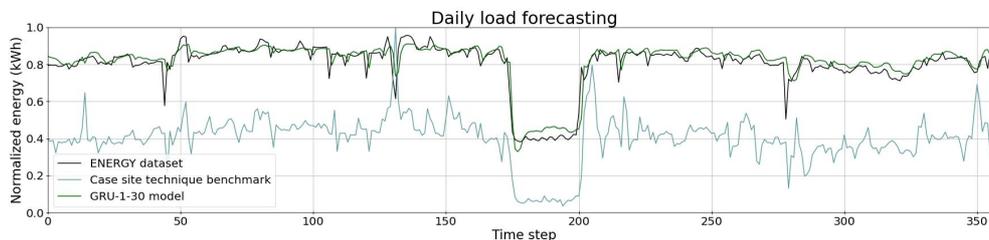
Random Forest 3-50	0.0561	4.94	0.0377	0.0043
Random Forest 6-50	0.0573	4.82	0.0356	0.0046
Random Forest 6-100	0.0580	4.83	0.0355	0.0080
Case site technique	0.4119	51.61	0.4039	-

Based on the RMSE metric, the deep learning models outperformed the machine learning models, and the manual and ARIMA benchmarks. This behaviour can be explained by the ability of deep learning models have to achieve insights outside of the domain of training data. The GRU model presented the best performance of all models tested as well as reducing the complexity inherent in the other deep learning models analysed; the simple RNN models presented the worst performance. In contrast, based on MAPE and MAE, the ARIMA model outperformed the deep learning models, the machine learning models, and the legacy manual approach. As the values of the RMSE, MAPE, and MAE are very similar, we used the Diebold-Mariano [15] test, a hypothesis test used to compare the significance of differences in two different prediction models and to confirm the results (GRU-1-30 presented de best performance).

With inference times of 0.8 and 0.0085, respectively, the deep learning and machine learning models performed best as a whole; standard deviations were insignificant. Random Forest-3-50 is the model with the shortest average inference time of those compared while the ARIMA model is the worst performing when compared to the machine learning and deep learning models. Although achieving good RMSE, MAPE and MAE results, the ARIMA inference time is much longer than the deep learning models, a significant limitation for practical use.

Figure 3 illustrate the daily load forecasts for the best deep learning model (GRU-1-30), the manual benchmark compared against the ground truth data. These clearly show that the proposed deep learning model is similar to the ground truth data compared to the existing manual technique used at the case site.

Figure 3. Daily load forecasting using GRU- 1-30, case site technique benchmark and ENERGY dataset.



Conclusion and future work

Based on both the grid search results and Diebold-Mariano test results, we found that the all deep learning and machine learning models outperformed the incumbent manual technique. Furthermore, the GRU model (GRU-1-30) outperformed the basic RNN and LSTM models in RMSE (0.0305), MAPE (4.33%), and MAE (0.0305), in a very short inference time (0.7058 seconds).

This paper highlights the potential of deep learning and ARIMA in energy-intensive manufacturing. The adoption of deep learning, like all data science technologies, requires overcoming human, organisational and technological challenges however against intense rivalry, firms may not have a choice.

For near real-time prediction, very short-term load forecasting (VSTLF) may be needed. In such use cases, rapid training times will be required. Furthermore, medium-term load forecasting (MTLF) may prove fruitful, deep learning training models may need to be augmented with historic trend data to account for longer seasonal cycles or predictable events. MTLF may enable new use cases including switches to the more sustainable or lower cost power supply. This may require ensemble solutions and is worthy of exploration.

The entire article has been published on <https://doi.org/10.3390/a13110274>.

References

1. [International Energy Outlook 2019](#). EIA. Retrieved 2020-11-19
2. Conti, J., Holtberg, P., Diefenderfer, J., LaRose, A., Turnure, J. T., & Westfall, L. . International energy outlook 2016 with projections to 2040; USDOE Energy Information Administration (EIA), Eds.; Office of Energy Analysis: DC (United States), 2016; pp. 113.
3. Seunghyoung Ryu; Jaekoo Noh; Hongseok Kim; Deep neural network based demand side short term load forecasting. *2016 IEEE International Conference on Smart Grid Communications (SmartGridComm) 2017*, 10, 308-313, [10.1109/smartgridcomm.2016.7778779](#).
4. Chong Chen; Ying Liu; Maneesh Kumar; Jian Qin; Energy Consumption Modelling Using Deep Learning Technique — A Case Study of EAF. *Procedia CIRP 2017*, 72, 1063-1068, [10.1016/j.procir.2018.03.095](#).
5. BERRIEL, Rodrigo F. et al. Monthly energy consumption forecast: A deep learning approach. In: 2017 International Joint Conference on Neural Networks (IJCNN). IEEE, 2017. p. 4283-4290.
6. Onat Güngör; Barış Akşanlı; Reyhan Aydoğan; Algorithm selection and combining multiple learners for residential energy prediction. *Future Generation Computer Systems 2019*, 99, 391-400, [10.1016/j.future.2019.04.018](#).
7. Weicong Kong; Zhao Yang Dong; Youwei Jia; David J. Hill; Yan Xu; Yuan Zhang; Short-Term Residential Load Forecasting Based on LSTM Recurrent Neural Network. *IEEE Transactions on Smart Grid 2017*, 10, 841-851, [10.1109/tsg.2017.2753802](#).
8. Jesus Lago; Fjo De Ridder; B. De Schutter; Forecasting spot electricity prices: Deep learning approaches and empirical comparison of traditional algorithms. *Applied Energy 2018*, 221, 386-405, [10.1016/j.apenergy.2018.02.069](#).
9. Bergstra, James and Bengio, Yoshua; Random Search for Hyper-Parameter Optimization. *J. Mach. Learn. Res.* **2012**, 13, 281–305, .
10. Jun-Mao Liao; Ming-Jui Chang; Luh-Maan Chang; Prediction of Air-Conditioning Energy Consumption in R&D Building Using Multiple Machine Learning Techniques. *Energies 2020*, 13, 1847, [10.3390/en13071847](#).
11. Heesung Yoon; Yongcheol Kim; Kyoochul Ha; Soo-Hyoung Lee; Gee-Pyo Kim; Comparative Evaluation of ANN- and SVM-Time Series Models for Predicting Freshwater-Saltwater Interface Fluctuations. *Water 2017*, 9, 323, [10.3390/w9050323](#).
12. Kadir Kavaklioglu; Modeling and prediction of Turkey's electricity consumption using Support Vector Regression. *Applied Energy 2010*, 88, 368-375, [10.1016/j.apenergy.2010.07.021](#).
13. R. Samsudin; A. Shabri; P. Saad; A Comparison of Time Series Forecasting using Support Vector Machine and Artificial Neural Network Model. *Journal of Applied Sciences 2010*, 10, 950-958, [10.3923/jas.2010.950.958](#).
14. YOUNG, Steven R. et al. Optimizing deep learning hyper-parameters through an evolutionary algorithm. In: Proceedings of the Workshop on Machine Learning in High-Performance Computing Environments. 2015. p. 1-5.
15. Francis X. Diebold; Comparing Predictive Accuracy, Twenty Years Later: A Personal Perspective on the Use and Abuse of Diebold–Mariano Tests. *Journal of Business & Economic Statistics 2015*, 33, 1-1, [10.1080/07350015.2014.983236](#).
16. Francis X. Diebold; Comparing Predictive Accuracy, Twenty Years Later: A Personal Perspective on the Use and Abuse of Diebold–Mariano Tests. *Journal of Business & Economic Statistics 2015*, 33, 1-1, [10.1080/07350015.2014.983236](#).

Retrieved from <https://encyclopedia.pub/entry/history/show/10242>