Evaluation of Probabilistic Models of Wind Plant Power Output Characteristics

Henry Louie
Department of Electrical and Computer Engineering
Seattle University
Seattle, Washington, USA 98122–1090
Email: louieh@seattleu.edu

***Abstract*—The power output by weather-driven renewable resources such as wind energy conversion systems can be appropriately described as being stochastic. To manage these resources, probabilistic models of wind power are being increasingly employed by power system stakeholders in applications such as stochastic unit-commitment programs and wind power forecast systems. This paper evaluates probabilistic models—specifically the probability density functions—of aggregate wind plant power output and conditional and unconditional variations of aggregate wind plant power output. The parameters of the models are fit to historical aggregate wind plant power data from three large North American systems. Parametric and non-parametric evaluations of the suitability of the models are performed in the form of χ^2 goodness-of-fit tests and through the inspection of probability plots and histograms. It is shown that Beta distributions are appropriate models for the aggregate power output and Laplace distributions are appropriate models for wind power variability. Conditional wind power variation follows a generalized extreme value distribution.**

I. Introduction

In 2009 the U.S. became the world leader in total wind plant capacity with over 25 000 MW installed, 8 500 MW of which were installed in 2008 alone [1]. Though the energy converted by these wind plants is approximately one percent of the total electrical energy demand in the U.S., operational concerns are growing as system operators must integrate the power from these inherently uncertain and variable resources [2], [3]. For example, in February of 2008 a weather system swept through Texas causing a significant down ramp of wind power which, along with other non-wind related factors, led to the use of an emergency electric curtailment plan [4]. In July of 2008, the Bonneville Power Administration balancing authority in the Pacific Northwest of the U.S. had an unmanageable surfeit of wind power that caused curtailment of the wind plants [5]. Similar challenges have arisen in other parts of world, for example the 2001 blackout in Crete [6].

To manage wind power variability and uncertainty, new probabilistic tools are being sought, developed and utilized by various power system stakeholders. Examples of such tools are stochastic unit-commitment programs, wind power forecast systems and monte carlo-based simulations for resource planning [7]–[9]. Underlying each tool is a probabilistic model or models, whose accuracy is paramount to the performance of the tool.

This paper proposes and evaluates the performance of several probabilistic models of aggregate wind plant power output. The focus is on large systems with total wind plant capacity levels in the range of 1 000 MW to 5 000 MW. Specifically examined is the appropriateness of utilizing common probability distributions such as the Beta, Gamma, Generalized Extreme Value (GEV), Laplace, Normal and Weibull distributions to model aggregate wind power output and one-hour wind power variations. The distributions are fit to model wind power data from three North American balancing authorities: Bonneville Power Administration (BPA), Electric Reliability Council of Texas (ERCOT) and the Midwest Independent System Operator (Midwest ISO). The models are evaluated using parametric and non-parametric techniques and their performance is examined.

This paper is arranged as follows. In Section II the probabilistic model evaluation methodology is given. The evaluations are preformed in Section III. Conclusions and future outlook are discussed in Section IV.

II. Methodology

The scope of this paper is to evaluate the skill—by parametric evaluation and non-parametric inspection—of various probability distributions to model three characteristics of system-wide normalized aggregate wind plant power output. The characteristics considered are the distributions of: wind power; hourly wind power variation; and conditional hourly wind power variation. The models are compared to data sets from BPA, ERCOT and Midwest ISO. The data are normalized to account for the increases of capacity that occurred within each data set. The parameters of the distributions are allowed to vary for each data set to evaluate the distribution's robustness.

A. Data Considerations

Each data set contains the aggregate hourly-averaged wind power in each system as well as the reported total wind plant capacity for each hour. The BPA and ERCOT data set covers January 1, 2007 through December 31, 2007 [10], [11]. The Midwest ISO began reporting aggregate wind power data in 2009 and so the data set is limited to January 1, 2009 through June 30, 2009 [12]. The characterizations of the wind power in these systems are detailed in [13], [14] and are not discussed in length in this paper.

978–1–4244–5721–2/10/$26.00 © 2010 IEEE

B. Problem Formulation

The general problem formulation for the evaluation is as follows. Let J be the total number of individual data sets that are included in the analysis. Let $\boldsymbol{x}_j$ denote a vector of real numbers whose elements' values are derived from the jth data set. For example, the elements in $\boldsymbol{x}_j$ could be the normalized wind power or the difference in normalized wind power from one hour to the next as computed from the jth data set. If there are G_j elements in $\boldsymbol{x}_j$, then $\boldsymbol{x}_j = [x_{j,1}, \dots, x_{j,g}, \dots, x_{j,G_j}]$ where g refers to the gth element of $\boldsymbol{x}_j$.

Let there be I unique probability distribution functions that are to be fit to the distribution of the values in $\boldsymbol{x}_j$ where i refers to the ith distribution. Let $\boldsymbol{\theta}_i$ be a vector whose elements are the parameters of distribution i. If there are L_i parameters for distribution i, then $\boldsymbol{\theta}_i = [\theta_{i,1}, \dots, \theta_{i,L_i}]$. For example, if the ith distribution function is the GEV distribution, then $\boldsymbol{\theta}_i$ has three elements corresponding to the shape, scale and location parameters.

The parameters for each distribution are estimated for each $\boldsymbol{x}_j$ using the maximum likelihood estimation method [15] where $\hat{\boldsymbol{\theta}}_{j,i}$ is a vector of the estimated parameters for distribution i to $\boldsymbol{x}_j$. An L_i-dimensional space bounded by the 95 percent confidence intervals of the elements of $\hat{\boldsymbol{\theta}}_{j,i}$ is then created. Within this space the set of parameters that minimizes the χ^2 goodness-of-fit test statistic [15] is sought. This is done by descretizing the search space into 10 values for each dimension and exhaustively searching the resulting 10^{L_i} combination of parameters. The set of parameters $\hat{\boldsymbol{\theta}}^*_{j,i}$ that result in the lowest test statistic is selected as the best fit. For example, for the GEV distribution, test statistics are computed for 10^3 combinations of parameters. The χ^2 statistic that corresponds to $\hat{\boldsymbol{\theta}}^*_{j,i}$ is denoted $q_{j,i}$.

The parametric evaluation proceeds by comparing $q_{j,i}$ to the critical value $\chi^2_\gamma(d) = p_{j,i}$, which is calculated as per the χ^2 test based on the significance level α and the number of degrees of freedom d. In this paper, a significance level of $\gamma = 0.05$ is used. The number of degrees of freedom is computed as $d(j,i) = n_j - L_i - 1$, where n_j is the number of bins used in the χ^2 test. The number of bins used is determined by: $n_j = \text{round}(\log_2(G_j))$ [16].

The χ^2 statistic is in itself useful in quantifying the goodness of the fit of a distribution. The χ^2 test is conducted according to: if $p_{j,i} \geq q_{j,i}$, then the fit passes the test and the hypothesis that the ith distribution characterizes the distribution of the values in $\boldsymbol{x}_j$ is accepted. However it must be acknowledged that the test is sensitive to the sample size and number and locations of the bins used. The large sample sizes used in this paper—several thousand hourly data points for each data set—and the fact that several exogenous factors external to the wind energy conversion process such as losses, outages, curtailment, seasonal weather patterns and capacity increases throughout the data sets tempers the expectation that a distribution will be found to fit the data with high accuracy. For these reasons non-parametric evaluations are also performed in the form of the inspection of probability plots or histograms [17], [18].

The application of the discussed general problem formulation to the wind power characteristics investigated in this paper is done by defining how $\boldsymbol{x}_j$ is derived from the data sets, which is discussed in the following subsections.

1) Application to Wind Power: With respect to the preceding general problem formulation, the derived values used to characterize wind power are the normalized wind power $P_j[h]$ values for each hour h, that is:

$$x_{j,h} = P_j[h] = \frac{\bar{P}_j[h]}{\bar{C}_j[h]} \quad (1)$$

where $\bar{P}_j[h]$ and $\bar{C}_j[h]$ are the reported aggregate power output and total capacity of the wind plants in data set j at hour h.

2) Application to Wind Power Variation: With respect to the general problem formulation, the derived values used to characterize hourly wind power variation are computed as:

$$x_{j,h} = \Delta P_j[h] = P_j[h+1] - P_j[h] \quad (2)$$

where $\Delta P_j[h]$ is the variation of the normalized wind power for hour h for data set j. Note that $\Delta P_j[h]$ does not exist for the last hour in the data set.

3) Application to Conditional Wind Power Variation: The hourly variation is further examined using conditional distributions. Let $\mathcal{W} = [\mathcal{W}_1, \dots, \mathcal{W}_m, \dots, \mathcal{W}_M]$ be a partition of the space of normalized wind power values over [0, 1]. The partition is done so that the subsets $\mathcal{W}_m = \left\{[\frac{m-1}{M}, \frac{m}{M})\right\} \forall m \neq M$ and $\mathcal{W}_M = \left\{[\frac{M-1}{M}, 1]\right\}$.

With respect to general problem formulation, let $\boldsymbol{x}_{j,m}$ be a vector of wind power variations derived from the jth data set such that the normalized power at hour h belongs to the mth subset:

$$\boldsymbol{x}_{j,m} = \{\Delta P_j[h] : P_j[h] \in \mathcal{W}_m\}. \quad (3)$$

In this analysis, there are 15 subsets used; if any subset contains less than 20 elements, then there is not sufficient data to perform the χ^2 test and it is omitted from the analysis.

C. Distributions Considered

The normalization process induces bounds on the distributions: [0, 1] for the wind power and [-1, 1] for the wind power variation. These limits are not enforced on the distributions considered under the assumption that the user of the model would be easily capable of discarding values outside the appropriate range.

1) Wind Power Distributions: In selecting the probability distributions to consider modeling the distribution of wind power, it is beneficial to examine the range of shapes that could be expected. In Fig. 1 a hypothetical histogram of the normalized wind power of a single wind plant is shown. Without detracting from the analysis, this example ignores losses, outages and curtailment. As more wind plants are added to the system, assuming independence of the wind speeds at each wind plant, the distribution of aggregate wind plant power output will approach a Normal distribution centered at the capacity factor of the wind plants in the system according to the Central Limit Theorem [15].

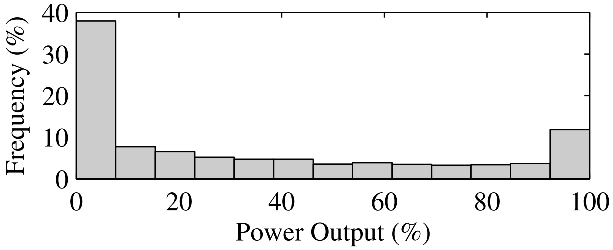

Fig. 1. Hypothetical distribution of normalized wind power from one wind plant.

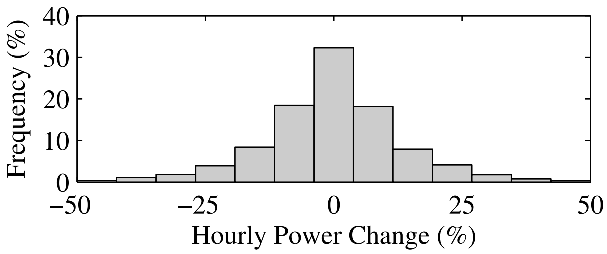

Fig. 2. Normalized histogram of the hour-to-hour variation in wind power.

A robust model will be capable of suitably approximating the aggregate wind power through the continuum of distributions—from that represented in Fig. 1 to a Normal distribution. Therefore, the probability distributions to be considered are first vetted by their ability to model a range of distribution shapes. The distributions considered to model wind power output are: Beta, Gamma, GEV, Normal and Weibull distribution. The fitting of Beta distributions in particular is investigated as it has been proposed—but not rigorously evaluated—in [13] to model wind power.

2) Wind Power Variation Distributions: The shape of the hourly wind power variation distribution has been empirically shown to be similar to that in Fig. 2. As wind plants with independent wind speeds are added to the system, the distribution will approach a Normal distribution per the Central Limit Theorem. The general symmetry of the distribution around zero power variation that is expected warrants that the GEV, Laplace and Normal distributions be considered.

3) Conditional Wind Power Variation Distributions: The distributions considered for each of the 15 subsets $\mathcal{W}_m$ are the same as the unconditional variation model: GEV, Laplace and Normal.

III. Model Evaluations

A. Wind Power

1) Parametric Evaluation: The computed χ^2 statistic q and critical value p for each distribution are provided in Table I. Recall that the value of χ^2 statistic is dependent on the number of bins and sample size considered so that vis-à-vis comparison of χ^2 statistic values of the same distribution among the different data sets is not meaningful; rather, the values should be compared within the same data set. From Table I, the use of the Beta distribution resulted in the lowest χ^2 statistic for each data set, out-performing the other distributions by an

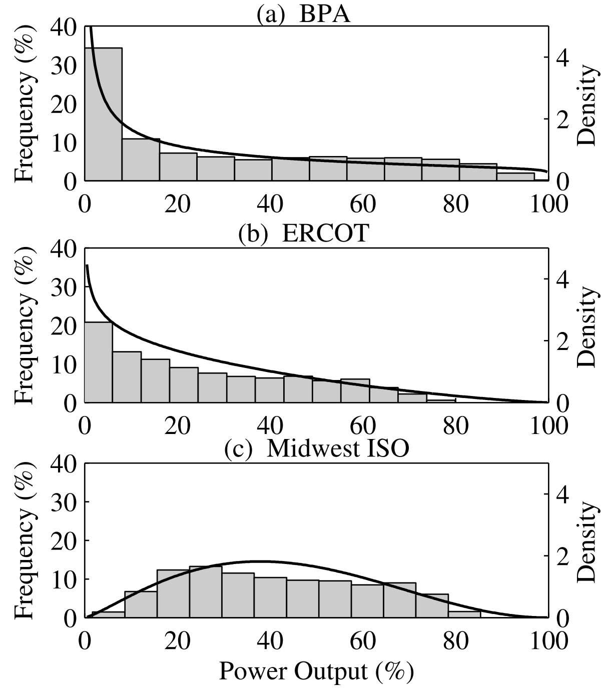

Fig. 3. Normalized histograms of the wind power in each data set with corresponding fit Beta distributions.

TABLE I
Wind Power χ^2 Test Results

	BPA		ERCOT		Midwest ISO	
Distribution	q	p	q	p	q	p
Beta	292	18.3	257	18.3	85	15.5
Gamma	2478	18.3	1298	18.3	180	15.5
GEV	3486	16.9	1043	16.9	163	14.1
Normal	5598	18.3	4113	18.3	1316	15.5
Weibull	2930	18.3	1117	18.3	115	15.5

order of magnitude. Since the shape of the distributions vary considerably, as shown in Fig. 3, the Beta distribution exhibits considerable robustness in its ability to model the data sets. In Fig. 3, for illustrative purposes, the fit Beta distributions are shown. The best fit α and β parameters for the Beta distribution are provided in the second and third columns of Table II.

The Normal distribution had the highest χ^2 statistic for each data set, indicating that the wind plants do not likely have independent wind speeds and therefore the influence of the Central Limit Theorem is limited. Each distribution failed the χ^2 test indicating that none of the distributions considered matches the data within the level of significance, for reasons discussed in Section II.

TABLE II
Best Fit Parameters

	Wind Power		Wind Power Variation	
Data Set	α	β	a	b
BPA	0.46	1.09	-0.00025	0.033
ERCOT	0.81	2.29	-0.00030	0.033
Midwest ISO	2.29	3.12	-0.00065	0.024

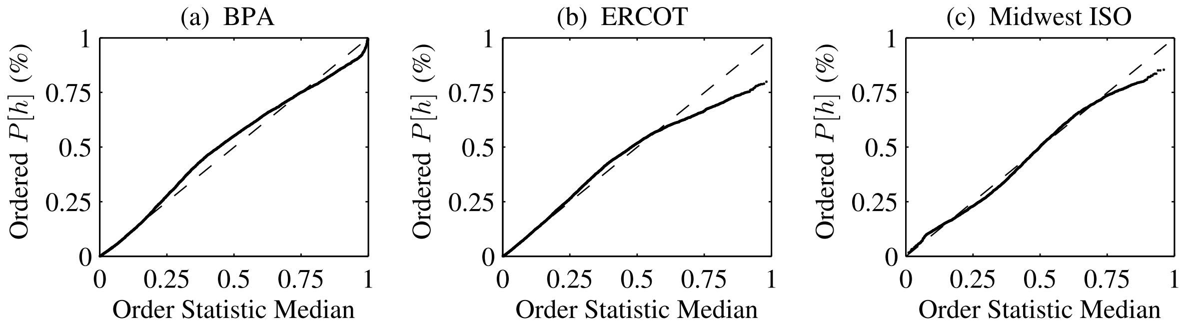

Fig. 4. Probability plots of the BPA, ERCOT and Midwest ISO wind power with respect to the Beta distribution.

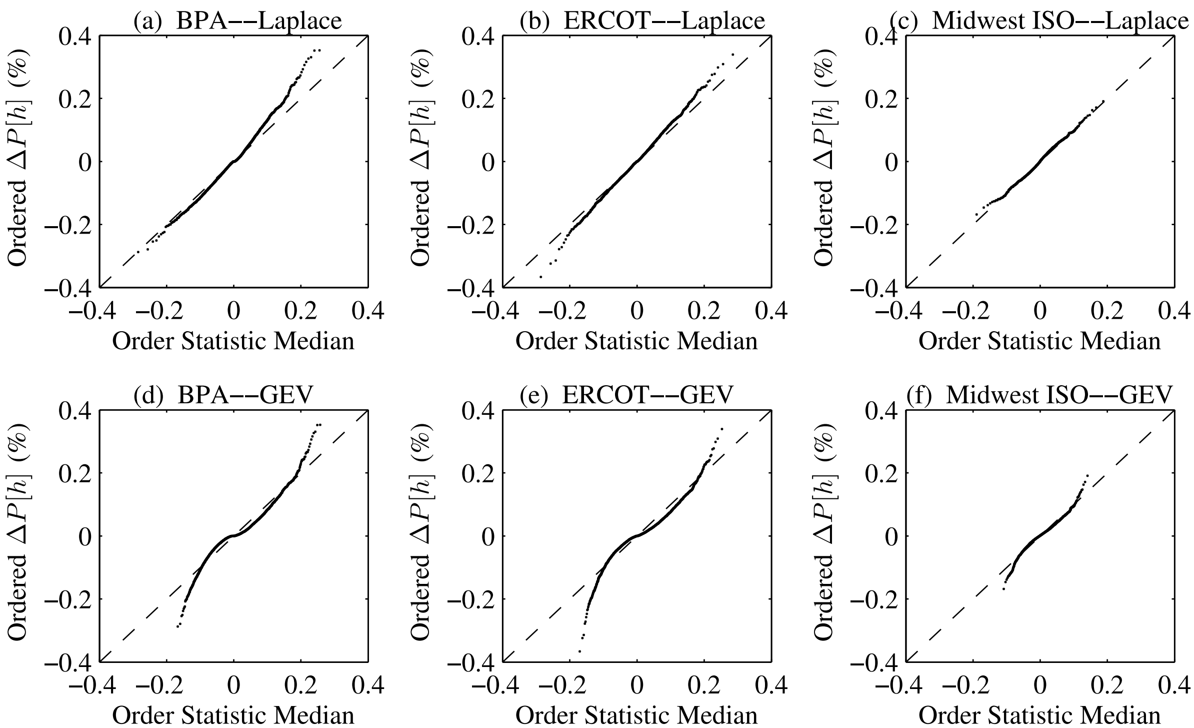

Fig. 5. Probability plots of the BPA, ERCOT and Midwest ISO data sets with respect to the Laplace and GEV distributions.

2) Non-Parametric Evaluation: The probability plots for the Beta distribution for each data set are provided in Fig. 4 for a non-parametric inspection. The abscissa of these plots is the order statistic median [18], [19]. The ordinate of the plots corresponds to the normalized wind power values arranged in ascending order. In the plots in Fig. 4 the wind power values are represented by the black dots which appear almost as a continuous line. The dashed diagonal line is provided for reference.

In inspecting probability plots the size of the deviation of the wind power values from the diagonal line represent the error of the fit distribution. In areas in which the line formed by the data points are steeper than the diagonal reference line, the observed distribution of the data are more dispersed than the Beta distribution; whereas if the line is shallower than the reference line, the Beta distribution is more dispersed than the observed distribution over that range.

In inspecting the plots, it is seen that in general the Beta distribution approximates the distribution of data. The deviations for ERCOT are pronounced for the higher order statistic medians because the ERCOT data set has a maximum normalized value of 0.80, whereas the Beta distribution is defined to 1. A similar feature is exhibited in plot (c) as the maximum normalized value of the Midwest ISO data is 0.85.

B. Unconditional Wind Power Variability

1) Parametric Evaluation: The computed χ^2 statistic and critical values for the wind power variations are provided in Table III. The Laplace distribution most closely approximates the distribution of wind power variation as evaluated by the χ^2 statistic. The GEV distribution had the next lowest average test statistic, followed by the Normal distribution. Overall, the test statistics are higher than those computed for the wind

TABLE III
HOURLY POWER VARIATION χ^2 TEST RESULTS

	BPA		ERCOT		Midwest ISO	
Distribution	q	p	q	p	q	p
GEV	3234	16.9	2083	16.9	159	14.1
Laplace	832	18.3	137	18.3	39	15.5
Normal	6387	18.3	3688	18.3	657	15.5

TABLE IV
NUMBER OF SUBSET χ^2 TEST PASSES

Distribution	BPA	ERCOT	Midwest ISO
GEV	12	3	11
Laplace	3	5	9
Normal	1	1	0

power in the previous subsection, indicating a less accurate approximation. Even the Laplace distribution did not pass the χ^2 test, for reasons discussed in Section II. The best fit location a and scale b parameters for the Laplace distribution for each data set are provided in the final two columns of Table II.

2) Non-Parametric Evaluation: The non-parametric evaluation of the fit of the distributions is performed by inspecting the probability plots in Fig. 5 for the Laplace and the GEV distributions. The Laplace distribution is in line with the diagonal reference line for most of the ordered data points between ± 0.2, indicating skill at modeling normalized power deviations between ± 20 percent of rated power. The skill at larger deviations is diminished. The GEV distribution is nearly symmetric about the origin in manner indicating that the distribution does not capture the peak exhibited of the data around the zero variation and has fatter tails. This non-parametric evaluation reinforces that the Laplace distribution is the most skilled at modeling one hour wind power variation.

C. Conditional Wind Power Variability

1) Parametric Evaluation: The number of subsets $\mathcal{W}_m$ in which the GEV, Laplace and Normal distribution passed the χ^2 test are summarized in Table IV. Note that for ERCOT and Midwest ISO the number of subsets with a sufficient number of elements to perform the χ^2 test were 13 and 14, respectively. Also note that due to the nature of the χ^2 test, it is possible for two different distributions to pass the test for the same set of data. The GEV exhibited the greatest skill in the BPA and Midwest ISO data sets, whereas the Laplace distribution was superior at modeling the ERCOT data. It is interesting that the Laplace distribution outperformed the GEV distribution when modeling the unconditional wind power variation, but the GEV recorded more χ^2 test subset passes. The reason for this is the robustness in the shapes that can be represented by the GEV distribution.

Examination of the GEV shape k, scale σ and location μ, parameters show how the best fit distribution change with subset number, as shown in Fig. 6. The trend of the shape parameter k is to decrease, indicating asymmetry toward negative power variations as $P[h]$ increases. Since k remains negative, the shape of the GEV is indicative of a reversed

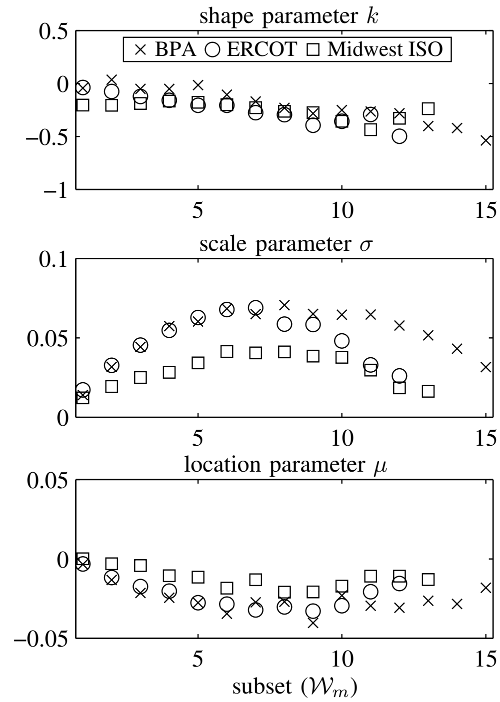

Fig. 6. Best fit parameters for the GEV distribution for each data set.

Weibull distribution. The parameter σ is concave, indicating a larger spread in the data as $P[h]$ approaches 50 percent of the normalized capacity. The convex shape of μ means the average value tends to become negative as $P[h]$ increases.

2) Non-Parametric Evaluation: Due to page limitations, probability plots of the conditional data are not examined. Rather, linearized histograms of the variation data are inspected. The plots for BPA, ERCOT and Midwest ISO data sets are displayed in Fig. 7 through Fig. 9, respectively. In these plots, the x-axis is the range of normalized power within each subset, the y-axis is the computed normalized variation of power and the z-axis is the normalized frequency of occurrences, normalized to the number of elements within each subset $\mathcal{W}_m$. Note that in the subsets $\mathcal{W}_{13}$ through $\mathcal{W}_{15}$ are empty for the ERCOT data set, and $\mathcal{W}_{14}$ through $\mathcal{W}_{15}$ are empty for the Midwest ISO data set.

In comparing the figures, it is apparent that they each follow a similar saddle-like shape, with high probability of zero variation within each subset, which is most pronounced in the first and last subset. In each subset, the frequency of deviation decayed with variation magnitude as variations greater than ± 20 percent rarely occurring. The distributions are asymmetrical within each subset: increases in power are more likely to occur if the power output is low and decreases in power are more likely to occur if the power output is high. These observations are in accordance with the parameters of the GEV distribution and so reinforces the conclusion of its ability to model this characteristic.

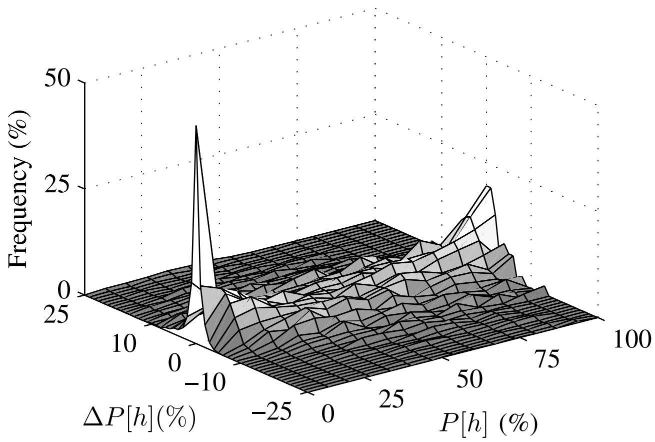

Fig. 7. Linearized histogram of the wind power variation in BPA.

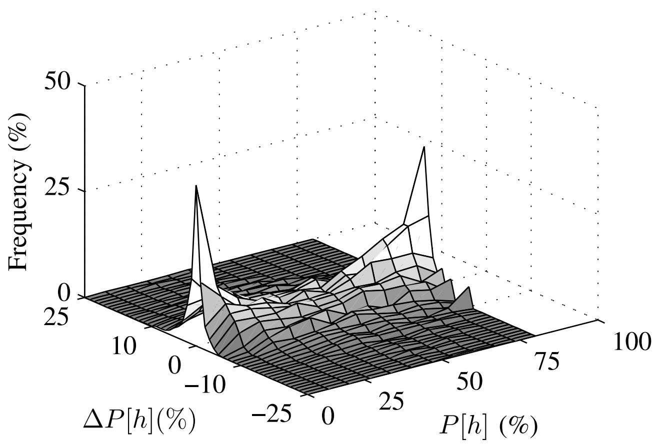

Fig. 8. Linearized histogram of the wind power variation in ERCOT.

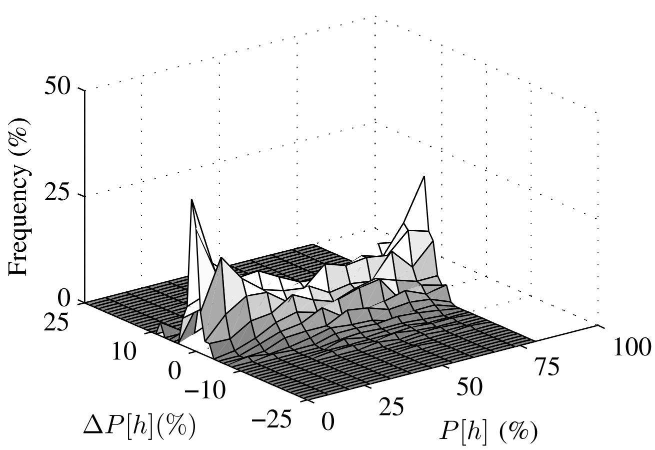

Fig. 9. Linearized histogram of the wind power variation in Midwest ISO.

IV. Conclusions

Probabilistic models of the system-wide wind power and wind power variations are becoming increasingly important to power system stakeholders in order to manage the uncertainty and variability imposed by weather-driven renewable resources. This paper employed parametric and non-parametric evaluation techniques on several probability distributions to determine their suitability as models. The analysis indicated that Beta distribution had the lowest χ^2 statistic when modeling aggregate wind power and the Laplace distribution was superior by that measure when modeling hour-to-hour variations. When conditionally modeling the variations, the robustness of the generalized extreme value distribution was superior. It is clear from these results, however, that since few of the distributions passed the goodness-of-fit test that diligence and prudence must be used in applying these models.

References

[1] American Wind Energy Association, "Annual wind industry report," Tech. Report, Apr. 2009.

[2] J. C. Smith, R. Thresher, R. Zavadil, E. A. DeMeo, R. Piwko, B. Ernst, and T. Ackerman, "A mighty wind," *IEEE Power Energy Mag.*, vol. 7, no. 2, pp. 41–51, Mar./Apr. 2009.

[3] J. C. Smith, M. Milligan, E. A. DeMeo, and B. Parsons, "Utility wind integration and operating impact state of the art," *IEEE Trans. Power Syst.*, vol. 22, no. 3, pp. 900–908, Aug. 2007.

[4] Electric Reliability Council of Texas, "ERCOT's operations report on EECP event," Project No. 27706, Mar. 2008. [Online]. Available: http://interchange.puc.state.tx.us/WebApp/Interchange/Documents/27706_114_577769.PDF

[5] Associated Press, "NW wind power was almost too much of a good thing," KTRV news report, 2008.

[6] T. Papazoglou and A. Gigandidou, "Impact and benefits of distributed wind generation on quality and security in the case of the Cretan EPS," in *CIGRE/IEEE Int'l Symposium on Quality and Security of Electric Power Delivery Systems*, Montréal, QC, Canada, Oct. 2003, pp. 193–197.

[7] A. Tuohy, P. Meibom, E. Denny, and M. O'Malley, "Unit commitment for systems with significant wind penetration," *IEEE Trans. Power Syst.*, vol. 24, no. 2, pp. 592–601, May 2009.

[8] P. A. Ruiz, C. R. Philbrick, and P. W. Sauer, "Wind power day-ahead uncertainty management through stochastic unit commitment policies," in *Proc. Power Systems Conf. and Exposition*, Seattle, WA, Mar. 2009.

[9] Seattle City Light, "Integrated resource plan 2008," 2008.

[10] Bonneville Power Administration. (2009, Apr.) BPA balancing authority load and total wind generation: 2007. [Online]. Available: http://www.transmission.bpa.gov/Business/Operations/Wind/default.aspx

[11] Electric Reliability Council of Texas. (2008, Oct.) ERCOT hourly wind output-2007. [Online]. Available: http://planning.ercot.com/data/hourly-windoutput/2007/

[12] Midwest ISO. (2009, Aug.) Historical hourly wind data reports. [Online]. Available: http://www.midwestiso.org/publish/Document/

[13] H. Louie, "Statistical analysis of the effects of geographic diversity on aggregate wind power generation," *IEEE Trans. Power Syst.*, submitted for publication.

[14] ——, "Characterizing and modeling aggregate wind plant power output in large systems," in *Proc. PES General Meeting*, Minneapolis, MN, Jul. 2010, to be published.

[15] A. Papoulis and S. U. Pillai, *Probability, Random Variables and Stochastic Processes*, 4th ed. New York, NY: McGraw-Hill, 2002.

[16] T. P. Ryan, *Modern Engineering Statistic*. Hoboken, NJ: John Wiley & Sons, Inc., 2007.

[17] J. Chambers, W. Cleveland, B. Kleiner, and P. Tukey, *Graphical Methods for Data Analysis*. Stamford, CT: Duxbury Press, 1983.

[18] NIST and SEMATECH. (2009, Jul.) NIST/SEMATECH e-Handbook of Statistical Methods. [Online]. Available: http://www.itl.nist.gov/div898/handbook/

[19] J. J. Filliben, "The probability plot correlation coefficient test for normality," *Technometrics*, vol. 17, no. 1, pp. 111–117, Feb. 1975.