



A Comparison of Wind Power and Load Forecasting Error Distributions

Preprint

Bri-Mathias Hodge, Anthony Florita, Kirsten Orwig,
Debra Lew, and Michael Milligan

*Presented at the 2012 World Renewable Energy Forum
Denver, Colorado
May 13-17, 2012*

NREL is a national laboratory of the U.S. Department of Energy, Office of Energy Efficiency & Renewable Energy, operated by the Alliance for Sustainable Energy, LLC.

Conference Paper
NREL/CP-5500-54384
July 2012

Contract No. DE-AC36-08GO28308

NOTICE

The submitted manuscript has been offered by an employee of the Alliance for Sustainable Energy, LLC (Alliance), a contractor of the US Government under Contract No. DE-AC36-08GO28308. Accordingly, the US Government and Alliance retain a nonexclusive royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for US Government purposes.

This report was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or any agency thereof.

Available electronically at <http://www.osti.gov/bridge>

Available for a processing fee to U.S. Department of Energy and its contractors, in paper, from:

U.S. Department of Energy
Office of Scientific and Technical Information

P.O. Box 62
Oak Ridge, TN 37831-0062
phone: 865.576.8401
fax: 865.576.5728
email: <mailto:reports@adonis.osti.gov>

Available for sale to the public, in paper, from:

U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
phone: 800.553.6847
fax: 703.605.6900
email: orders@ntis.fedworld.gov
online ordering: <http://www.ntis.gov/help/ordermethods.aspx>

Cover Photos: (left to right) PIX 16416, PIX 17423, PIX 16560, PIX 17613, PIX 17436, PIX 17721



Printed on paper containing at least 50% wastepaper, including 10% post consumer waste

A Comparison of Wind Power and Load Forecasting Error Distributions

Bri-Mathias Hodge
Anthony Florita
Kirsten Orwig
Debra Lew
Michael Milligan
National Renewable Energy Laboratory
Transmission and Grid Integration Group
1617 Cole Blvd.
Golden, CO 80401
e-mail: bri-mathias.hodge@nrel.gov
anthony.florita@nrel.gov
kirsten.orwig@nrel.gov
debra.lew@nrel.gov
michael.milligan@nrel.gov

ABSTRACT

The introduction of large amounts of variable and uncertain power sources, such as wind power, into the electricity grid presents a number of challenges for system operations. One issue involves the uncertainty associated with scheduling power that wind will supply in future timeframes. However, this is not an entirely new challenge; load is also variable and uncertain, and is strongly influenced by weather patterns. In this work we make a comparison between the day-ahead forecasting errors encountered in wind power forecasting and load forecasting. The study examines the distribution of errors from operational forecasting systems in two different Independent System Operator (ISO) regions for both wind power and load forecasts at the day-ahead timeframe. The day-ahead timescale is critical in power system operations because it serves the unit commitment function for slow-starting conventional generators.

1. INTRODUCTION

Increasing levels of power are being provided by variable and uncertain power sources, such as wind power, leading to concerns about the viability of current power system operations practices in future high-penetration scenarios. However, it is important to realize that power system operations have been developed to meet variable and uncertain load. The variable and uncertain nature of load is one of the reasons for the predominance of the day-ahead unit commitment followed by hourly (or sub-hourly)

economic dispatch model of unit scheduling. An estimate of the power requirements must be made so that slow-starting thermal units are available to meet the anticipated load, but an update to the load forecast in the dispatch timeframe allows for more accurate scheduling, once some of the uncertainty has been resolved. The same procedures may be utilized to aid in the scheduling of wind power output, and increasingly accurate wind power forecasts may be incorporated at shorter timescales, such as six and four hours ahead. However, in order for these practices to incorporate increased amounts of wind power, the nature of the forecast error distribution should be well known, and accurately modeled. Accurately representing the forecast error characteristics is especially important in wind integration studies, where an inaccurate modeling of the forecast errors may overestimate or underestimate the costs of system operations in scenarios with large amounts of wind.

There have been a number of large-scale studies of wind power integration in the United States in the last ten years (1-6). These studies all recognize the capability of wind power forecasting to reduce the operating costs in systems with large wind penetration; however, they often make simplifying assumptions about the nature of wind power forecasting due to the unavailability of actual data on such large penetration systems. Previous studies have often assumed that wind forecasting errors follow a Gaussian distribution (7, 8). While most wind power integration studies assume that load is perfectly forecasted to isolate the costs of wind power forecasting, those that do include load forecasting may assume

that load forecast errors also follow a Gaussian distribution (7).

In this work we examine the statistical nature of wind power and load forecasting error distributions from operational systems, to better inform wind integration studies. The day-ahead forecast horizon will be the focus of this study due to its critical nature in the unit commitment process. Because power systems operations are accustomed to dealing with load forecasting errors, we compare them with those seen in wind power forecasting to see if lessons may be learned that will help enable the economic incorporation of larger amounts of wind power.

2. METHODS AND DATA

In this section we provide background on the statistical distributions discussed in the study. In addition, we present the datasets that will be examined in further detail. The analysis was performed using the *R* statistical software environment (9), using functions from the *HyperbolicDist* (10) and *MASS* (11) packages.

2.1 STATISTICAL BACKGROUND

Because we will be discussing statistical distributions from a perspective not normally utilized in the power systems community, we will provide some background and definitions of terminology that will appear in what follows. The Gaussian distribution that is commonly assumed to model both load and wind forecast errors may be fully described by mean and standard deviation values. While this simplicity is useful, it masks instances where very dissimilar distributions may have similar means and standard deviations. These two parameters correspond to the first two statistical moments; however, more information about the shape of the distribution can be extracted by examining the third and fourth moments, skewness and kurtosis, respectively. Skewness can be thought of as the symmetry of the distribution, while kurtosis is a measure of the relative weighting of the peak and tails of the distribution. A distribution with a large kurtosis value is known as leptokurtic, while one with a small kurtosis value is known as platykurtic. In what follows, kurtosis will refer to excess kurtosis, i.e. the kurtosis above that of the Gaussian distribution. The examination of these two additional moments allows for a more accurate representation of the operational errors witnessed in the datasets described in the next section.

2.2 DATASETS

The load data and forecasts examined are from two different independent system operators (ISOs) in the

United States, representing the states of California (CAISO) and New York (NYISO). The CAISO day-ahead load forecasts and actuals were taken from the CAISO OASIS system (12). The dataset used is hourly averaged load from 2011. The maximum load during the period is 45,569 MW, the minimum load 18,605 MW, and the mean load is 26,297 MW. The NYISO hourly load forecasts and actuals from 2010 were obtained from the NYISO website (13). The mean load in the period under consideration was 18,664 MW, with a maximum of 33,452 MW and a minimum of 11,859 MW.

The wind forecasting data that was examined also comes from two different ISOs in the United States, California and Texas (ERCOT). The ERCOT interconnection dataset covers a 13-month period and represents the combined output of approximately 9,000 MW of installed wind power capacity. The day-ahead forecast is made at 16:00 the previous day. The CAISO data is the aggregated wind power output of 16 different wind plants over a one year period, with a total capacity of approximately 940 MW. These forecasts are produced at 05:30 the previous day for each hour of the next day.

3. DAY-AHEAD LOAD FORECASTS

The process of scheduling generating units to meet expected demand is known as the Unit Commitment and Economic Dispatch (UCED) process. Due to the fact that large thermal units often have long start-up times, the unit commitment decisions (i.e. whether a unit will be on or off during the specified period) have traditionally been performed in the day-ahead timeframe. Because load is variable and uncertain, day-ahead forecasts are required to ensure that sufficient generating capacity is available to meet the expected load. This forecast can be a critical factor in ensuring near-optimal system operations. For example, if the load forecast is significantly lower than the realized load, too little baseload capacity may be operational at the needed time, and fast-starting, more expensive units will be required to fulfill the load. The economic cost of this deviation from the optimal dispatch stack can be directly attributed to the inaccurate load forecast. The main cause of additional costs in the day-ahead timeframe is the commitment or de-commitment of large thermal units. This fact, along with the ability of the economic dispatch process to handle smaller forecast errors, means that large magnitude, but relatively rare, forecasting errors are the most important for system operations. For this reason, accurately modeling the tails of the forecast error distribution is critical for understanding their impacts on system operations. Therefore, an accurate comparison between the forecast error distributions observed in real system operations and those commonly assumed in power system operation studies has the potential to increase the fidelity of the study outcomes.

Load generally follows a familiar pattern, reaching its peak during the day and into the evening, with a nighttime nadir. In most of the United States, temperatures tend to significantly influence the load shape, with hot summer days requiring more air conditioning during the day, and cold winter nights increasing the minimum load. Figure 1 shows a week's worth of hourly load requirements, plus the day-ahead load forecasts from the CAISO system. Weekend days tend to have lower load than weekdays, as seen in the two lower peaks of the figure 1. Because of the load forecast's strong dependence on the temperature forecast, significant load forecast errors are often autocorrelated. As seen in the load shape, if the load is over-forecast for one hour during the morning ramp, the error tends to persist throughout the day.

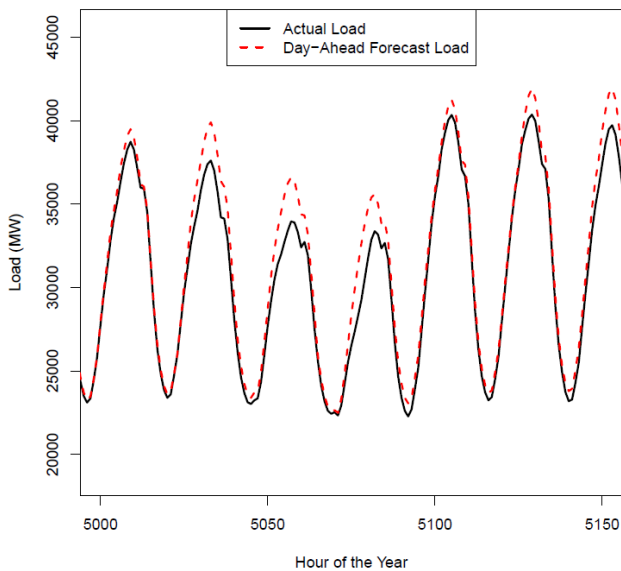


Fig. 1: Hourly load and forecast load values for one week in the CAISO system.

3.1 GAUSSIAN COMPARISON

The Gaussian distribution is often assumed for distributions where many phenomena are at play, often with an invocation of the central limit theorem. However, this assumption should be checked against real data before the results of models incorporating the assumption may be verified and validated.

Figure 2 displays a histogram of the day-ahead load forecasting errors from the CAISO system, normalized by the yearly average load. One important thing to notice is the long right tail of the distribution, highlighted by the numerous errors between 10% and 20% of the average load. The distribution also shows a significant positive skewness and is leptokurtic (i.e. narrower, more

prominent peak and fatter tails) when compared to a Gaussian distribution with the same mean and standard deviation as the observed errors. While the assumed Gaussian distribution can match the mean and standard deviation of the observed errors, it does not represent the skewness and kurtosis observed, creating a significantly different distribution shape. One implication of this is that the tails and the peak of the distribution are underrepresented.

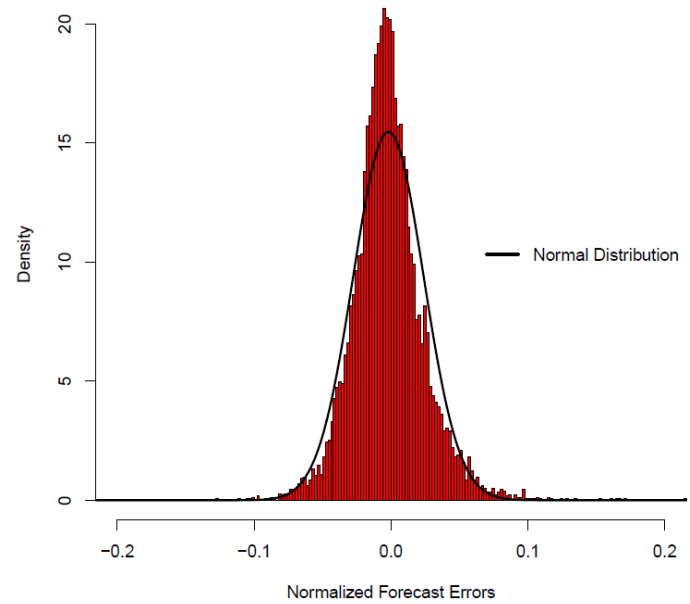


Fig. 2: A histogram of the distribution of day-ahead load forecasting errors for the CAISO system, normalized by the yearly average load. $\mu = -0.002$; $\sigma = 0.026$; $\gamma = 0.715$; $\kappa = 4.725$.

While the histogram and the statistical moment calculations seem to indicate that the load forecast error distribution is poorly represented by the Gaussian distribution, additional assurance is provided by a normal quantile-quantile (Q-Q) plot. Figure 3 shows a Q-Q plot of the CAISO day-ahead load forecast errors. The line in the figure goes through the first and third quantiles, and should pass through most of the data points if the observed distribution is Gaussian. However, we observe significant deviations, especially in the tails of the distribution.




Fig. 3: A normal quantile-quantile plot of the distribution of day-ahead load forecasting errors for the CAISO system, normalized by the yearly average load. The line runs through the first and third quantiles of the observed distribution.

3.2 DISTRIBUTION MODELING

Having established that the Gaussian distribution is a poor fit for the observed day-ahead load forecast errors, we now propose an alternative distribution. The hyperbolic distribution is proposed to more accurately model the semi-heavy tails observed in the load forecast error distributions. Essentially, the hyperbolic distribution is a Laplace, or double exponential, distribution that allows asymmetry to capture skewness. Figure 4 shows the normalized day-ahead load forecast errors for the NYISO system. Also included are a Gaussian distribution with the same mean and standard deviation, and a hyperbolic distribution fit to the observed errors. The distribution has a significant bias, larger standard deviation than the CAISO errors, positive skew, and is leptokurtic. A comparison between the two model distributions shows that the hyperbolic distribution does a better job than the Gaussian distribution in representing the pronounced peak and slimmer shoulders of the observed distribution. Additionally, the hyperbolic distribution does do a slightly better job than the Gaussian distribution in representing the semi-heavy tails; however, they are still underrepresented in the model distribution. This is also seen in figure 5 where the hyperbolic distribution covers the observed distribution for most of the plot, the only exceptions being small deviations in the tails, and a slight mismatch in the right shoulder of the distribution.




Fig. 4: A histogram of the distribution of day-ahead load forecasting errors for the NYISO system, normalized by the yearly average load. $\mu = -0.024$; $\sigma = 0.036$; $\gamma = 0.379$; $\kappa = 3.799$. A Gaussian distribution with the same mean and standard deviation is shown along with a hyperbolic distribution fit to the observed data.




Fig. 5: A cumulative distribution plot of the NYISO day-ahead load forecast errors, along with the Gaussian and hyperbolic model distributions.

4. DAY-AHEAD WIND POWER FORECASTS

In areas where wind power generators participate in the day-ahead market, the forecasted wind power during the next day is an important variable in deciding what other generators need to be made available. Because wind does not have any

fuel costs associated with its production, it tends to have a lower bid cost than most thermal generators. Thus, the forecasted wind production reduces the amount of thermal generation that will be necessary to meet the forecasted load. However, similar to load forecast errors, large errors in the day-ahead wind forecast can have economic consequences in the unit commitment process.

Although wind power output may display some daily and seasonal characteristics, it follows much less regular patterns than does load. Generally speaking, wind power output tends to be higher during nighttime periods, though times can be found when there is no output. This, of course, makes wind power more difficult to forecast than load. The greater range of variability experienced by even aggregations of wind power plants also adds to the difficulty of forecasting its output at the day-ahead timescale. Figure 6 shows one week worth of wind power output and day-ahead wind power forecasts for the ERCOT system. Over this short period the total wind output varies very significantly, from almost 8,000 MW to near zero output. From a first glance at figure 6, the wind forecast does a fairly good job of anticipating large changes in wind power. However, during a period of large variability, even slight errors like the slight forecast phase error during the large down ramp shown, can have large consequences on system operations. This phase error creates an hourly forecast error of approximately 2,000 MW that must be compensated for in system operations. It is important to note that the error shown is at the day-ahead timescale, and so only impacts unit commitment decisions. Updated forecasts may be incorporated into the economic dispatch process and could eliminate or reduce the error before the dispatch timeframe. The only possible cost then associated with the error would be the difference between supplying that energy with a mid-merit unit that may need to be started and a baseload unit that might have otherwise supplied the required energy. However, even this may not always be the case for large errors, and depends significantly on the current state of the system when the error occurs.

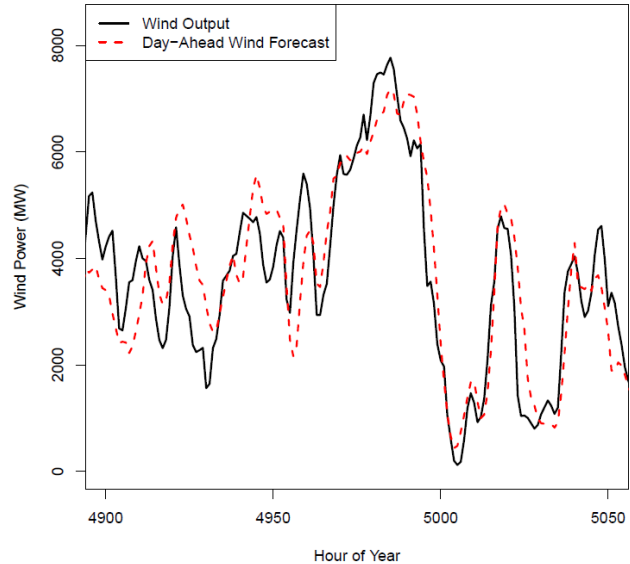


Fig. 6: Hourly wind power output and forecast wind power values for one week in the ERCOT system.

4.1 GAUSSIAN COMPARISON

Though the current level of wind energy penetration in the United States is fairly low, many studies have been performed on the impact of future high penetrations of wind energy on system operations. Because wind is a much larger component of the generation fleet in these studies, the impact of wind forecasting errors is much greater. For this reason it is critical that the distribution of wind power forecast errors is accurately represented in these studies. Perhaps the most important component is correctly characterizing the tails of the distribution, as they represent the largest forecast errors that will be seen in the study.

As with our examination of the load forecasting errors, normal Q-Q plots will provide a means by which the observed wind power forecast error distributions may be compared with the Gaussian distribution. Figure 7 shows a normal Q-Q plot of the ERCOT day-ahead wind power forecasting errors, with significant deviations from normality in the tails of the distribution. This is also the case, though to a slightly larger extent, for the CAISO day-ahead wind power forecasting errors, displayed in figure 8. Previous work (14) seems to indicate that the larger the geographic diversity of the system under study, the more accurate the Gaussian representation will be, hence the larger deviations from the CAISO system, which is significantly smaller in terms of wind power capacity.




Fig. 7: A normal quantile-quantile plot of the distribution of day-ahead wind power forecasting errors for the ERCOT system, normalized by the installed wind power capacity.




Fig. 8: A normal quantile-quantile plot of the distribution of day-ahead wind power forecasting errors for the CAISO system, normalized by the installed wind power capacity.

4.2 DISTRIBUTION MODELING

After demonstrating the poor representation that the Gaussian distribution provides for the day-ahead wind forecast errors distribution, we examine the applicability of the hyperbolic distribution. Fig. 9 shows a histogram of the day-ahead wind power forecasting errors for the

CAISO system, along with a Gaussian distribution with the same mean and standard deviation, as well as a hyperbolic distribution fit to the data. The hyperbolic distribution more accurately represents the leptokurtic nature and skewness of the observed distribution, and provides a more accurate representation of the semi-heavy distribution tails, than does the Gaussian distribution.

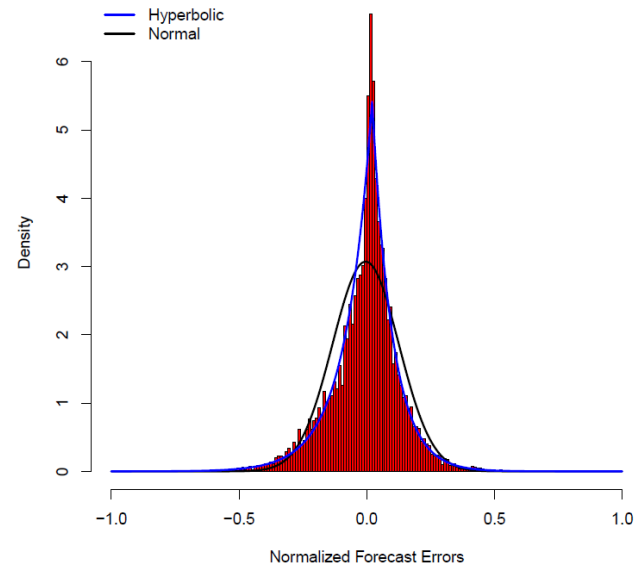


Fig. 9: A histogram of the distribution of day-ahead wind power forecasting errors for the CAISO system, normalized by the installed wind capacity. $\mu = -0.004$; $\sigma = 0.130$; $\gamma = -0.393$; $\kappa = 1.503$. A Gaussian distribution with the same mean and standard deviation is shown along with a hyperbolic distribution fit to the observed data.

The superior representation of the hyperbolic distribution is also seen in the cumulative distribution plot, shown in figure 10. The hyperbolic distribution mirrors the observed errors very closely, with only small deviations in the left shoulder of the distribution. On the other hand, the Gaussian distribution shows significant deviations, most clearly in the shoulders and tails of the distribution.

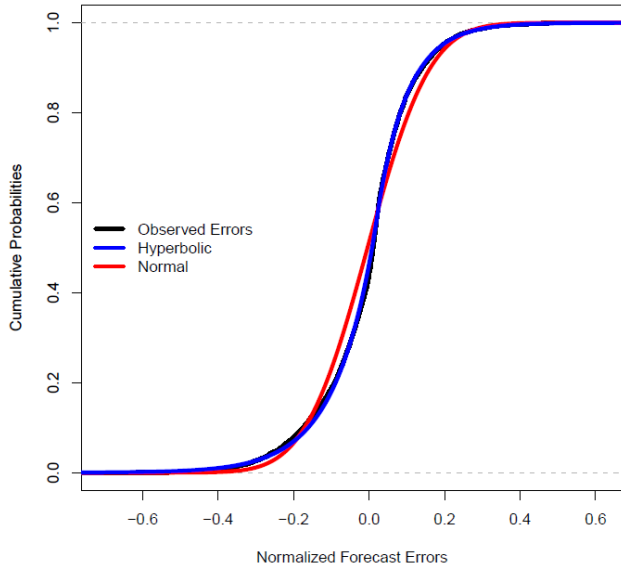


Fig. 10: A cumulative distribution plot of the CAISO day-ahead wind power forecast errors, along with the Gaussian and hyperbolic model distributions.

Fig. 11 shows a histogram of the day-ahead wind power forecasting errors for the ERCOT system, along with a Gaussian distribution with the same mean and standard deviation, and a hyperbolic distribution fit to the observed errors. The hyperbolic distribution provides a better fit for the pronounced peak and slimmer shoulders seen in this leptokurtic distribution.

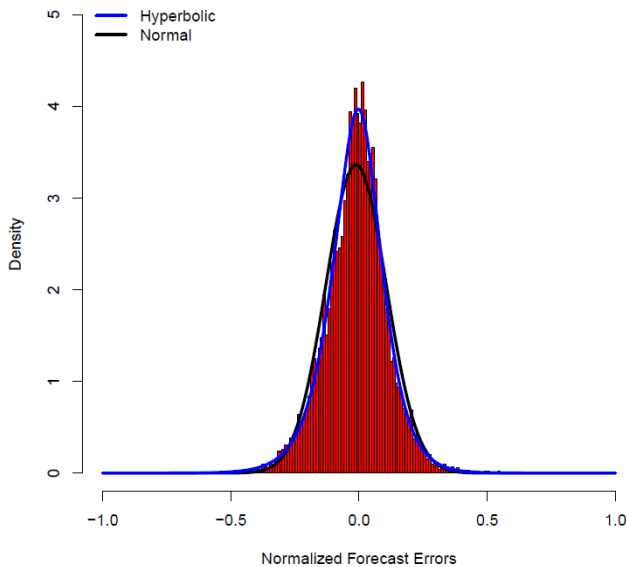


Fig. 11: A histogram of the distribution of day-ahead wind power forecasting errors for the ERCOT system, normalized by the installed wind capacity. $\mu = -0.012$; $\sigma = 0.119$; $\gamma = -0.062$; $\kappa = 1.030$. A Gaussian distribution

with the same mean and standard deviation is shown along with a hyperbolic distribution fit to the observed data.

5. COMPARISON

After examining the day-ahead load and wind power forecasting error distributions, it is clear that there are some noticeable similarities, but also important differences. Perhaps the most important similarity between the wind power and load forecasting errors is that they are leptokurtic distributions at the geographic scales of a single ISO, and are thus poorly represented by the Gaussian distribution. The load forecasting errors have larger kurtosis values, though kurtosis for wind power forecasting errors are very strongly dependent on the timescale of the forecast (15). The pronounced peaks that are one feature of the leptokurtic distribution are the result of a large number of very small forecast errors, indicating some forecasting skill. The regular daily pattern of load helps to explain the more leptokurtic distributions observed for load than for wind power forecasting. Another similarity between the observed error distributions is the applicability of the hyperbolic distribution as a model distribution. This distribution is able to more accurately capture the pronounced peak, slim shoulders, and semi-heavy tails seen in both wind power and load forecasting error distributions.

An important difference between the load and wind power forecast error distributions is the range of observed values, in terms of their normalized value. The load forecasting errors shown have a smaller range of values than the wind power forecasting errors; however, this is also a function of the values chosen for normalization. Although the total wind power capacity seems a clear choice for normalizing the wind power forecasting error values, a number of different values may be used for the load errors. In this work the average yearly load is used, though a case could also be made for the yearly maximum load, the yearly minimum load, or a number of other values. In terms of power, the load forecast errors are much larger, but this is to be expected because the total installed wind capacity is much less than the maximum load in current systems. It is also important to note that load follows more distinct daily patterns, and has a longer history of being forecast. It is not unreasonable to think that wind power forecasting could improve significantly as its importance in power systems operations increases.

Another difference between the two forecasting error distributions lies in the geographic diversity of the resource they are attempting to forecast. Load is a function of the weather and usage patterns observed at millions of households and businesses, spread across wide geographic areas for all of the ISOs considered. This dispersion creates a smoothing of the overall load profile that aids in its forecasting. At current penetration levels, wind power forecasting at the ISO level is

the aggregation of tens of wind plants, or hundreds of individual turbines. With larger penetrations, and more geographic diversity, the same smoothing trends will tend to decrease the total spread of forecasting errors, because there tends to be less correlation across larger geographic domains (16).

6. CONCLUSION

In this work we have analyzed, modeled, and compared the error distributions that arise from operational day-ahead wind power and load forecasting systems currently in use in three different ISOs in the United States. Although increasing amounts of wind generation incorporate additional variability and uncertainty into power system operations, systems without wind power already have large amounts of variability and uncertainty because of load. The shape of day-ahead wind power forecasting errors is similar to those of day-ahead load forecasts, and the goal of both forecasts should be to limit the amount of large errors, so as to enable efficient system operations.

7. ACKNOWLEDGMENTS

The authors would like to thank David Maggio at ERCOT and James Blatchford at CAISO for providing the wind power forecasting data used in this study.

8. REFERENCES

- (1) NYISO, Growing Wind - Final Report of the NYISO 2010 Wind Generation Study, 2010
- (2) EnerNex, Xcel Energy and the Minnesota Department of Commerce Wind Integration Study - Final Report, 2004
- (3) EnerNex, Wind Integration Study for Public Service Company of Colorado, 2006
- (4) EnerNex, Eastern Wind Integration and Transmission Study, 2010
- (5) GE, Analysis of Wind Generation Impact on ERCOT Ancillary Services Requirements, 2008
- (6) GE, Western Wind and Solar Integration Study, 2010
- (7) Ortega-Vazquez, M. and Kirschen, D., Estimating Spinning Reserve Requirements in Systems With Significant Wind Power Generation Penetration, IEEE Transactions on Power Systems, 2009
- (8) Doherty, R. and O'Malley, M., A New Approach to Quantify Reserve Demand in Systems with Significant Installed Wind Capacity, IEEE Transactions on Power Systems, 2005
- (9) R: A Language and Environment for Statistical Computing, 2010
- (10) Scott, D., HyperbolicDist: The hyperbolic distribution, 2009

(11) Venables, W.N. and Ripley, B.D., Modern Applied Statistics with S. Springer, 2002

(12) CAISO, CAISO Open Access Same-time Information System (OASIS), 2011

(13) NYISO, Market & Operation Data, 2011

(14) Hodge, B.-M. and Milligan, M., Wind Power Forecasting Error Distributions over Multiple Timescales, IEEE Power & Energy Society General Meeting, 2011

(15) Hodge, B.-M., Ela, E. and Milligan, M., The Distribution of Wind Power Forecast Errors from Operational Systems, 10th International Workshop on Large-scale Integration of Wind Power, 2011

(16) Wan, Y.-H., A Primer on Wind Power for Utility Applications, 2005