

Measuring Error in Wind Power Forecasting Using a New Forecasting System

W. David Lubitz* and Bruce R. White†

*Department of Mechanical & Aeronautical Engineering, University of California, Davis
One Shields Ave., Davis, California, USA*

The sources of error in wind power forecasting were investigated by constructing an experimental near-real time wind power forecasting system that operates on a desktop PC and forecasts 12 to 48 hours in advance. The system uses model output of the Eta or COAMPS regional scale forecasts (RSFs), and is currently configured to forecast the power production of a case study wind farm in the Altamont Pass, California, USA from 12 to 48 hours in advance. Archived RSF data can be used, allowing comparison of different methods of forecasting using the same input data with known actual power production. Several different forecasting approaches, including deriving predictive equations using multiple linear regression, wind-tunnel studies, and a "forecast matching" scheme were investigated using the system. It was found it is difficult to predict near surface conditions from forecasted upper level weather conditions. If the wind speed at the wind farm can be accurately forecasted, several methods can be used to predict power production with a minimal loss of accuracy. The "forecast matching" method was found to be relatively accurate at wind power forecasting, especially given its simplicity.

Nomenclature

A, B	= constants
C	= binary variable ($C = 1$ if wind speed is greater than wind turbine cut-in speed; otherwise $C = 0$)
f	= Coriolis parameter
F_i	= value of variable i in current forecast
H_i	= value of variable i in archived forecast
i	= variable number
n	= number of variables
P	= predicted power production of wind farm in units of percentage of wind farm observed capacity
R_i	= normalizing factor for i^{th} variable
S	= matching score between two forecasts
U	= wind speed
U_G	= wind speed at geostrophic height
U^*	= friction velocity
Y	= decimal time of year (At 0:00 Jan. 1, $Y = 0$. At 23:59 Dec. 31, $Y = 1$)
z	= height above ground
z_o	= surface roughness
α	= angle of wind direction difference between geostrophic height and near-surface height
k	= von Kármán's constant (≈ 0.4)

I. Introduction

Generally, the results and the accuracy of most wind power forecasting are not made publicly available since most forecasting is done by companies contracted by clients, and both the forecasting methods and the

* Doctoral Candidate, Department of Mechanical & Aeronautical Engineering, University of California, Davis, CA, USA, 95616, AIAA Member.

† Professor, Department of Mechanical & Aeronautical Engineering, University of California, Davis, CA, USA, 95616, AIAA Member.

Copyright © 2005 by W. D. Lubitz and B. R. White. Published by the American Institute of Aeronautics, Inc. or the American Society of Mechanical Engineers with permission.

forecasts themselves are sensitive commercial information. This makes it difficult to determine where to focus research work on wind power forecasting, and is the main reason an in-house forecasting system has been developed at UC Davis.

One case where two forecasting systems were directly compared was the Electric Power Research Institute's (EPRI) California Wind Energy Forecasting Study, a multi-year wind forecasting study¹² that concluded in 2003. Two different forecasting systems, Risoe's *Prediktor* and TrueWind Solution's *eWind* were applied to two wind farms in California, including one in the Altamont Pass, and forecasted power production for one year. Interestingly, the average forecasting error of the two systems was nearly identical, remarkable given that the forecasting systems implemented were completely different, and even used different input RSFs from the United States National Weather Service (NWS) (Risoe used Eta, TrueWind used AVN).

An Altamont Pass wind farm was selected for use as the case study. The study area consists of approximately 100, 100 kW turbines installed along ridgelines with hub heights of about 20 meters, situated in complex terrain (Fig. 1). The Altamont Pass region consists of a series of ridgelines with heights up to several hundred meters. Winds in the Pass are strongly influenced by the local terrain, with a wind constancy at the wind farm site of 0.65. Sea breeze effects significantly affect local flow, and high winds are noted to occur at night during periods of atmospheric stability. There is also significant seasonal variability in the winds. As shown in Figs. 3 and 4, summers are characterized by strong winds from the southwest, while in the winters winds are lighter and of variable direction. All of this makes the Altamont Pass a challenging wind resource to forecast accurately.

II. The Forecasting System

The wind power forecasting system implemented at UC Davis is diagramed in Fig. 2. The system operates on a desktop PC. The system automatically downloads Eta and COAMPS regional scale model output as it becomes available. (Archived model output data also can be used.) Using this downloaded model output, a forecast of the wind speed and direction at the wind farm meteorological tower is produced by one of several methods. The results of this step can be checked for accuracy by comparing to recorded wind speed and direction data from the wind farm meteorological tower provided by the wind farm operator. Next, the wind speed and direction at the tower is used to predict the wind speed and power production of the wind farm. The power production results can be checked for accuracy by comparing the forecasted power to the power production of the wind farm provided by the farm operator. The accuracy of just this step also can be determined by using the recorded wind speed and direction at the meteorological tower as the input, instead of the forecasted wind speed and direction, and comparing the resulting predictions to the recorded power production. Each of the first three steps in Fig. 2 is a source of additional potential error in the final forecast. The level of error for each of these steps will be investigated below using results from the case study wind farm.

III. Uncertainty of the Regional Scale Forecast (RSF)

It is generally believed that the majority of the uncertainty in commercial wind power forecasting results is due to the uncertainty in the output of the RSFs that are used as inputs to the forecasting systems, however the measures




Fig. 1. Map showing Altamont Pass site ("Site") and Oakland radiosonde site ("OAK"). Elevations in meters above sea level. Map is 240 km wide.

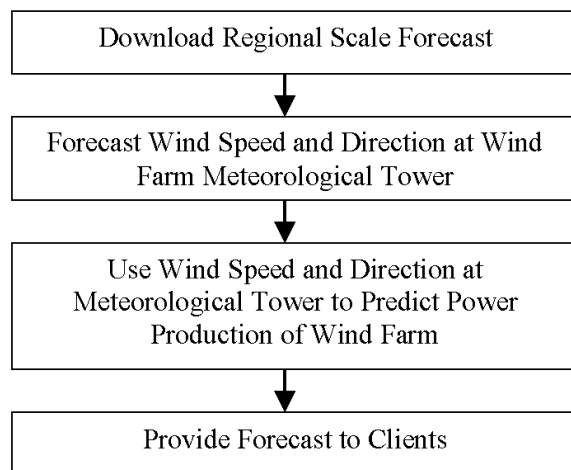


Fig. 2. Schematic diagram of experimental wind power forecasting system developed at UC Davis.

of error made available by RSF operators rarely assess the error in the specific input variables used by wind power forecasters (such as wind speeds at geostrophic levels), or only provide averages of the error levels in a variable over many forecasts, and therefore cannot be compared directly to specific wind power forecasts.

The RSF used in this analysis is the 40 km grid Eta forecast, a regional-scale, non-hydrostatic operational weather forecasting system maintained by the National Center for Environmental Prediction (NCEP), part of the NWS. The model is run twice a day, at 0:00 GMT and 12:00 GMT, and produces forecasts at three-hour intervals from zero to 84 hours in the future. Immediately after a forecast is completed, it is made available on several FTP servers in GRIB format. Forecast output is given on 40 km square grid overlaid across the continental United States and surrounding regions, using a Lambert Conformal projection. The domain is divided vertically into 39 levels,

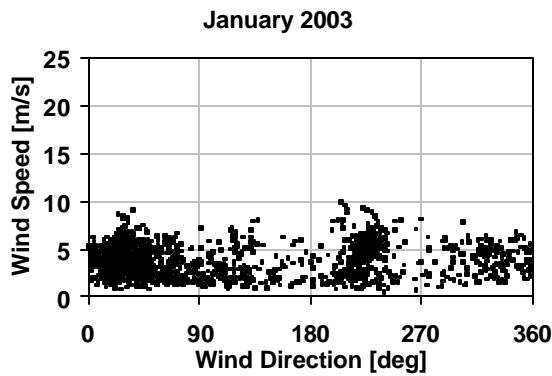


Fig. 3. Wind speed vs. direction at Altamont Pass wind farm during January, 2003, showing typical winter winds.

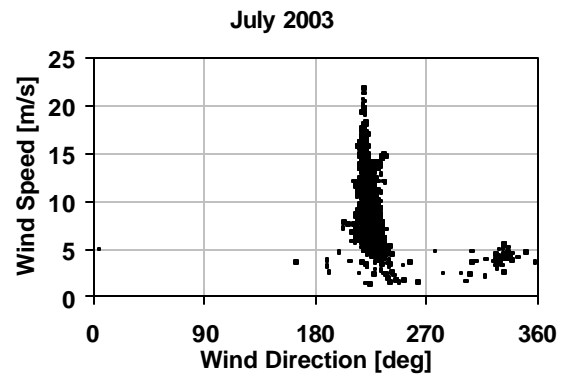


Fig. 4. Wind speed vs. direction at Altamont Pass wind farm during July, 2003, showing typical summer winds.

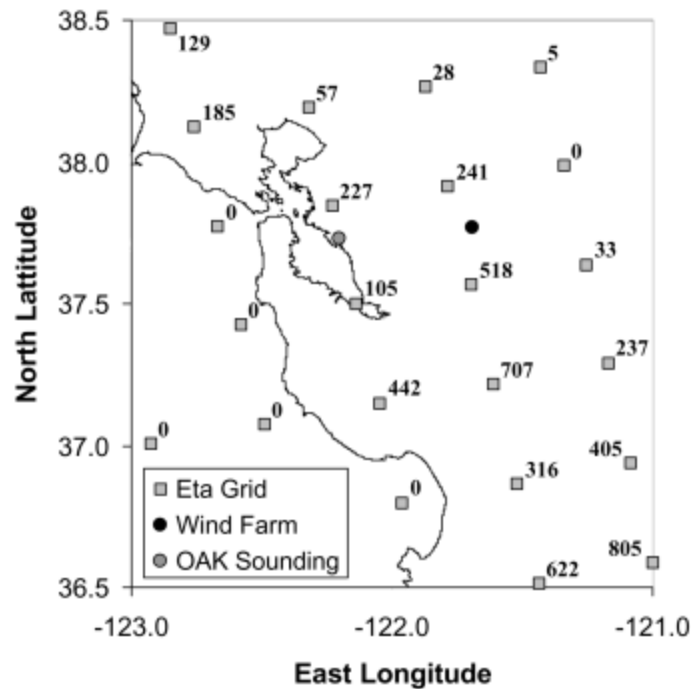


Fig. 5. Location of Eta model grid points relative to the case study wind farm and Oakland radiosonde launch site. Each Eta grid point is labeled with the surface height in meters above sea level at that point in the Eta model.

along surfaces of constant pressure, with the lowest level at 1000 mb and the highest at 50 mb. Since pressure changes with time, the absolute elevation of a level at a specific grid point is variable. (The average height of 1000 mb was found to be 140 m ASL; 900 mb was 1030 m ASL, 800 mb was 2010 m ASL and 700 mb was 3020 m ASL.)

As part of the UC Davis case study, weather conditions at a nearby radiosonde launching site (Oakland, CA, USA, 45 km west of Altamont Pass. WMO ID# 72943) also are forecasted. The radiosonde is launched twice daily at 0:00 GMT and 12:00 GMT, and reports wind speed and direction, temperature and dewpoint at various altitudes from the surface to well into the stratosphere.

Initially, bicubic interpolation was used to extrapolate forecast weather conditions from the RSF grid points to the sounding location. However, it was suspected that the interpolation method itself also introduces error. Therefore, bilinear interpolation and a closest grid point method also were implemented in order to investigate the sensitivity of the error levels to the interpolation method. Any or all of these interpolation methods can be used in the UC Davis forecasting system to extrapolate gridded RSF data to wind farm meteorological tower locations. By interpolating the RSF output to the radiosonde location, many upper level variables forecasted by the RSF can be compared to field measurements on a forecast by forecast basis, and the level of uncertainty in the data from the RSF can be estimated. The three interpolation methods were compared by using archived Eta forecasts for a 21 month period from November 30, 2002 through July 31, 2004. The Eta forecasted wind speed at each of the surrounding grid points was determined from the horizontal components of the wind speed at each grid point. This data was then interpolated to the sounding location using bicubic or bilinear interpolation. The third “interpolation” method involved using the data from the closest grid point to the sounding location. The radiosonde takes frequent measurements as it rises, typically one or more every few hundred meters. Readings were linearly interpolated to the pressure of Eta levels for analysis. No statistical or other corrections were applied.

The overall error level, quantified as a mean absolute error (MAE), the absolute value of the forecasted minus the actual value of a variable, was almost identical for the three interpolation methods implemented. MAE results for wind speed at several different levels and at various forecast intervals, are given in Fig. 3. As expected, error levels increase over longer forecast intervals, and higher levels tend to be forecasted more accurately than lower levels. In absolute terms, the data given in Fig. 6 represents MAE values between 1.5 m/s (for the initialization runs at 700, 800 and 900 mb), to 2.5 to 3 m/s (for the 48 hour forecasts). The relatively high levels of MAE observed in the initialization runs (“forecasting 0 hours ahead”) were somewhat unexpected, since the sounding data is used to initialize the Eta model.

The explanation is that the measured data must be interpolated to the grid points and modified to agree with other available data, thus increasing the discrepancy between actual and interpolated data. There are also significant differences between the terrain at the sounding location as modeled by Eta and the actual terrain. Compare the heights of the grid points near the sounding location that range as high as 227 m, as shown in Fig. 5, to the actual site elevation of 6 m above sea level. While the actual Oakland sounding site is adjacent to the San Francisco Bay on a several kilometer wide coastal plain, the local Eta grid points have averaged this region with nearby ridges; the complex terrain of the San Francisco Bay area can only be modeled as a few hills. Another cause of discrepancies is that each model run starts at minus 12 hours, (to ensure stability and consistency by hour zero), so only earlier sounding data is used, and giving additional opportunity for the initialization run to diverge from the measured data at hour zero.

For the three interpolation methods, wind speed MAEs between 1.5 and 3 m/s were observed. It should be noted that in EPRI², both forecasters forecasted the wind speed at a meteorological tower on a ridge in the Altamont Pass with an anemometer height of 18 m AGL. Wind speed MAEs for the one year period were 2.0 m/s and 2.2 m/s. These results suggest that when constructing a wind power forecast, it is important to start with the most accurate data available from the RSF (typically higher elevation data). It also suggests the importance of the statistical corrections which are typically included in wind power forecasting models, even ones based on physical modeling, in reducing the overall error. Since the most accurately forecasted variables are at levels above the atmospheric boundary layer, methods of predicting surface weather conditions based on those variables in the upper atmosphere were investigated.

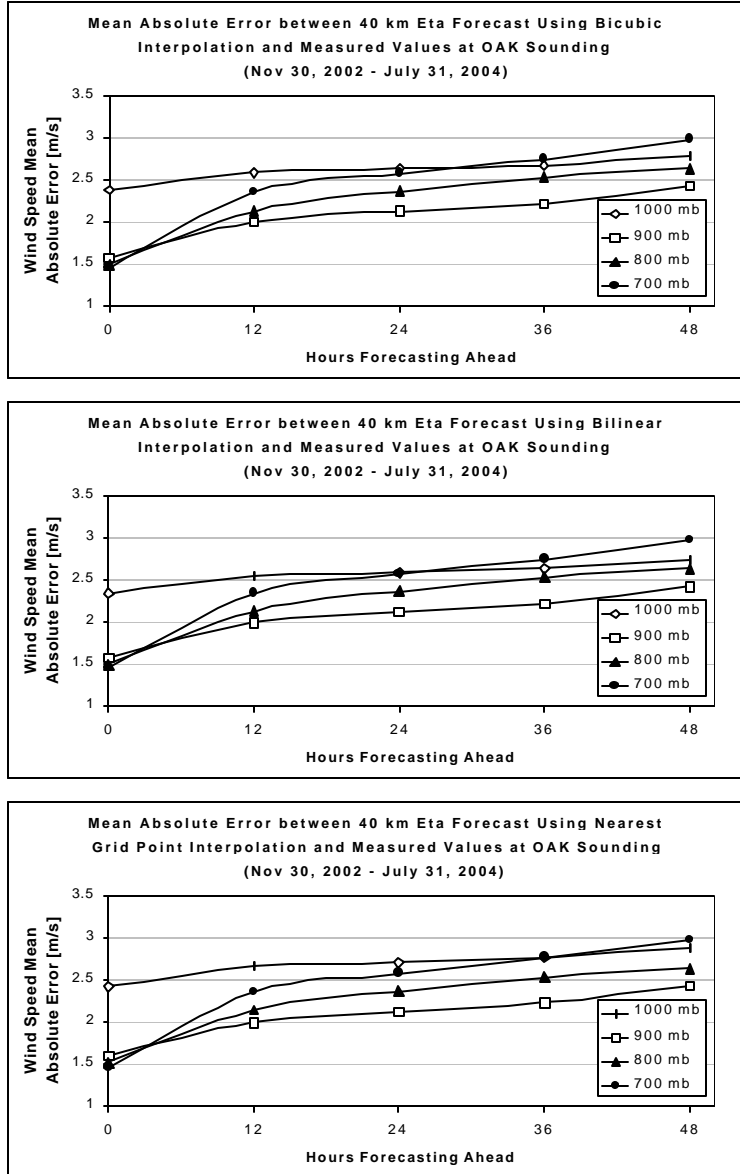


Fig. 6. Mean absolute error in Eta forecasted wind speed at Oakland sounding location for four different levels, as a percentage of the average measured wind speed observed at that level. Forecasted wind at the sounding location was determined from gridded Eta model output using bicubic interpolation (top), bilinear interpolation (middle) and the grid point closest to the sounding location (bottom). Results are given for forecast intervals from 0 hours ahead (initialization run) to 48 hours ahead.

IV. Forecasting a Reference Near-Surface Wind Speed

Once the RSF has been obtained, it is necessary to use the forecast data to forecast a near-surface wind speed, typically the wind speed at the meteorological tower associated with the wind farm. This is necessary because the RSF models have such large grid spacing (Eta's grid spacing is 40 km) that only the largest terrain features, such as large mountain ranges, can be reproduced. As shown in Fig. 5, the terrain model used by Eta does not reproduce a pass in the Altamont area; instead the Altamont Pass and surrounding region is modeled as a slope. Even if model data with a much finer grid resolution is used that can capture the individual ridges and valleys near a wind farm, the region near the surface (where wind turbines are) is the most difficult region to model accurately. The increased turbulence and shear in the surface layer, combined with factors such as the flow around local terrain and heating or cooling of the surface, make modeling of the near-surface region difficult.

A. Geostrophic Drag Law

Landberg and Watson³ used the geostrophic drag law to forecast near surface wind speeds from upper air wind speeds taken from an RSF. This method involves the use of a few equations that can be iteratively solved with little computational effort. The neutral geostrophic drag law is a method of predicting surface shear stress, in the form of the friction velocity U_* , based on the wind speed at geostrophic height U_G :

$$U_G = \frac{U_*}{\mathbf{k}} \sqrt{\left[\ln\left(\frac{U_*}{fz_o}\right) - A \right]^2 + B^2} \quad (1)$$

where f is the Coriolis parameter, z_o is the surface roughness and \mathbf{k} is von Kármán's constant (≈ 0.4). A and B are constants. Zilitinkevich⁴ found that the median values of A and B reported in 16 studies gave $A = 1.7$ and $B = 4.5$, with standard deviations of 0.9 and 1.1 respectively. The angle of the surface shear stress direction from the direction of the geostrophic wind \mathbf{a} is given by

$$\tan \mathbf{a} = \frac{-B}{\ln\left(\frac{U_*}{fz_o}\right) - A} \quad (2)$$

Landberg and Watson³ used Eqns. (1) and (2) in conjunction with the logarithmic law

$$\frac{U}{U_*} = \frac{1}{\mathbf{k}} \ln\left(\frac{z}{z_o}\right) \quad (3)$$

to predict a reference wind speed at turbine hub height from upper atmosphere winds forecasted by HIRLAM.

While Landberg and Watson³ reported success applying this method to sites in Denmark and other parts of northern Europe, the reliability of the geostrophic drag law appears to break down in the more complex terrain of the Altamont Pass, and this approach did not accurately predict the near-surface winds at the case study wind farm.

To determine if this method could be applied to the Altamont Pass region, archived initialization runs (i.e., “zero hour forecasts”) of Eta model output from November 30, 2002 to July 31, 2004 was used to predict the wind speed at the case study wind farm meteorological tower. To use this method, first a vertical level in the Eta output had to be chosen as being at the geostrophic height. Six different levels evenly spaced between 700 and 950 mb were used as geostrophic level on a trial basis. For a specific forecast, the overall wind speed at the geostrophic height was taken from the Eta model output. Eqns. (1), (2) and (3) then were used to predict the wind speed that would be seen at the meteorological tower if it were situated in flat terrain. A speed-up factor of 1.5 was then applied to the forecasted wind speed to account for the ridge-top location of the meteorological tower. The speed-up factor was determined by comparing wind speeds at simulated anemometer height in a wind-tunnel test of the complex terrain around the meteorological tower site, with a test of flat terrain. Details of the wind-tunnel testing can be found in Cheng et al.⁵ and Lubitz and White⁶.

The results of the analysis are given in Table 1. Results are given assuming geostrophic height at six different levels. Input wind speeds taken from geostrophic height in Eta model initialization runs (“0 hour forecasts”) between

	Geostrophic Level					
	700 mb	750 mb	800 mb	850 mb	900 mb	950 mb
ME	-0.86	-1.76	-2.46	-2.88	-3.03	-3.02
MAE	4.57	4.41	4.34	4.39	4.23	3.90
R	-0.04	-0.02	0.00	0.04	0.14	0.34

Table 1. Mean Error (ME) and Mean Absolute Error (MAE), in units of m/s, and dimensionless correlation coefficient (R) of predictions of wind speed at case study wind farm meteorological tower using geostrophic drag law.

November 30, 2002 and July 31, 2004. Average measured wind speed at the meteorological tower was 7.46 m/s. The R values are the correlation coefficients between the Eta model wind speed at each given geostrophic height and the measured wind speed at the meteorological tower.

The average forecasted wind speed at the meteorological tower anemometer was low for all of the trial geostrophic heights, even with the application of the speed-up factor. A closer look at the data revealed a significant lack of correlation between the wind speed at geostrophic height, and the wind speed at the meteorological tower anemometer location, with correlation coefficients ranging from 0.04 when compared to the 700 mb wind, to 0.34 when compared to the 950 mb wind. It is apparent that this lack of correlation represents a potentially very large source of error in a forecast for the case study site.

B. Multiple Linear Regression

Burda et al.⁷ reported success predicting the *average daily* wind speed at an Altamont Pass site with an equation derived using multiple linear regression, using inputs of the zonal component of the 950 mb Oakland sounding wind speed and the maximum daily air temperature in San Francisco. Unfortunately, attempts to use multiple linear regression to derive equations to predict the wind speed at the meteorological tower anemometer *at a specific time* using variables from the Eta model output were unsuccessful. None of the derived equations had an R^2 value greater than 0.5, even equations derived for specialized cases (such as for specific times of day or seasons). This was attributed primarily to the lack of correlation between the meteorological tower wind speed and the upper level wind speeds, temperatures and pressures. For example, Fig. 7 shows the relation between the wind speed at 700 mb in the Eta initialization runs ("zero hour forecasts") and the wind speed as measured at the meteorological tower.




Fig. 7. Comparison of the wind speed as measured at the wind farm meteorological tower versus the wind speed at the 700 mb level of the Eta initialization ("zero hour forecast") runs.

C. Forecast Matching

It was decided to use the archives of Eta forecasts and wind farm information compiled at UC Davis to directly predict the wind farm conditions, due to the lack of success predicting the wind conditions at the wind farm meteorological tower from the Eta output using the above methods.

At UC Davis, the information from each forecast is archived in a separate file by the forecasting system after the forecast is downloaded. For this analysis, the corresponding information from the wind farm (wind speed and direction, and power production) at the time being forecasted was appended to each forecast file, so that each archive forecast file contained the measured wind farm conditions at the time being forecasted, for future error analysis. A "matching forecast" can be performed by downloading a current Eta forecast and then searching the

archive for a "matching" archived forecast; one with conditions closest to those predicted by the current forecast. Since the two "matched" forecasts are for two times with very similar weather, it is anticipated that the wind farm conditions for the current forecast will be close to those associated with the matched archived forecast: the archived wind speed and direction (and the power production) are used as a new forecast.

A means of quantifying the apparent similarity of key conditions between two Eta forecasts was necessary to find a match. A scoring system was devised to quantify the degree of difference between a set of corresponding variables in the current forecast and an archived forecast. A set of 18 variables were chosen ("U" and "V" components of wind speed, temperature and level height at 1000 mb, 900 mb, 800 mb and 700 mb, and the horizontal components of the pressure gradient, determined from bicubic interpolation) for comparison between the current forecast and an archived one.

The "matching score" S of the two forecasts was then calculated as

$$S = \sum_{i=1}^n \left(\frac{H_i - F_i}{R_i} \right)^2 \quad (4)$$

where n is the number of variables (18 in this case), H_i and F_i are the archived and current forecast values of variable i (where $1 \leq i \leq n$), and R_i is a normalizing factor used to weight the importance of the variable. (Initially for each variable, R was arbitrarily set equal to the standard deviation of the variable observed in the dataset, divided by four.) A score is calculated for each archived forecast, with a lower score indicating greater similarity between an archived forecast and the current forecast. For a given current forecast, the wind farm conditions are taken from the file of the lowest scoring archived forecast.

The accuracy of the matching method was determined by forecasting the wind farm conditions corresponding to each of the archived forecasts. When a specific forecast was being determined, archive forecasts within four days of the one being used were excluded from the matching analysis. This ensures that the forecast is not matched to itself, or to a forecast very close in time, representing the identical weather conditions in the forecast. Both wind speed at the meteorological tower anemometer and wind farm power production were forecast. ME and MAE results, summarized in Table 2, were encouraging.

	Wind Speed [m/s]		Power Production [% Wind Farm Cap.]		Obs.
	ME	MAE	ME	MAE	
All Data	-0.05	2.96	-0.46	18.0	12813
Power < 20%	0.84	2.53	8.33	10.6	7632
20% < Power < 80%	-0.78	3.26	-6.56	29.5	3293
Power > 80%	-2.38	4.22	-25.33	28.1	1888

Table 2. Observed accuracy of the matching method at forecasting wind speed at the case study wind farm meteorological tower anemometer and the wind farm power production. Forecasts were prepared for each of the Eta forecasts in the UC Davis archive (consisting of forecasts from 0 to 48 hours in the future, initialized at 0:00 and 12:00 GMT, between Nov. 30, 2002 and July 31, 2004)

V. Predicting Power Production from the Reference Wind Speed

A. Predicting Power Using A Statistically Derived Equation

Multiple linear regression (MLR) was applied to derive an equation to predict the power production of the case study wind farm using the wind speed at the wind farm meteorological tower anemometer and the date and time as input data. (For use in a forecasting system, the wind speed at the meteorological tower anemometer first would have to be forecast.) A common statistical approach to predicting meteorological conditions at a specific location based on other measured data is to generate a predictive equation for each predictand, or measurement to be predicted, using multiple linear regression, iteratively choosing from the potential predictors to minimize the number of predictors while maximizing accuracy. Wilks⁸ gives a good description of the statistical processes that were used. MLR was used to generate an equation to predict the power production of the case study wind farm. Time data and the wind speed at the meteorological tower from two years of historical data (July 1, 2001 to June 30, 2003) were

used to generate a dataset with 13 potential predictor variables, including the logarithm and square root of the wind speed, wind speed to the first, second and third power, and sine and cosine functions with periods of half day, one day and one year. Wind direction was not included in this equation (or used in the following “median-based” power curve) due to the wind farm anemometer apparently malfunctioning and sticking in a single position for periods ranging from hours to months, and within the two year dataset, accordingly some times of the year contained wind direction data considered unreliable.

Screening was applied to remove the least important potential predictors, and the following equation was derived to predict power P , expressed here as a percentage of the observed wind farm capacity

$$P = -11.8278U + 2.281434U^2 - 0.07626U^3 - 7.74768C - 4.59537 \cos(2\pi Y) + 17.50517 \quad (5)$$

where U is the wind speed at the meteorological tower in m/sec, $C = 1$ if U is greater than the cut-in wind speed of the turbine or $C = 0$ otherwise, and Y is the time of the year in decimal form (i.e., at 0:00 on Jan. 1, $Y = 0$, and at 23:59 on Dec. 31, $Y = 1$).

B. Predicting Power Using A Power Curve Derived by Curve Fitting Historical Data

If historical power production information is available for a wind farm, it is also possible to generate a wind farm power curve by performing a best fit to the power production as a function of wind speed.

For the case study wind farm, a power curve was generated using two years (July 2001 to June 2003) of power production data. Each observation was binned by wind speed, and then for each wind speed bin, the median value of power production was taken as the amount of power that would be expected to be produced at the bin wind speed. Median values were used instead of mean values in order to reduce the effects of outlying data points. The resulting power curve is presented in Fig. 8.

C. Predicting Power Using Wind-Tunnel Modeling

A method was derived using measurements taken in an atmospheric boundary layer wind tunnel (ABLWT) to predict the power production of wind turbines when the wind speed at a local meteorological tower is either known or forecasted. This method involves measuring the wind speed distributions over a scaled model of the terrain on which the turbines and the meteorological tower are located, and compiling this information into a database that is effectively a "power curve". The primary advantage with this method is that once the data is collected, it is easy to use, like the power curves based on historical data. An advantage of this method is that no historical data is needed to generate an ABLWT database, thus it could be applied to new wind farms, or even be used to site turbines for maximum power production on undeveloped terrain. For a more detailed explanation of the ABLWT based prediction method, see Cheng et al.⁵ and Lubitz and White⁶. For the case study wind farm, a series of wind-tunnel tests were performed, and a database compiled that allows the prediction of power as a function of wind speed and direction.

D. Comparison of the Power Curve Methods

Each of three methods described above produces a power curve for the case study wind farm. For comparison, the three power curves are shown in Fig. 8, along with the actual measurements of power production as a function of wind speed for the period July 1, 2003 to June 30, 2004. It is noted that the actual power production data exhibits a rather large amount of scatter relative to other wind farm power curves, perhaps due to a higher level of uncertainty in the wind speed measurements.

To quantify the power predicting accuracy of the three methods, each was used to predict the power of the case study wind farm for the one year period between July 1, 2003 and June 30, 2004. For each case, the ME and MAE between the actual power production and the production predicted by the method was determined. See Table 3. Results also were determined separately for the cases of actual power production of less than 20% of wind farm capacity, between 20% and 80%, and exceeding 80% of wind farm capacity. This was done to observe the degree of accuracy of the methods under different wind regimes, since it might be expected that it would be more difficult to forecast the mid-range power production, when power output varies rapidly as a function of wind speed, and less difficult to forecast the power production in low wind conditions.

VI. Forecasting Power: Overall Accuracy

Using the Eta model output, the "matching" method to forecast meteorological tower weather conditions, and the power curve for the wind farm derived using the medians of binned historical production data, an overall power




Fig. 8. Wind farm power production, as a percentage of the observed wind farm capacity, July 1, 2003 to June 30, 2004. Also includes power curve derived using multiple linear regression (Eq. 5, assuming time of January 1), power derived using binned, median power production values, and a power curve derived from wind-tunnel data assuming a wind direction of 240°.

	All Data 16859 obs.		Power < 20% 10006 obs.		20% < Power < 80% 4541 obs.		Power > 80% 2312 obs.	
	ME	MAE	ME	MAE	ME	MAE	ME	MAE
Median-Based Power Curve	0.1	5.9	1.9	3.2	-0.7	10.4	-6.1	8.7
MLR Power Curve	-0.2	6.7	2.5	5.0	-1.7	9.0	-8.7	10.0
Wind Tunnel Power Curve (w/o Wind Direction)	2.7	6.7	3.7	4.2	4.2	11.5	-4.5	8.3
Wind Tunnel Power Curve (w/ Wind Direction)	1.1	6.2	2.6	3.4	0.9	11.3	-4.8	8.5
Wind Tunnel Power Curve (w/ Wind Dir. & Density Cor.)	1.0	6.3	2.6	3.4	0.7	11.4	-5.1	8.7

Table 3. ME and MAE between actual power production and predicted production for each power curve. Wind-tunnel results include (1) power predicted without using wind direction (assuming wind from 240°, the most common wind direction on site), (2) incorporating the wind direction, and (3) incorporating both wind direction and an air density correction. Results are given for all observations, and for three categories of actual power production as a percentage of wind farm capacity.

forecast MAE of 18% of the observed wind farm capacity for the case study wind farm over the period of November 2002 to July 2004.

It is important to look carefully at the error statistics given for a wind power forecast. For example, at the case study wind farm, winds strong enough for power production only occur roughly half the time: this fact can make certain types of forecasts (i.e., those that under-predict the wind, and therefore the power) look "better" than others. Consider a very simple forecast: no power produced, at any time. Applying this to the case study wind farm, at the times corresponding to the Eta forecasts in the UC Davis archive (currently, November 30, 2002 to July 31, 2004), the power production MAE would be 27% of the observed wind farm capacity, which would seem respectable for a simple forecasting method. (It is noted that the ME also would be 27%.)

VII. Conclusions

A fairly significant amount of uncertainty enters the wind power forecast with the regional scale forecast (RSF), with the levels of the RSF model above the atmospheric boundary layer generally being more accurate than those near the surface. In complex terrain, predicting the surface wind conditions from the upper level meteorological conditions in the RSF also is challenging, more so than in simpler terrain. For the Altamont Pass case study, attempts to do this by applying the geostrophic drag law or using multiple linear regression to derive predictive

equations were found to be unsuccessful. Matching the current RSF to the most similar historical RSF in an archive was found to be a promising method of forecasting a near surface wind speed (and the wind farm power production).

With sufficient historical meteorological data available, a power curve can be derived to predict the power production of the wind farm from the meteorological tower wind speed with an expected MAE of about 5% of the wind farm capacity. If historical meteorological data is not available, the same level of accuracy can still be achieved using a database of wind-tunnel measurements to predict the power production of the wind farm based on meteorological conditions recorded at the farm meteorological tower^{5,6}. Since this method can be used without site specific adjustment (“tuning”), good predictions can be obtained even if no historical data is available. It also could be applied to siting of wind turbines at potential wind farm locations.

It appears that the majority of the error introduced into a wind power forecast comes from the RSF used as input data. Additionally, most of the remaining error arises from using the RSF data to predict reference wind conditions near the surface. Once near surface wind conditions are forecast, several methods can be used to forecast the wind power with only a small additional uncertainty (e.g., an MAE of roughly 6% of the wind farm capacity was found for the Altamont Pass site).

Acknowledgements

This work is being carried out as part of a multi-year wind power forecasting project sponsored by the Electric Power Research Institute (EPRI) and the California Energy Commission (CEC), involving the California Independent System Operator (CalISO), Lawrence Livermore National Laboratory (LLNL), TrueWind Solutions LLC, and the University of California, Davis (UCD).

The authors wish to thank Robert Szymanski of Powerworks, Inc., Charles “Chuck” McGowin of EPRI, George Simons and Dora Yen-Nakafuji of CEC, and Steve Chin of LLNL. This research was financially supported by the EPRI Wind Power Forecasting Project through the CEC PIER Program.

References

¹EPRI Technical Report 1007338. *California Wind Energy Forecasting System Development and Testing. Phase 1: Initial Testing*. Palo Alto, CA, USA, January 2003(a).

²EPRI Technical Report 1007339. *California Wind Energy Forecasting System Development and Testing. Phase 2: 12-Month Testing*. Palo Alto, CA, USA, July 2003(b).

³Landberg, L., Watson, S. J. "Short-Term Prediction of Local Wind Conditions." *Boundary-Layer Meteorology*. Vol. 70, pp. 171-195. 1994.

⁴Zilitinkevich, S. S. "Velocity Profiles, the Resistance Law and the Dissipation Rate of Mean Flow Kinetic Energy in a Neutrally and Stably Stratified Planetary Boundary Layer." *Boundary-Layer Meteorology*. Vol. 46, pp. 367-387. 1989.

⁵Cheng, J., Lubitz, W., White, B. "Wind-Tunnel Prediction of Wind Power Production in Complex Terrain." AIAA Paper No. 2004-1360. Proc. of AIAA/ASME Wind Energy Symposium 2004. Reno, NV, USA. Jan. 2004.

⁶Lubitz, W. D., White, B. R. "Prediction of Wind Power Production Using Wind-Tunnel Data: A Component of a Wind Power Forecasting System." *Proceedings of American Wind Energy Association Global Windpower 2004*. Chicago, IL, USA. Mar. 28-31, 2004.

⁷Burda, T. J., Wilson, J. W., Hiester, T. R., Germain, A. "Significance of Upper Air Meteorological Conditions in Altamont Pass Wind Resource Assessment." *Proc. Windpower '85*. San Francisco, CA, USA. Aug. 27-30, 1985.

⁸Wilks, D. S. *Statistical Methods in the Atmospheric Sciences*. Academic Press. 1995.