

IBM Research Report

Gust Speed Forecasting Using Weather Model Outputs and Meteorological Observations

Honggei Li, Fei Liu, Jonathan R. Hosking, Yasuo Amemiya

IBM Research Division

Thomas J. Watson Research Center

P.O. Box 218

Yorktown Heights, NY 10598



Research Division

Almaden - Austin - Beijing - Cambridge - Haifa - India - T. J. Watson - Tokyo - Zurich

Gust Speed Forecasting Using Weather Model Outputs and Meteorological Observations

Hongfei Li, Fei Liu, Jonathan R. M. Hosking, Yasuo Amemiya

IBM Thomas J. Watson Research Center

Yorktown Heights, NY

Abstract

Strong wind gust can cause severe damages, such as power outages, which are major concerns for emergency management of utility companies. Yet, gust speed is not a standard output from weather forecasting models. In this paper, we proposed a Bayesian hierarchical model which combines historical meteorological observations and weather model outputs to forecast hourly gust speed. The exploratory analysis suggested a two-step sequential spatio-temporal model as follows. In the first step, we calibrated sustained wind speed forecast and in the second step, we forecasted gust speed based on calibrated wind forecast. We demonstrated the use of our model by a real application concerning gust-caused damages for a local power utility company from New York area.

Key Words: Gust, Model calibration, Spatio-temporal, Wind

1 Introduction

This paper describes a statistical model to forecast gust speed, the maximum wind speed in every two seconds. Gust is a major factor to cause damages and thus, its prediction is critical in many real applications. For example, power outages due to strong wind gust are major concerns for the emergency management of electric utility companies. Strong wind gust blows off trees and branches. Consequently, the wires and poles fall down, leading to power failures and extensive damages including threat to human lives. One severe event was reported on September 16, 2010 in New York City. A high wind gust speed of 100 mph was recorded, causing major transportation issues and power outages in the metropolitan area. In the city alone, Con-Ed reported more than 4,500 outages on Staten Island and 27,000 in Queens (<http://newyork.cbslocal.com/2010/09/16/severe-storm-rips-through-new-york-city/>).

In the example of emergency management for electric utility companies, an effective gust forecast can help to gain extra time for proactive allocation and deployment of resources to minimize time to repair (Li et al., 2010). A new IBM weather modeling capability, dubbed "Deep Thunder" (DT) (Treinish and Praino, 2006), provides high-resolution forecasts for a region ranging from a metropolitan area up to an entire statewide calculations as fine as every mile. Forecasts are made up to one day ahead of time. The reports include various weather elements, such as wind speed, temperature, humidity and pressure. However, currently DT system does not forecast gust speed. Driven by this business concern of gust-caused damages, in this paper, we developed a statistical method to forecast hourly instantaneous gust speed, which integrates the DT weather forecast and the historical weather data.

Previous researches have been done on modeling sustained wind speed (e.g., Brown et al., 1984; Haslett and Raftery, 1989; Genton and Hering, 2007). Despite the amount of work on wind-forecast post process and ensemble forecasting (e.g., Slughter et al., 2010; Thorarinsdottir and Gneiting, 2010), very few work was devoted to forecasting gust speed which provides direct information to address the issues in emergency management. In meteorology, wind gust (shorted as gust) is a different concept from sustained wind (shorted as wind). For comparison, wind speed is the average of sampled wind speed in two minutes while gust speed is the

maximum sampled wind speed in two seconds. Even though gust and wind are two distinctive meteorological measures, some studies have suggested a linear relationship between sustained wind speed and gust speed (Deacon, 1965). This motivated our method to forecast gust speed based on sustained wind speed forecast.

Unlike the wind speed recorded instantaneously, the gust observations herein are collected accumulatively, *i.e.*, the current gust speed record is updated, if it exceeds the gust speed at previous time point; otherwise, the gust speed will not be updated. In other words, the gust observations are censored. This is an example of record values (Lin, 1987), where the research focus is to forecast the occurrence time of the next record. Different from previous work, we are interested in the inference of the underlying random process, from which the incomplete-observed record value is obtained. To be specific, we aim at forecasting the true hourly instantaneous gust speed based on its record value history. To achieve this goal, we developed the following two-step sequential hierarchical model as suggested by exploratory analysis. In the first step, we built a linear model to calibrate DT wind forecast using wind speed observations. We then, in the second step, forecasted the gust speed based on the calibrated DT wind speed. The resulting model, which incorporated information from DT wind forecast, wind speed and gust observations, allowed us to forecast the instantaneous gust speed.

The remainder of the paper is organized as follows. Section 2 introduces the data sets and the exploratory analysis which justifies the model we developed. Section 3 describes the two-step sequential model in a Bayesian hierarchical framework. The forecasting results of a real application are shown in Section 4 and the discussion in Section 5.

2 Data and Exploratory Analysis

2.1 Forecast and Observation Data

Deep Thunder (DT) system developed at the IBM Thomas J. Watson Research Center is a service that provides local high-resolution weather predictions customized to business applications for weather-sensitive operations. It uses a capability at the meso-scale (which corresponds to 1- to 2-km grid spacing, or cloud scale) to predict the

combination of weather conditions across the northeastern area of US. The service area is shown as blue grids on the left map of Figure 1. DT operates one-day-ahead forecast initialized at 00 hours UTC in international standard time, 8 p.m. eastern time in summer when daylight saving time operates, and 7 p.m. eastern time otherwise. The IBM DT system has been used for a utility company in the New York metropolitan area, to provide weather forecast which initiates the power outage damage forecast (Li et al., 2010). Our focus is to forecast the severe weather conditions that can disrupt the electrical distribution network of overhead lines. The wind gust forecast is critical for this purpose, which is, unfortunately, not forecasted directly from IBM DT system.

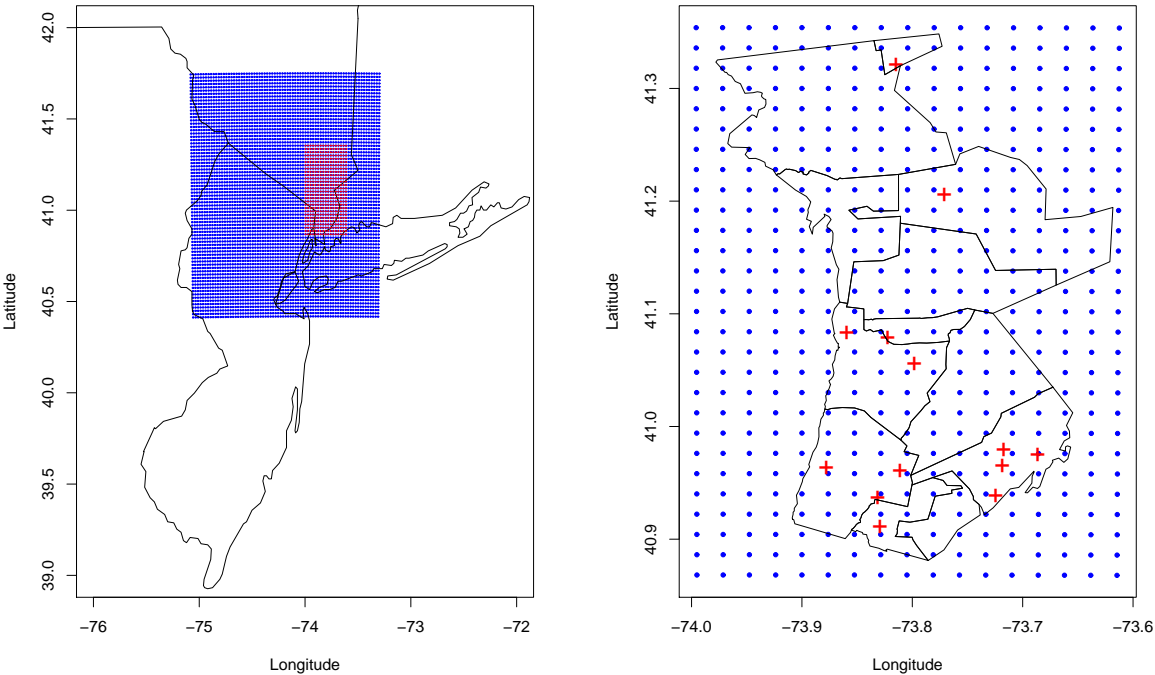


Figure 1: The left panel shows DT forecast domain (blue grids) in the northeastern area of the US and the subset of the grids (red grids) that covers the utility company service area that we are interested in. The right panel zooms in the grid area of interests. 13 weather stations where weather observations are collected are plotted as red pluses.

We extracted the DT weather forecast outputs at the selected grids that cover the service area of the utility company. These grids are marked red on the left panel of Figure 1. The DT model outputs report various weather elements, such as wind speed, temperature, humidity and pressure. The wind forecast was recorded

in knots (1 knot = 1.51 mile per hour (mph)) and converted to mph for analysis. The weather forecast data between July 1, 2009 and January 31, 2010 were selected to develop the model and the data between February 1, 2010 and February 17, 2010 were used for testing and validation.

Historical weather observations were provided by AWS Convergence Technologies, Inc., Germantown, MD, including hourly observations of surface wind speed (mph), accumulative gust speed (mph), temperature and etc. Weather observations were made at 13 weather stations located in the service area of the utility company. The stations are shown in the right panel of Figure 1, labeled by red crosses. To compare the DT weather forecast with real observations by locations, we interpolated the weather forecast data to those 13 weather station locations using bilinear interpolator, a common practice in the meteorological community.

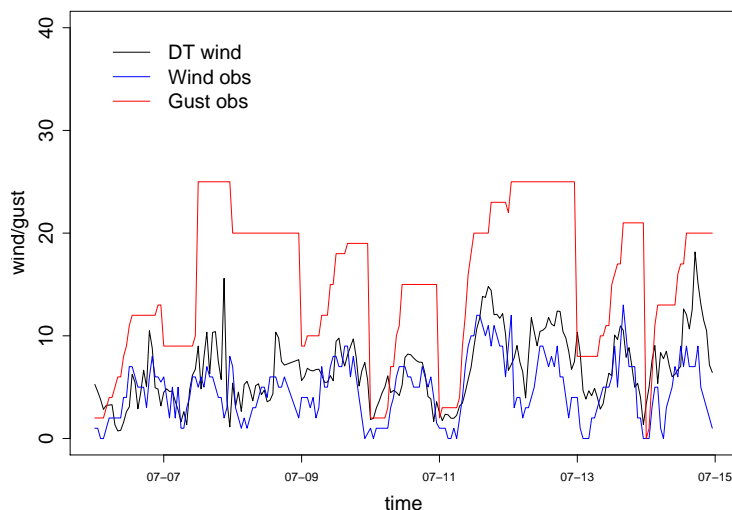


Figure 2: Plots of the hourly values of the accumulative gust speed observations (red curve), forecasted (black curve) and observed wind speeds (blue curve) from July 1, 2009 to July 15, 2009.

In Figure 2, we showed the hourly values of the accumulative gust speed observations (red curve), forecasted (black curve) and observed (blue curve) wind speeds from July 1, 2009 to July 15, 2009 as an illustration. From the plot, we can see the record value property of gust observations. Gust speed is an increasing step-function curve within 24 hours. Accumulative gust is only updated when the gust speed at current time exceeds the speed at previous time point or at midnight of each day. In our situation, the record value is reset at midnight for every 24 hours. Compared to gust observations, the DT wind forecast and wind observations are

instantaneously recorded.

2.2 Exploratory Analysis

In this section, we showed some results from exploratory data analysis. These results were illustrative and helpful to provide insights to our model building. Since gust observations are censored, there is no direct way to explore the relationship between gust speed and other variables. For this reason, we only picked the observed gust values which are record values and removed those censored gust values. Accordingly, we picked the observed wind and DT wind values at the time points that the gust speed was observed rather than censored.

- Exploration of relationship between wind speed and gust speed suggests that we can predict gust speed using a linear model based on wind observations if they are available.

We showed the scatterplots of wind observations and wind forecast versus observed gust speed in Figure 3. Through comparison, gust observations have stronger correlation with wind observations (left panel) than with wind forecast values (right panel). In addition, a strong linear relationship between gust speed and wind speed can be observed from these two plots.




Figure 3: Scatterplots of gust observations vs. wind observations (left panel) and gust observations vs. DT wind forecast (right panel).

- Exploration of relationship between DT wind forecasts and wind observations suggests a two-step sequential model.

Wind observations, in principle, are not available at the stage of predicting gust speed. Thus, we need to rely on DT wind forecast to some extent. We showed the scatterplot of the wind observations and the DT wind forecasts in Figure 4. DT wind forecast represents a strong linear relationship with wind observations which suggests a linear relationship between them. However, we also noticed the potential biases in the DT wind forecast. As we can see from the plot, the DT wind forecast is usually higher than the wind observations. For this reason, direct use of the DT wind forecasts to predict gust speed will lead to biased gust speed forecasts. More importantly, the gust speed is record valued and not completely observed. The inference of the true gust speed process heavily depends on the accuracy of the predictors. This exploratory analysis suggests the following two-step sequential modeling. In the first step, we calibrated the DT wind forecast against the historic wind observations. Then, given the calibrated forecast wind values combined with historical gust observations, we provided predication of gust speed.




Figure 4: Scatterplots of wind observations vs. DT wind forecast.

- Exploration of seasonality patterns of wind speed suggests inclusion of seasonality factors in the model.

In Figure 5, we compared the 24-hour patterns of DT wind speed forecasts (left panel) and the wind observations (right panel) through boxplots. As shown in the plots, it is apparent that wind speed is higher during the day, *i.e.*, 9am to 6pm, while the DT wind forecasts do not show such pattern. This

suggests we include 24-hour seasonality factors in the statistical model as means to capture the seasonal difference between the wind observations and DT forecasts.

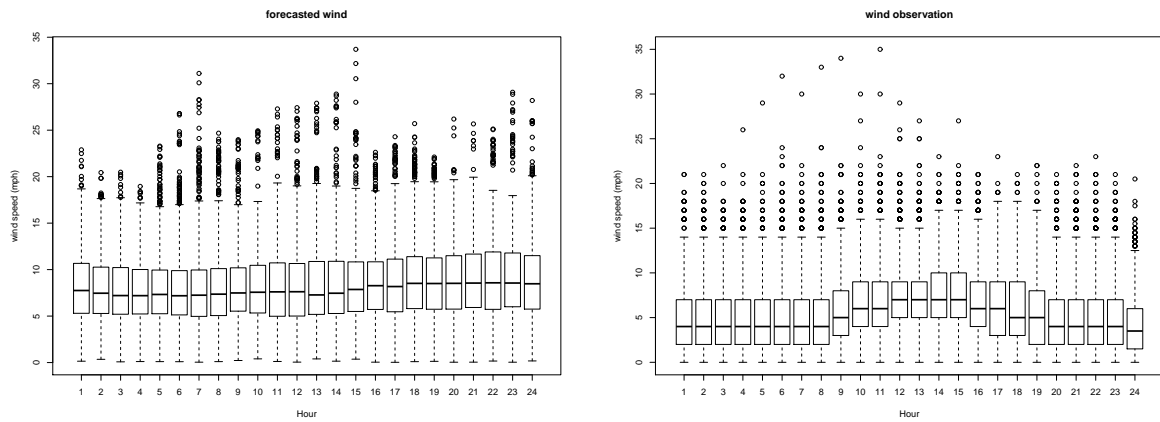


Figure 5: Exploration of 24-hour seasonality patterns of DT wind forecasts (left panel) and wind observations (right panel).

- Exploration of other meteorological variables suggests they are not significant.

We also considered including more meteorological forecast variables, such as temperature, humidity and pressure. Exploratory analysis showed that after including sustained wind speed as a predictor in the gust forecasting model, these variables are not significant and thus will not be considered in the model.

- Exploration of spatial heterogeneity suggests spatial varying coefficients.

Figure 6 explores whether the relationship between observed wind speed and DT forecasted wind speed varies over space. We fitted linear regression models of wind observations to DT wind forecasts for the data collected at each of the 13 different locations, respectively. Figure 6 showed that there were substantial differences among the resulting regression slope coefficients. To test the heterogeneity of location effects, we used Breusch-Pagan (BP) test (Breusch and Pagan, 1979). The BP test gave a very small p-value, which confirmed the heterogeneity due to spatial location impact. The same procedure was also performed to explore the relationship between gust speed and wind speed, with similar conclusion on the spatial heterogeneity.

- Exploration of spatial correlation of the regression coefficients at different locations suggests the existence




Figure 6: Exploration of spatial heterogeneity. A linear regression model of wind observations to DT wind forecasts was fitted for the data collected at each of the 13 different locations, respectively. This plot shows the estimates (circles in the plot) of the slope coefficients for each location and the corresponding 95% confidence levels (dotted lines) of the estimates.

of a spatial covariance structure.

Followed by the analysis of spatial heterogeneity, we studied the spatial correlation of the regression coefficient estimates at different locations. Figure 7 shows the semivariograms of the 13 slope estimates. As can be seen from the plot, there is a clearly spatial pattern between the coefficients, which motivated us to build a model with spatial varying coefficients (Gelfand et al., 2003).

3 Method

Throughout this paper, we are going to use the following notations:

- $t \in \mathcal{T} = \{1, \dots, T\}$: $t = 1$ refers to midnight of July 1, 2009, and T is the total number of hours that data covered;
- $\mathcal{U} = \{u_1, \dots, u_M\}$: the weather station locations and M is the number of weather stations;
- $x_t(u)$: the hourly wind speed forecast at location u and time t , which is interpolated using bilinear interpolation of the nearby forecast values, $u \in \mathcal{U}$ and $t \in \mathcal{T}$;




Figure 7: Semivariogram of the slope estimates from the linear regression of wind observations versus wind forecasts at 13 locations.

- $y_t(u)$: the wind observations at location u and time t , $u \in \mathcal{U}$ and $t \in \mathcal{T}$;
- $z_t(u)$: the original accumulative gust speed records at location u and time t , $u \in \mathcal{U}$ and $t \in \mathcal{T}$;
- $g_t(u)$: the true gust speed process at location u and time t , $u \in \mathcal{U}$ and $t \in \mathcal{T}$.

The gust speed record $z_t(u)$, collected hourly at different weather stations, was updated only at the midnight of each day or when the gust speed at time t exceeded the one at time $t - 1$. This implied the following relationship between $z_t(u)$ and $g_t(u)$ with $z_1(u) = g_1(u)$ at starting time $t = 1$,

$$z_t(u) = \begin{cases} g_t(u) & \text{if } g_t(u) \geq g_{t-1}(u) \text{ or } [(t-1)/24] = 0, \\ z_{t-1}(u) & \text{otherwise,} \end{cases}$$

where $[A/B]$ means the remainder after A is divided by B . As shown in the exploratory analysis in Section 2.2, the observed gust speed $g_t(u)$ is linearly related to the wind speed $y_t(u)$. This suggests the following model

$$g_t(u) = \delta + \eta(u)y_t(u) + \epsilon_t^g(u), \tag{1}$$

with $\epsilon_t^g(u) \sim N(0, \sigma_g^2)$. The exploratory analysis in Section 2.2 indicated that the slope coefficients vary over space. Therefore, we model $\boldsymbol{\eta} = \{\eta(u), u = 1, \dots, M\}$ as spatially varying coefficients,

$$\boldsymbol{\eta} \sim N(e\mathbf{1}_M, \boldsymbol{\Sigma}(\phi_\eta, \sigma_\eta^2)), \quad (2)$$

where $\boldsymbol{\Sigma}(\phi_\eta, \sigma_\eta^2)$ is the covariance matrix with the first parameter specifying spatial correlation and the second parameter specifying variance. The intercept coefficient, δ , does not show heterogeneity in space from our exploratory analysis and thus, we modeled it as a scalar parameter.

To guarantee that the gust speed value is nonnegative, we put constraints to parameters δ and $\eta(u), u = 1, \dots, M$ so that the estimates are positive. In addition, we investigated residual Q-Q plot of model (1) and found out that the normality assumption works reasonably well.

In reality, the wind speed observation can never be available in advance. Therefore, we had to rely on the DT wind forecasts. To predict gust speed, we proposed a two-step sequential modeling. In the first step, the wind forecast was calibrated using wind observations (Equation 3). Then given the calibrated wind forecast, we predict gust speed according to model (1). We modeled the wind forecast calibration process as

$$y_t(u) = \alpha_t + \gamma_{h(t)} + \beta(u)x_t(u) + \epsilon_t^y(u), \quad (3)$$

where $\epsilon_t^y(u) \sim N(0, \sigma_y^2)$ and $\gamma_{h(t)}$ is the hourly discrepancies, with $h(t) = [(t-1)/24]$ and $h(t) = 1, \dots, 24$.

In above, α_t incorporated temporal autocorrelation besides the seasonality and we assigned the following autoregressive structure as

$$\alpha_t = \rho\alpha_{t-1} + \epsilon_t^\alpha,$$

where $\epsilon_t^\alpha \sim N(0, \sigma_\alpha^2)$. Similar to the model of $\boldsymbol{\eta}$, we further modeled $\boldsymbol{\beta} = \{\beta(u), u = 1, \dots, M\}$ as spatially varying coefficients,

$$\boldsymbol{\beta} \sim N(b\mathbf{1}_M, \boldsymbol{\Sigma}(\phi_\beta, \sigma_\beta^2)), \quad (4)$$

where $\Sigma(\phi_\beta, \sigma_\beta^2)$ is the covariance matrix with the first parameter specifying spatial correlation and the second parameter specifying variance.

We run MCMC algorithm to estimate parameters. Within each iteration of the MCMC algorithm, we imputed the hourly gust speeds according to their full conditional posterior distributions. The algorithm works as follows:

Algorithm of imputation

At MCMC sampling iteration r ,

- If $z_t(u)$ is an observed value, $g_t^{(r)}(u) = z_t(u)$, where we use $\cdot^{(r)}$ superscript to represent the MCMC samples of the parameters at iteration r .
- If $z_t(u)$ is unobserved, given the MCMC samples of the unknown parameters at iteration r , sample $g_t^{(r)}(u)$ from truncated normal distribution, $TN(\delta^{(r)} + \eta^{(r)}(u)y_t(u), \sigma_g^{2(r)})$ with lower bound equal to 0 and upper bound equal to $z_t(u)$.
- Treat the sampled $g_t^{(r)}(u)$ as the true value, and use it to update the model parameters.

Forecast

To forecast the gust speed at the grids over the entire region, we need to obtain the estimates of the parameters at any location. Some of the parameters vary over space, such as $\beta = (\beta(u_1), \dots, \beta(u_M))'$. To provide interpolation for the parameter surface, taking the β as an example, we sampled $\beta(s^*)$ at a new location s^* from the posterior of the β process. To be specific,

$$f(\beta(s^*)|\mathbf{x}, \mathbf{y}) = \int f(\beta(s^*)|\beta, \sigma_\beta^2, \phi_\beta) f(\beta, \sigma_\beta^2, \phi|\mathbf{x}, \mathbf{y}). \quad (5)$$

The first density under the integral is a univariate normal that can be derived from (4). The second density is the posterior distribution of the parameters from which posterior samples can be obtained.

To predict the true wind, or calibrated wind, $y_t(s^*)$, given the wind forecast $x_t(s^*)$ and the past true wind till time t , denoted by $\mathbf{y}_{(t)} = (\mathbf{y}_{(t)}(u_1), \dots, \mathbf{y}_{(t)}(u_M))'$, we require

$$\begin{aligned} f(y_t(s^*)|x_t(s^*), \mathbf{y}_{(t)}) &= \int f(y_t(s^*)|\alpha_t, \gamma_{h(t)}, \beta(s^*), \sigma_y^2, x_t(s^*)) \\ &\quad \times f(\beta(s^*)|\boldsymbol{\beta}, \sigma_\beta^2, \phi_\beta) \\ &\quad \times f(\boldsymbol{\beta}, \alpha_t, \gamma_{h(t)}, \sigma_y^2, \sigma_\beta^2, \phi_\beta|\mathbf{x}, \mathbf{y}). \end{aligned} \quad (6)$$

The first density under the integral is given by Equation 3. Similar to the strategy of sampling $\beta(s^*)$, the samples of the other parameters can be obtained from the posterior distributions, the second and the third terms under the integral.

To predict $g_t(s^*)$ given the wind forecast $x_t(s^*)$, we have

$$\begin{aligned} f(g_t(s^*)|x_t(s^*)) &= \int f(g_t(s^*)|\delta, \eta(s^*), \sigma_g^2, y_t(s^*)) \\ &\quad \times f(y_t(s^*)|\alpha_t, \gamma_{h(t)}, \beta(s^*), \sigma_y^2, x_t(s^*)) \\ &\quad \times f(\beta(s^*)|\boldsymbol{\beta}, \sigma_\beta^2, \phi_\beta) \times f(\eta(s^*)|\boldsymbol{\eta}, \sigma_\eta^2, \phi_\eta) \\ &\quad \times f(\boldsymbol{\beta}, \alpha_t, \gamma_{h(t)}, \boldsymbol{\eta}, \boldsymbol{\delta}, \sigma_y^2, \sigma_\eta^2, \sigma_\beta^2, \phi_\beta, \phi_\eta|\mathbf{x}, \mathbf{y}). \end{aligned} \quad (7)$$

Again, it is straightforward to obtain samples from this predictive distribution.

4 Results

The model was developed using forecast and observations from July 1, 2009 to January 31, 2010. The data from February 1, 2010 to February 17, 2010 are used for testing and validation purpose. Given the sustained wind forecast, we calibrated the forecasted wind speed and then forecasted gust speed. In the left panel of Figure 8, we showed the comparison of the hourly forecasted values and observations at one weather station as an illustration. We can see that the wind forecast (blue curve) is usually higher than the wind observations (green

curve) which has been discussed in Section 2.2. The calibrated wind forecast is shown as a light blue curve which is closer to the real observations (green curve). The gust observations are plotted in black curve which shows the record value property and the instantaneous gust forecast is represented by a red curve. Comparing these two curves, we can see that when gust speed is unobserved, the gust forecast is usually lower than its value which is due to the fact that the gust observations are censored; when gust speed is updated, our gust model provides pretty good forecast compared to the observations.

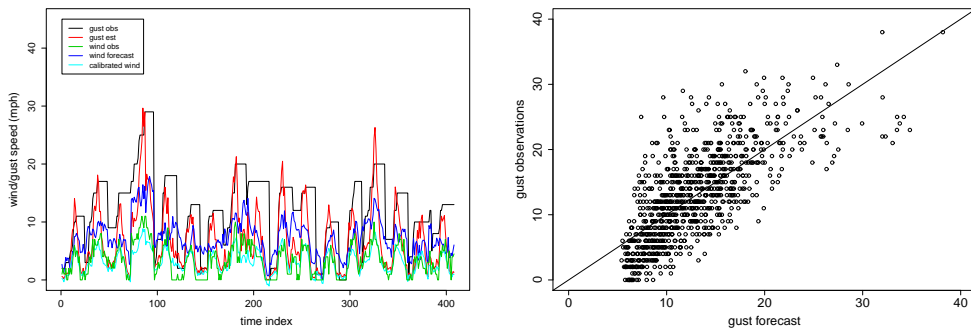


Figure 8: Evaluation of gust forecast. Left panel: comparison of gust forecast (red curve) with gust observations (black curve) and wind calibration (light blue curve) with wind observations (green curve) and DT wind forecast (blue curve) at one weather station. Right panel: scatterplot of gust forecast versus observed gust speed.

In addition to the comparison of gust forecast and gust records by locations, we also compared them by choosing the observed gust speed and the corresponding gust forecast for all of the locations in the right panel of Figure 8. Overall, the scatterplot shows that gust forecast lines up well with gust observations despite that the forecast is somewhat positively biased when the gust speed is low (e.g., less than 10mph). The bias is in part due to the fact that the DT model tends to overforecast when the wind speed is low. Furthermore, the forecast accuracy of low gust speed is not our major concern, since in the context of power-outage emergency management, the interest is the forecast accuracy for moderate to strong gust that is directly related to the damages. Our method provides reasonable accuracy in forecasting gust for this purpose.

In the left panel of Figure 9, we showed the estimated η surface. Recall that in (1), η are spatial varying regression coefficients to connect true wind speed with the gust speed. From the MCMC algorithm, we obtained the posterior samples of η at the 13 weather stations. Then we interpolated them to the whole η

surface and estimated the η surface using the posterior means from its posterior distributions. The plot shows that the coefficients vary substantially across space. Indeed, the estimates range between 1.1 and 1.9. Outside the hull of 13 weather stations, the η surface is sort of constant. This is expected due to the fact that the weather stations are located scarcely in that region. We can further obtain the gust forecast surface by plugging the interpolated η and β in (7). As an illustration, we showed the snapshot of gust forecast surface at 3am on Feb 12, 2010 in the right panel of Figure 9. Clearly, the gust speed varies dramatically over space. Our method provides a high-resolution gust forecast, which is critical information for dispatching facilities and personnel under extreme weather conditions. With our high-resolution gust forecast, utility companies can allocate more resources to the high risk areas. This provides an efficient solution to the emergency management.

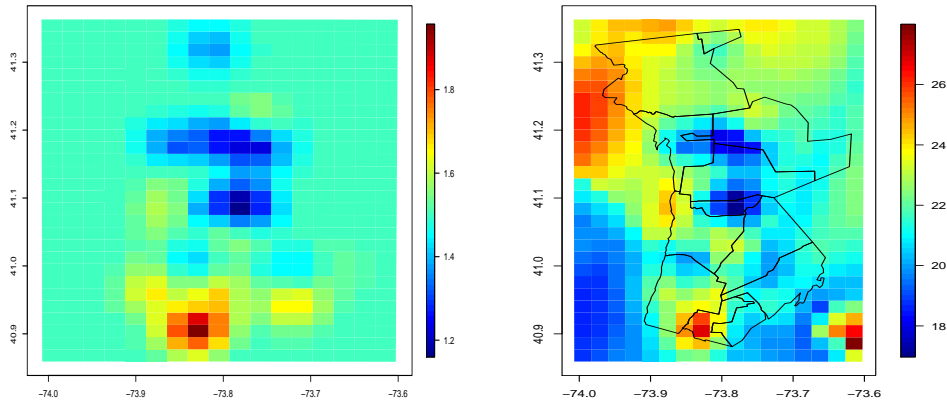


Figure 9: Maps of η surface (left panel) and gust forecast (right panel) at 3am on Feb 12, 2010.

5 Discussion

Motivated by the business need to forecast instantaneous gust speed using IBM Deep Thunder system, we developed a Bayesian hierarchical model to forecast gust speed. Our method combines information from multiple sources including the DT forecast data and historical weather information. The proposed model provides effective local high-resolution gust forecast, which are critical for weather-sensitive operations, such as emergency management for power outages caused by strong storm events. Our approach is currently being used by the

utility company in its distribution operations.

When the gust forecast is available, it can be fed into the damage forecasting model to provide outage prediction. It will be another interesting topic to model damage forecast while combined with the gust forecast in one framework. We will leave it for future work.

References

- Breusch, T. S. and Pagan, A. R. (1979). “A Simple Test for Heteroscedasticity and Random Coefficient Variation.” *Econometrica*, 47, 1287-1294.
- Brown, B. G., Katz, R. W., and Murphy, A. H. (1984). “Time Series Models to Simulate and Forecast Wind Speed and Wind Power.” *Journal of Climate and Applied Meteorology*, 23, 1184-1195.
- Deacon, E. L. (1965). “Wind Gust Speed: Averaging Time Relationship.” Tech. rep., C.S.I.R.O., Division of Meteorological Physics, Aspendale, Victoria, Australia.
- Gelfand, A. E., Kim, H., Sirmans, C. F., and Banerjee, S. (2003). “Spatial Modeling With Spatially Varying Coefficient Processes.” *Journal of the American Statistical Association*, 98, 387–396.
- Genton, M. and Hering, A. (2007). “Blowing in the Wind.” *Significance*, 4, 11–24.
- Haslett, J. and Raftery, A. E. (1989). “Space-Time Modelling with Long-Memory Dependence: Assessing Ireland’s Wind Power Resource.” *Applied Statistics*, 38, 1–50.
- Li, H., Treinish, L., and Hosking, J. (2010). “A statistical Model for Risk Management of Electric Outage Forecasts.” *IBM Journal of Research and Development*, 54.
- Lin, G. D. (1987). “On Characterizations of Distributions via Moments of Record Values.” *Probability Theory*, 74, 479–483.
- Sloughter, J. M., Gneiting, T., and Raftery, A. E. (2010). “Probabilistic Wind Speed Forecasting Using Ensembles and Bayesian Model Averaging.” *Journal of the American Statistical Association*, 105, 25–35.
- Thorarinsdottir, T. L. and Gneiting, T. (2010). “Probabilistic Forecasts of Wind Speed: Ensemble Model Output Statistics by Using Heteroscedastic Censored Regression.” *Journal of Royal Statistical Society, A*, 173, 371–388.

Treinish, L. A. and Praino, A. (2006). "The Role of Meso-Scale Numerical Weather Prediction and Visualization in Weather-Sensitive Decision Making." In *Forum Environ. Risk Impacts Society: Successes Challenges*. [Online]. Available: http://ams.confex.com/ams/Annual2006/techprogram/paper_102538.htm.