Deep Neural Network Regression as a Component of a Forecast Ensemble

Gregory D. Merkel
Department of Electrical and Computer Engineering
Marquette University
Milwaukee, WI

Richard J. Povinelli
Department of Electrical and Computer Engineering
Marquette University
Milwaukee, WI

Ronald H. Brown
Department of Electrical and Computer Engineering
Marquette University
Milwaukee, WI

***Abstract*— Marquette University's GasDay Project specializes in short-term load forecasting of natural gas demand. Traditionally, this forecasting is done using artificial neural networks and linear regression. This paper examines the viability of using DNNs as component models in the GasDay ensemble. The ensemble of interest is evaluated using weighted MAPE on 88 natural gas data sets and compared to the current GasDay ensemble. The DNN-enhanced ensemble performs better on this metric than the current ensemble.**

***Keywords*— Deep learning, deep neural network, regression, short-term load forecasting, time-series analysis, ensemble forecasting.**

1. Introduction

Short-term load forecasting of natural gas is important for several reasons. Many purchasing and operation decisions are made by natural gas utilities using these forecasts, and there are high costs to utilities and their customers if the forecast is inaccurate.

Figure 1 shows a time series of daily natural gas load. The data shown in Figure 1 is from several metropolitan areas in northern United States.

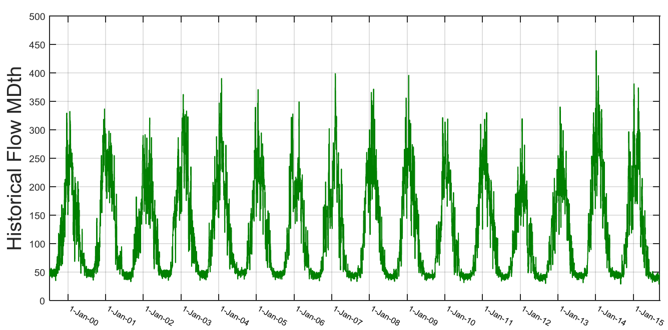

Figure 1: Natural gas load versus time for weighted combination area (1999-2015).

As can be seen in Figure 2, the daily natural gas data has a roughly linear relationship with temperature. Because of this, traditionally, linear regression (LR) and autoregressive integrated moving average (ARIMA) models are used for short-term load forecasting of natural gas [1]. These traditional models perform well on the linear stationary time-series, and thus have been used successfully in the short-term load forecasting problem [2].

Unfortunately, gas demand contains nonlinearities. Some of these are easy for a proficient forecaster to capture using an LR model by, for instance, using heating degree days as an input instead of temperature. However, natural gas demand also contains many nonlinearities that either are difficult for forecasters to glean from the data themselves or cannot be easily captured with LR or ARIMA models. One of the forecasting community's answers to this problem for many years has been using artificial neural networks (ANN) as ANNs are able to capture nonlinearities [1], [3], [4].

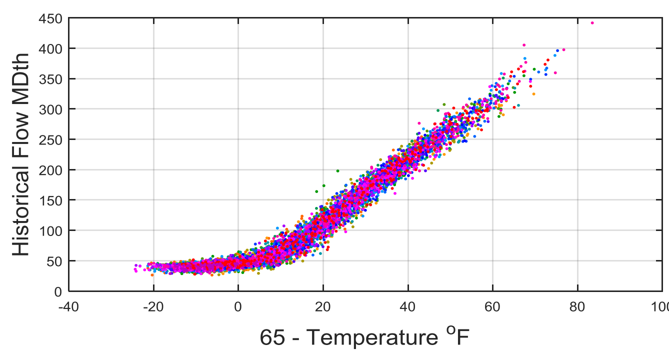

Figure 2: Natural gas load versus temperature for weighted combination area (1999-2015).

Recently, the machine learning community has successful replaced nonlinear models, including ANNs with deep neural networks (DNN) [5]. Längkvist discusses trends in DNNs and ANNs for many time-series problems including video analysis, motion capture, speech processing, and music recognition [5]. DNNs are not simply large ANNs. The term DNN refers to neural networks trained using special techniques that allow them to be much larger than shallow ANNs. Traditional ANNs are trained using gradient descent. Large neural networks trained by gradient descent also are prone to overfitting data sets. The DNN technique used in this paper avoids both of these problems by using a restricted Boltzmann machine training algorithm to “pre-train” the model, followed by a few epochs of gradient descent [6]. More information on RBMs and how to train them can be found in [6], [7] and [8].

In a related paper, we discuss how well the deep neural network performs on its own as a forecaster of natural gas demand and found that it perform better than other individual models [9]. In this paper, we will be using the DNN as a component of an ensemble of models. The ensemble will be evaluated with and without the DNN component. This experiment will be run on 73 different natural gas data sets and it will be shown that on average, the addition of the DNN component results in better forecast.

2. The GasDay ensemble: dynamic post processor

The GasDay dynamic post processor is an ensembling method that adjusts the weights given to each component forecast based on its recent performance. This is described in detail in [10]. If one model consistently outperforms the other then the dynamic post processor will weight it more heavily. Additionally, if both models are consistently high or low, which would indicate that the natural gas system is changing, the dynamic post processor will adjust accordingly.

3. The experiment

This experiment is run using the entire 2015-2016 heating seasons worth of forecasts for each of 67 operating areas. The results are also compared on all days as well as unusual days. It is expected that for most areas the ensembles will perform better with the additional deep neural network component.

The metric we will use to evaluate the performance of the ensemble models will be weighted mean absolute percent error (WMAPE):

$$WMAPE = 100 \times \frac{\sum_{n=1}^{N} \left| \hat{s}(n) - s(n) \right|}{\sum_{n=1}^{N} s(n)}$$

This metric was chosen because of the value it provides compared to other common metric. Root mean squared error (RMSE) is a powerful metric for short term natural gas load forecasting, because it naturally places more value on days with higher load, which are more important to natural gas utilities. Unfortunately, RMSE is dependent on the magnitude of the system and cannot be used to compare the performance of the technique between different systems. Mean absolute percent error (MAPE), on the other hand, is not magnitude dependent and can be used to compare performance between models of different systems. Conversely, MAPE naturally places higher value on days with lower load, which are less important to natural gas utilities. Therefore, we will use WMAPE as it avoids the problems of both of these metrics.

Another important criteria when evaluating a short term natural gas demand forecasting model is its performance on high demand and traditionally hard to forecast days. The highest demand days for natural gas are the coldest days when the heating load is the greatest. Prices are higher when the weather is colder so it is more important that a forecast be correct on these days. Other important days for forecasting natural gas are days much colder or warmer than the previous day, the first warm days after a heating season, and the first cold days leading into a heating season. A forecasting model that performs well on these days can help a local distribution company avoid penalties and having to buy gas on the spot market. As such, all of the models are evaluated on these day types individually as well as on all days.

4. Results

The first results, from comparing the GDDPP to the GDDPP with an additional DNN component are as expected. This can be seen in Figure 3. The p-values for unusual day types are included in Table 1. These results are clear, conclusive, and as expected.

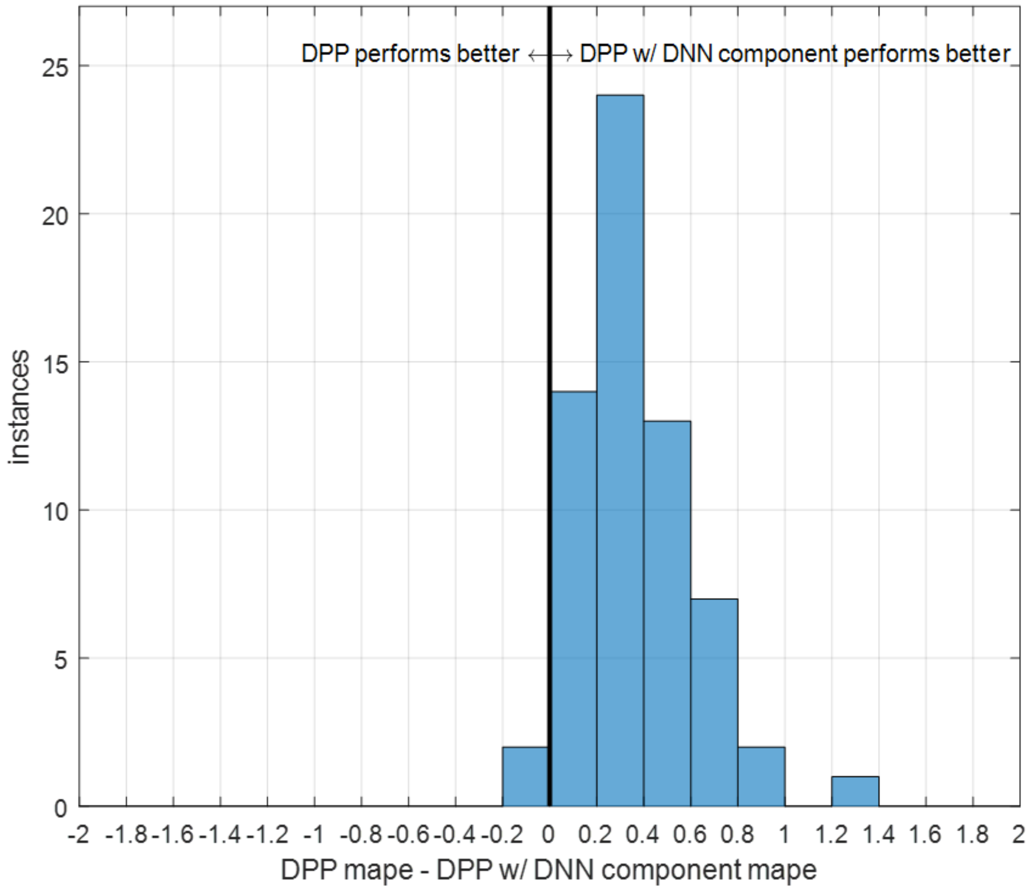

Figure 3: This is a histogram of the differences in WMAPE between the GasDay ensemble with the DNN component and the current GasDay ensemble. Values on the left of the thick line at 0 indicate areas where the current GasDay ensemble performs better. Those on the right indicate areas where the GasDay ensemble with the DNN component performs better.

Table 1: Right-tailed t-test comparing the current GasDay ensemble to the GasDay ensemble with the DNN component on each unusual day type. Values less than 0.05 indicate unusual day types on which the GasDay ensemble with the DNN component performs significantly better.

Unusual Day Type	**p-value**
All Days	**8.30×10^{-19}**
Coldest Days	**1.59×10^{-8}**
Colder Than Normal Heating Days	**2.58×10^{-5}**
Warmer Than Normal Heating Days	**1.61×10^{-9}**
Windiest Heating Days	**1.81×10^{-7}**
First Heating Days	**4.65×10^{-6}**
First Non-Heating Days	**3.65×10^{-3}**

5. Conclusions

It is also concluded from this paper that including a DNN component provides significant value to many areas and some value to almost all areas. Additionally, this improvement is most apparent in the coldest days and other important to forecast day types. This work is elaborated on in [11].

6. References

[1] S. R. Vitullo *et al*, "Mathematical models for natural gas forecasting," *Canadian Applied Mathematics Quarterly*, vol. 17, *(7)*, pp. 807-827, 2009.

[2] T. Haida and S. Muto, "Regression based peak load forecasting using a transformation technique," *IEEE Trans. Power Syst.*, vol. 9, *(4)*, pp. 1788-1794, 1994.

[3] K. Hornik, M. Stinchcombe and H. White, "Multilayer feedforward networks are universal approximators," *Neural Networks*, vol. 2, *(5)*, pp. 359-366, 1989.

[4] D. C. Park *et al*, "Electric load forecasting using an artificial neural network," *IEEE Trans. Power Syst.*, vol. 6, *(2)*, pp. 442-449, 1991.

[5] M. Längkvist, L. Karlsson and A. Loutf, "A review of unsupervised feature learning and deep learning for time-series modeling," *Pattern Recognition Letters*, vol. 42, pp. 11-24, 2014.

[6] Geoffrey E. Hinton, Simon Osindero and Yee-Whye Teh, "A Fast Learning Algorithm for Deep Belief Nets," *Neural Computation*, vol. 18, *(7)*, pp. 1527-1554, 2006.

[7] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks." in *Aistats*, pp. 249-256. 2010.

[8] G. Hinton, "A practical guide to training restricted Boltzmann machines," *Momentum*, vol. 9, *(1)*, pp. 926, 2010.

[9] G. Merkel, R. Povinelli and R. Brown, "Deep Neural Network Regression for Short-Term Load Forecasting of Natural Gas" *37th Annual International Symposium on Forecasting*, 2017.

[10] R. Brown *et al*, "Improving Daily Natural Gas Forecasting by Tracking and Combining Models," *37th Annual International Symposium on Forecasting*, 2017.

[11] G. Merkel, "Deep Neural Networks as Time Series Forecasters of Energy Demand," *ProQuest Dissertations Publishing*, 2017.