

Probabilistic Wind Gust Forecasting Using Nonhomogeneous Gaussian Regression

THORDIS L. THORARINSDOTTIR

Institute of Applied Mathematics, Heidelberg University, Heidelberg, Germany

MATTHEW S. JOHNSON

Department of Statistics, Oregon State University, Corvallis, Oregon

(Manuscript received 17 March 2011, in final form 18 July 2011)

ABSTRACT

A joint probabilistic forecasting framework is proposed for maximum wind speed, the probability of gust, and, conditional on gust being observed, the maximum gust speed in a setting where only the maximum wind speed forecast is available. The framework employs the nonhomogeneous Gaussian regression (NGR) statistical postprocessing method with appropriately truncated Gaussian predictive distributions. For wind speed, the distribution is truncated at zero, the location parameter is a linear function of the wind speed ensemble forecast, and the scale parameter is a linear function of the ensemble variance. The gust forecasts are derived from the wind speed forecast using a gust factor, and the predictive distribution for gust speed is truncated according to its definition. The framework is applied to 48-h-ahead forecasts of wind speed over the North American Pacific Northwest obtained from the University of Washington mesoscale ensemble. The resulting density forecasts for wind speed and gust speed are calibrated and sharp, and offer substantial improvement in predictive performance over the raw ensemble or climatological reference forecasts.

1. Introduction

Accurate forecasts of wind speed are essential in many applications, such as energy production, aviation, recreational boating, and agriculture. They are of particular importance for severe weather warnings issued for the general public. High winds are known to cause damage to infrastructure, especially sudden outbursts of wind gusts. Wind gusts also influence the design and structure of buildings and bridges, as wind-induced loads on the structure can be expected to cause damage when the gusts exceed a given threshold. However, available wind gust observations are usually sparse, given their small spatial and temporal scales, and wind gust forecasts are not a standard output from numerical weather prediction (NWP) models.

Historically, gust speed forecasts are obtained from sustained wind speed forecasts using gust factors (Durst 1960). The gust factor is defined as the ratio of the maximum gust speed to the sustained wind speed, and its

value may depend on the measure for sustained wind speed, the magnitude of the wind speed, as well as altitude and terrain (Ágústsson and Ólafsson 2004; Jungo et al. 2002; Weggel 1999). More general forecasting models have also been proposed, in which the gust speed is modeled as a linear function of the wind speed forecast with spatially varying parameters [see, e.g., Etienne et al. (2010) and Li et al. (2010)].

The aim of this paper is to show how, in everyday situations, a probabilistic forecast for wind gust may be obtained based on an ensemble forecast for maximum wind speed and a gust factor that provides the connection between the predictive maximum wind speed and the predictive gust speed. Probabilistic forecasts are exceptionally attractive for weather quantities such as gust, as wind gusts are closely related to various kinds of weather related risk. Different risk factors will depend on different gust speed thresholds, and, with a fully probabilistic forecast, the probability of the gust speed exceeding each threshold can easily be calculated. We propose that a probabilistic wind and gust speed forecast be issued in three parts, similar to a precipitation forecast. That is, at a given location, a forecast would be issued for the following three quantities:

Corresponding author address: Thordis L. Thorarinsdottir, Institute of Applied Mathematics, Heidelberg University, Im Neuenheimer Feld 294, D-69120 Heidelberg, Germany.
E-mail: thordis@uni-heidelberg.de

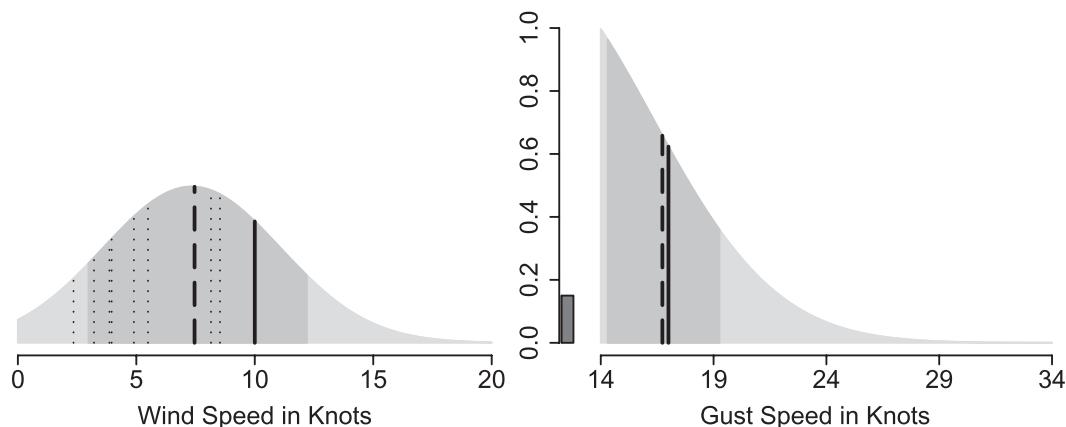


FIG. 1. The 48-h-ahead wind and gust speed forecasts valid on 19 Mar 2008, at Pearson Field Airport in Vancouver, WA. (left) Probabilistic forecast for wind speed with a median at 7.5 kt (dashed line), the eight ensemble forecasts ranging from 2.3 to 8.5 kt (dotted lines), and the observation at 10 kt (solid line). (middle) Probability of gust forecast at 15%. (right) Probabilistic forecast for gust speed with a median at 16.7 kt (dashed line) and the observation at 17 kt (solid line). The 80% prediction intervals for wind and gust speed are indicated with a darker shade of gray.

- 1) daily maximum wind speed;
- 2) probability of gust; and
- 3) conditional on gust being observed, the daily maximum gust speed.

An example of such a forecast with the corresponding observations is given in Fig. 1.

Two methods have recently been proposed for the first step: Sloughter et al. (2010) model the predictive wind speed distribution as a mixture of gamma distributions using Bayesian model averaging (BMA), while Thorarinsdottir and Gneiting (2010) employ nonhomogeneous Gaussian regression (NGR) using a Gaussian distribution truncated at zero. Both methods provide calibrated and sharp probabilistic forecasts for the maximum wind speed, and, for the application discussed in these papers, the predictive performance of the two methods is essentially equal. Here, we build upon the NGR framework of Thorarinsdottir and Gneiting (2010), as we can immediately derive a probabilistic forecast for gust speed from the wind speed forecast based only on the assumption of a multiplicative relationship between wind speed and gust speed. This connection is not immediate under the mixture model assumption of BMA.

The remainder of the paper is organized as follows. In section 2, we review the NGR technique and describe our extension of it to wind gust forecasting. In section 3, we present the results of a case study in which we issue daily 48-h-ahead forecasts of maximum wind and gust speed over the North American Pacific Northwest in 2008 based on the eight-member University of Washington mesoscale ensemble (Eckel and Mass 2005). Finally, a discussion is provided in section 4.

2. Nonhomogeneous Gaussian regression

The NGR or ensemble model output statistics (EMOS) approach was originally introduced by Gneiting et al. (2005) for temperature and sea level pressure. The NGR approach models the predictive distribution of the weather quantity of interest conditional on the ensemble forecast by a Gaussian distribution with a location parameter that is a linear function of the individual ensemble members and a spread parameter that is an affine function of the ensemble variance.

In Thorarinsdottir and Gneiting (2010), the NGR method is extended to allow for nonnegative variables such as maximum wind speed. Specifically, let Y denote wind speed and let X_1, \dots, X_k denote an ensemble of individually distinguishable ensemble member forecasts for Y . To address the nonnegativity of the predictand, the predictive distribution is modeled by a truncated Gaussian distribution with a cutoff at zero,

$$Y|X_1, \dots, X_k \sim \mathcal{N}^0(\mu, \sigma^2), \quad (1)$$

where the location parameter is given by $\mu = a + b_1X_1 + \dots + b_kX_k$ and the spread parameter by $\sigma^2 = c + dS^2$. Here, $S^2 = (1/k)\sum_{i=1}^k(X_i - \bar{X})^2$ denotes the ensemble variance with $\bar{X} = (1/k)\sum_{i=1}^kX_i$. The density of this distribution is given by

$$f^0(y) = \left[\frac{1}{\sigma} \varphi\left(\frac{y - \mu}{\sigma}\right) \right] / \Phi\left(\frac{\mu}{\sigma}\right) \quad (2)$$

for $y > 0$ and $f^0(y) = 0$ otherwise, where φ and Φ denote the standard Gaussian density and the standard Gaussian cumulative distribution function (CDF), respectively. For clarity, we have omitted the dependence on the forecast

ensemble in the definition of the density. The truncated Gaussian predictive distribution is conditional on the dynamical ensemble, and as such it differs from the unconditional (marginal) distribution of wind speed, which is typically modeled with Weibull densities (see, e.g., He et al. 2010).

As described in Thorarinsdottir and Gneiting (2010), the parameters of the predictive distribution in (1) can be estimated using a minimum continuous ranked probability score (CRPS) estimation over a rolling training period. That is, on a given day, the parameters of the predictive distribution are estimated using all available forecasts and observations from the previous n days. We here employ the regional version of the method, where the training data from all stations are pooled together and a single set of parameter estimates is obtained for a given day. This constitutes the first step in our forecasting scheme.

Furthermore, let Z denote gust speed and assume that the relationship between Y and Z can be modeled with a gust factor γ ,

$$Z = \gamma Y, \tag{3}$$

for $\gamma \geq 1$. By combining (1) and (3), we obtain a predictive distribution for Z given by

$$Z|X_1, \dots, X_k \sim \mathcal{N}^0(\gamma\mu, [\gamma\sigma]^2). \tag{4}$$

This can easily be seen by noting that the conditional predictive mean and variance of Y are given by $\mathbb{E}(Y|X_1, \dots, X_k) = \mu + \sigma\lambda(-\mu/\sigma)$ and $\mathbb{V}(Y|X_1, \dots, X_k) = \sigma^2[1 - \delta(-\mu/\sigma)]$, respectively, where $\lambda(x) = \varphi(x)/[1 - \Phi(x)]$ and $\delta(x) = \lambda(x)[\lambda(x) - x]$.

Steps two and three in our forecasting scheme depend on the definition of gust that is used to obtain the gust speed measurements. In our case study, we use observation data from the North American Automated Surface Observing System (ASOS) network. Our forecasts are therefore based on the definition of gust that is used in the ASOS algorithm (National Weather Service 1998). Under certain technical conditions, a 5-s average wind speed is stored in memory as a gust for 10 min. If the maximum gust in memory is greater than the current 2-min average wind speed by at least 3 kt (1 kt = 0.514 m s⁻¹), the current wind speed is more than 2 kt, and the maximum gust in memory is at least 10 kt greater than the minimum 5-s average wind speed over the last 10 min, then the maximum gust in memory is designated as the reportable gust. The minimum gust speed reported by ASOS is 14 kt.

The measurements we obtain from the ASOS network only allow us to reconstruct the last condition. We will thus focus on this condition and assume that gust is observed if the gust speed is at least 14 kt. However, we

also see that wind gust is only observed if the variation in wind speed between the peaks and the lulls reaches a certain limit. We therefore expect to observe lower gust factors when the wind gust is not observed, although it is difficult to investigate this issue in detail, as the high-resolution data needed for such an investigation are not saved in the ASOS network. Based on these considerations we estimate two separate gust factors, one for the probability of gust forecasts and another for the conditional gust speed forecasts, conditional on gust being observed.

From the predictive distribution in (4), the probability that a gust will be observed becomes

$$\mathbb{P}(Z \geq 14|X_1, \dots, X_k) = \left[1 - \Phi\left(\frac{14 - \gamma_1\mu}{\gamma_1\sigma}\right)\right] / \Phi\left(\frac{\mu}{\sigma}\right). \tag{5}$$

Given parameter estimates for the wind speed predictive distribution in (1), we estimate the parameter γ_1 by minimizing the Brier score (Brier 1950) for (5) over the training set. That is, we find the value of γ_1 that minimizes

$$\frac{1}{n} \sum_{j=1}^n \left\{ \left[1 - \Phi\left(\frac{14 - \gamma_1\mu_j}{\gamma_1\sigma_j}\right)\right] \Phi\left(\frac{\mu_j}{\sigma_j}\right)^{-1} - \mathbb{1}\{Z_j \text{ observed}\} \right\}^2,$$

where the sum extends over the forecast cases in the training set. Note that the location parameter μ_j and scale parameter σ_j depend on j only through the ensemble member forecasts.

Conditional on gust being observed, the predictive density for gust speed is given by

$$f^0(z)\mathbb{1}\{z \geq 14\} = f^{14}(z) = \left[\frac{1}{\gamma_2\sigma} \varphi\left(\frac{z - \gamma_2\mu}{\gamma_2\sigma}\right) \right] / \left[1 - \Phi\left(\frac{14 - \gamma_2\mu}{\gamma_2\sigma}\right) \right], \tag{6}$$

for $z \geq 14$ and $f^{14}(z) = 0$ if $z < 14$. This is the density of a truncated Gaussian distribution with a cutoff at 14. We estimate the parameter γ_2 by minimizing the CRPS (Matheson and Winkler 1976; Hersbach 2000) over the gust speed observations in the training set given the parameter estimate for the wind speed forecast. For a general predictive distribution F and an observation y , the CRPS is given by

$$\text{crps}(F, y) = \int_{-\infty}^{+\infty} [F(x) - \mathbb{1}\{x \geq y\}]^2 dx.$$

When F is a truncated Gaussian distribution with a cutoff at 14, this expression becomes

$$\text{crps}[\mathcal{N}^{14}(\alpha, \beta), y] = \beta \lambda(\alpha, \beta)^{-2} \left\{ \frac{y - \alpha}{\beta} \lambda(\alpha, \beta) \left[2\Phi\left(\frac{y - \alpha}{\beta}\right) + \lambda(\alpha, \beta) - 2 \right] + 2\phi\left(\frac{y - \alpha}{\beta}\right) \lambda(\alpha, \beta) - \frac{1}{\sqrt{\pi}} \left[1 - \Phi\left(\sqrt{2}\frac{14 - \alpha}{\beta}\right) \right] \right\},$$

where $\lambda(\alpha, \beta) = 1 - \Phi[(14 - \alpha)/\beta]$ and $y \geq 14$.

The parameter estimation for both γ_1 and γ_2 relies on numerical optimization, which is performed with the Broyden–Fletcher–Goldfarb–Shanno algorithm (Bertsekas 1999) as implemented in R (R Development Core Team 2010). Despite this, the computational cost is very low, as only one parameter is estimated in each numerical optimization process. An example of the resulting probabilistic forecasts is shown in Fig. 1.

3. Case study

a. Forecast and observation data

We consider forecasts and observation data from 83 observation stations in the ASOS network in the United States and Canada, as illustrated in Fig. 2. The forecasts in our database are 48-h-ahead forecasts of maximum wind speed, initialized daily at 0000 UTC, using the eight-member University of Washington Mesoscale Ensemble (UWME) system (Eckel and Mass 2005). Forecasts at observation sites are created by bilinear interpolation from the model grid. We use verifying observations of two different measures of wind speed, the daily maximum wind speed and the daily maximum gust speed. Daily maximum wind speed is defined as the maximum of an hourly 2-min average wind speed over 12 h of the day, prior to 0000 UTC. That is, once every hour, a 2-min average wind speed is recorded, and at 0000 UTC the maximum of the 12 most recent observations is recorded as the daily maximum wind speed. The daily maximum gust speed is also taken over 12 h of the day prior to 0000 UTC (its definition is given in section 2). All observations are rounded to the nearest whole knot when recorded at the meteorological station, except that wind speeds below 1 kt are recorded as zero. The data were subject to the quality control procedures described by Baars (2005).

b. Results

The results presented below are based on data from 1 January to 31 December 2008 from the 83 ASOS stations shown in Fig. 2. For this time period, data are available for 291 days, for a total of 22 863 individual forecast cases. The results are based on a rolling training period of length 20 days for all three forecasts. To test the robustness of our method, we have also performed

the same analysis using training periods of length 15, 20, . . . , 50 days for the wind speed forecasts, and for each of these values we have generated gust speed forecasts using training periods of length 15, 20, . . . , 50 days. The method seems very robust against these changes, as the performance scores presented in Table 1 change by less than 0.5% for wind speed and by less than 1.0% for gust speed, which is in accordance with the results of Thorarinsdottir and Gneiting (2010). Data from the calendar year 2007 were used to obtain rolling training periods of length 20 days for the first days in the test set.

1) GUST FACTORS

The NGR gust factor estimates for the predictive models in (5) and (6) are shown in Fig. 3. On all days, a lower gust factor estimate was obtained for the probability of gust forecasts in (5) than for the gust speed forecast in (6), which is conditional on wind gust being observed. This is in concurrence with the definition of wind gust and the related discussion in section 2. For the probability of gust forecast, the estimated values for γ_1 range from 1.12 to 1.33 with an average of 1.24, while the estimated values for γ_2 range from 1.45 to 1.69 with an average of 1.52. Furthermore, as can be seen in Fig. 3, there seems to be very low correlation between the daily estimates for the two gust factor parameters. A seasonal pattern is observable for γ_1 with the estimated values generally higher during the summer months than during winter. The opposite seems to hold for γ_2 , where the values are generally lower during the summer months with shorter periods of higher values during the colder seasons.

2) PROBABILITY OF GUST

For the calendar year 2008, gust was observed in 8324 forecast cases, or in 36% of all forecast cases. Figure 4 shows the reliability diagram (Wilks 2006) for the probability of gust forecast under the NGR model in (5) and for the UWME forecast. The UWME forecast is based on a consensus vote from the raw ensemble multiplied with a gust factor. That is, the probability of gust is forecasted as the proportion of ensemble members that exceed the value 14 kt after multiplication with a gust factor. The gust factor is estimated by minimizing the Brier score over the same training set as was used for the corresponding NGR forecast. The NGR forecasts are well calibrated for all except the very high probabilities,



FIG. 2. The locations of the 83 ASOS stations (small black dots) used in this study. The stations are located in the North American Pacific Northwest, including the Canadian provinces of British Columbia (BC) and Alberta (AB), and the U.S. states of Washington (WA), Oregon (OR), Idaho (ID), California (CA), and Nevada (NV).

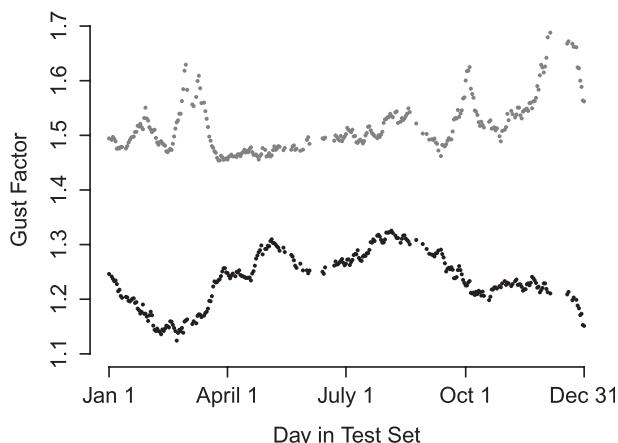


FIG. 3. Daily NGR gust factor estimates for the probability of gust predictive model in (5) (black dots) and the gust speed predictive model in (6) (gray dots) over all available days in 2008.

As an example of a daily forecast, Fig. 5 shows the probability of gust forecasts valid on 19 March 2008, under the NGR model. On this day, the forecast probabilities vary from 6% to 98%, and we see a general correspondence between a higher probability of gust forecast and gust being observed.

3) WIND SPEED AND GUST SPEED

In assessing the probabilistic forecasts of wind and gust speed, we follow Gneiting et al. (2007) and aim to

while the UWME is uncalibrated in that it underestimates the lower probabilities and overestimates the higher probabilities. This is in accordance with the results of Thorarinsdottir and Gneiting (2010), wherein it is shown that the UWME for wind speed is highly underdispersive in that too many observations fall either below or above all the ensemble member forecasts.

TABLE 1. Mean CRPS, MAE, coverage, and average width of the 77.8% prediction intervals for probabilistic forecasts of wind and gust speed over the Pacific Northwest in 2008. The upper part of the table shows the results for the wind speed forecasts and the lower part the results for the gust speed forecasts when gust was observed. The MAE refers to the point forecast given by the median of the respective predictive distribution.

Forecast	CRPS (kt)	MAE (kt)	Coverage (%)	Width (kt)
Wind speed climatology	2.79	3.94	80.7	11.5
UWME	2.64	3.23	45.3	4.9
NGR	2.20	3.06	78.6	9.2
Gust speed climatology	3.08	4.33	81.6	12.8
UWME with gust factors	3.72	4.68	52.9	8.6
NGR with gust factors	2.56	3.60	76.9	10.5

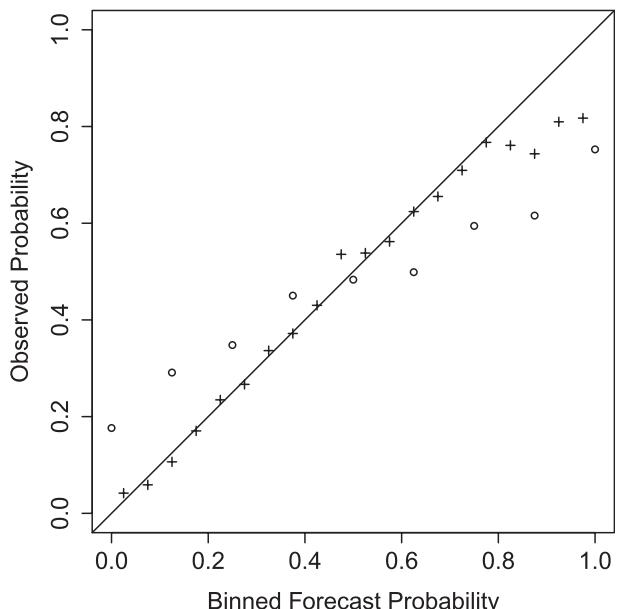


FIG. 4. Reliability diagram for the NGR forecast (plus symbol) and the UWME forecast combined with gust factors (open circle) of binned probability of gust forecasts vs observed relative frequency of gusts over the Pacific Northwest in 2008.

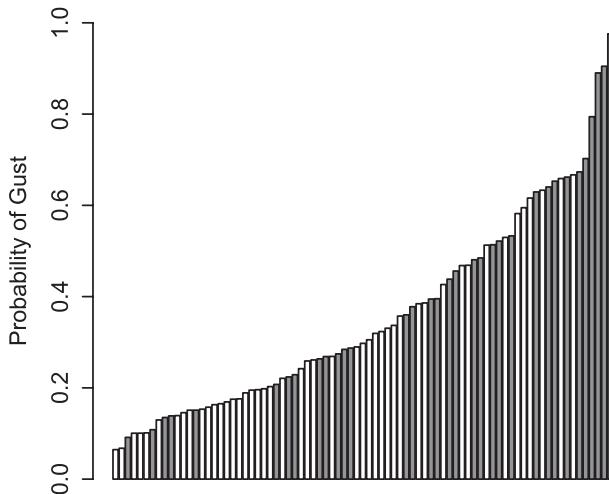


FIG. 5. The 48-h-ahead probability of gust forecasts under the NGR model over the Pacific Northwest valid on 19 Mar 2008, in increasing order. On this day, data and forecasts were available at 81 stations, and gust was observed at 34 stations. Bars corresponding to stations where gust was observed are shaded in dark gray.

maximize the sharpness of the predictive distributions subject to calibration. Calibration refers to the statistical consistency between the forecast distribution and the observations. For discrete distributions such as the UWME, calibration can be assessed with a verification rank histogram where the rank of each observed value is plotted relative to the eight ensemble member forecasts (Anderson 1996; Hamill and Colucci 1997). The verification rank histogram for gust speed forecasts over the Pacific Northwest in 2008 is shown in Fig. 6a, where the UWME forecast values have been multiplied with gust factors that are estimated over a training period using minimum CRPS estimation as described for the NGR model in section 2. The UWME is highly underdispersive, with too many observations falling below all the ensemble values; if the ensemble members and the observations were exchangeable, the histogram should be close to uniform.

The probability integral transform (PIT) histogram is the continuous analog of the verification rank histogram (Dawid 1984; Gneiting et al. 2007). The PIT is the value that the predictive cumulative distribution function attains at the observation, and, again, deviations from uniformity indicate lack of calibration. The PIT histogram for the NGR predictive distribution for gust speed given in (6) is shown in Fig. 6b. The NGR method significantly improves the calibration of the ensemble, and the PIT histogram is nearly uniform.

The calibration of the forecasts can further be assessed by the coverage of prediction intervals. In Table 1, we report the coverage of the 77.8% prediction intervals for

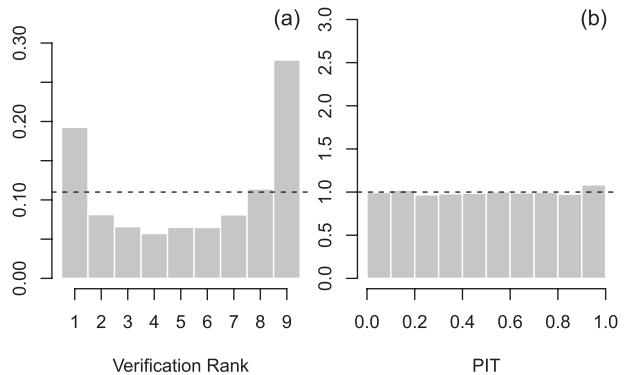


FIG. 6. Calibration checks for probabilistic forecasts of gust speed over the Pacific Northwest in 2008: (a) verification rank histogram for the UWME forecast multiplied with gust factors; (b) PIT histogram for the NGR probabilistic forecast given in (6).

both wind and gust speed forecasts aggregated over the Pacific Northwest in 2008. We have selected this particular prediction interval because it corresponds to the probability that the observation falls within the range of the ensemble for a perfectly calibrated ensemble. Similarly, the sharpness of a probabilistic forecast can be assessed by calculating the average width of the prediction intervals. In addition, to provide the results for the UWME forecasts and the NGR forecasts, Table 1 also shows the results for climatological reference forecasts for both wind and gust speed. The reference forecasts are obtained from the observed values for wind or gust speed, in the training period used to estimate the parameters of the NGR model. The climatological reference forecasts are well calibrated, yet the prediction intervals are very wide on average. The NGR forecasts show slightly better calibration than the climatological forecasts and are much sharper.

Scoring rules provide summary measures of predictive performance. In Table 1, we report two negatively oriented scores, the mean absolute error (MAE), based on the median of the predictive distribution, and the mean CRPS, aggregated over the Pacific Northwest and the year 2008. Both scores are attractive in that their values are given in the units of the observations. For wind speed, we see that the NGR method performs significantly better than the raw ensemble, which again significantly improves the climatology forecast. For gust speed, the NGR method shows the best predictive performance, which, when measured by the CRPS, is slightly better than the raw ensemble for wind speed, while the climatology outperforms the UWME.

Figure 7 shows the gust speed forecasts under the NGR model and the corresponding observations at the 34 stations in our test set where gust was observed on 19 March 2008. On this day, the observed gust speeds range

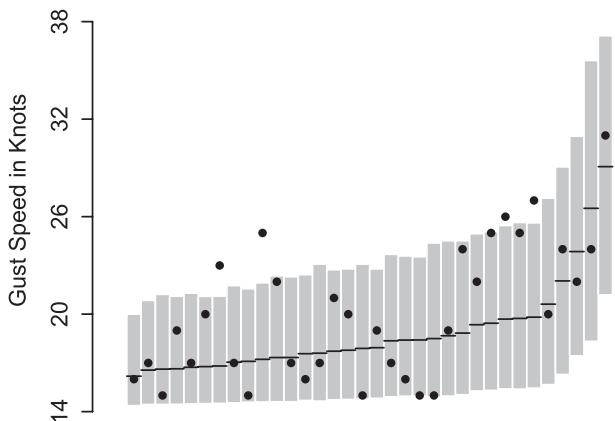


FIG. 7. The 48-h-ahead gust speed forecasts under the NGR model valid on 19 Mar 2008. Forecasts for the 34 stations in the Pacific Northwest where gust was observed are shown in increasing order. The gray bars show the 80% prediction intervals, and each forecast median is indicated with a solid line. Observed gust speeds are indicated with black dots.

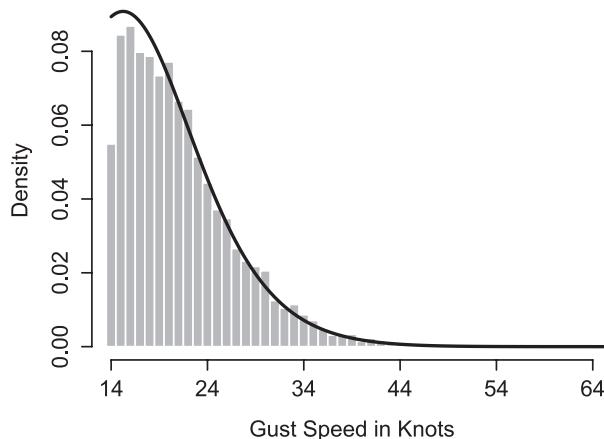


FIG. 8. The empirical density of observed gust speed over all stations in the Pacific Northwest in 2008 (in gray) and the marginal predictive density under the NGR model (solid line). Two observed values, at 79 and 151 kt, fall outside the interval shown.

from 15 to 31 kt, and the medians of the corresponding predictive distributions range from 16.2 to 29.1 kt. The 80% prediction intervals are generally wider for higher values of the predictive median, indicating increased forecast uncertainty for higher gust speeds. Figure 1 shows the full predictive distribution at Pearson Field Airport in Vancouver, Washington, valid on this day, and in Fig. 5 the corresponding probability of gust forecasts under the NGR model are given.

4) MARGINAL DENSITY FOR GUST SPEED

As mentioned in section 2, the NGR predictive density for gust speed given in (6) is conditional on the forecast ensemble, and it thus differs from the marginal density for gust speed. However, the marginal predictive density for gust speed under the NGR model can be constructed by calculating the average density forecast over all days in 2008 and all stations where gust was observed. In Fig. 8, the marginal density forecast for gust speed under the NGR model is shown together with the empirical distribution of the corresponding observed gust speeds. The marginal predictive density seems to exhibit a good fit to the empirical density. This conclusion is further supported by Fig. 9 which shows a quantile–quantile (Q–Q) plot comparing the quantiles of the two distributions shown in Fig. 8.

4. Discussion

We have shown how the nonhomogeneous Gaussian regression method to statistically postprocess ensemble forecasts for wind speed (Thorarinsdottir and Gneiting 2010) can be combined with a gust factor to obtain full

predictive distributions for gust speed. The predictive distributions for gust speed are based on the NGR parameter estimates for wind speed, and only the estimation of the gust factor parameters is required. The gust speed forecasts are thus obtained at a very low additional computational cost. This follows from the distribution assumptions of the NGR method, which are unaffected by a multiplication with a constant. Using the UWME (Eckel and Mass 2005), we apply the NGR

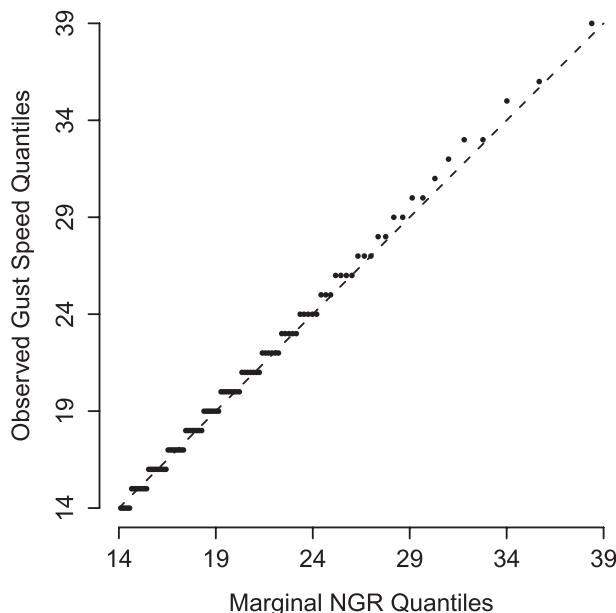


FIG. 9. A Q–Q plot comparing the quantiles from the marginal predictive distribution under the NGR model and the observed quantiles for gust speed over all stations in the Pacific Northwest in 2008.

method to create 48-h-ahead forecasts for wind and gust speed over the North American Pacific Northwest. In our forecasting framework, we issue three different forecasts: a probabilistic forecast for daily maximum wind speed; a forecast for the probability of wind gust being observed; and, conditional on wind gust being observed, a probabilistic forecast for the daily maximum gust speed. The resulting forecasts are well calibrated, and the predictive densities are sharp in that the prediction intervals are much shorter on average than prediction intervals based on climatological reference forecasts.

In his seminal work, Durst (1960) showed that the gust factor may depend on the magnitude of the mean wind speed, and similar results have been obtained by, for example, Ágústsson and Ólafsson (2004) and Jungo et al. (2002). Here, we estimate a constant daily gust factor over all locations within the study area, which makes such stratification by magnitude difficult. A different option would be to estimate a location-specific gust factor at each observation site separately; this would allow for the estimated gust factor to depend on the magnitude of the wind speed forecast ensemble. However, we have found this local approach to yield lower predictive performance, both in terms of performance scores and calibration, than the approach applied in section 3. While Thorarinsdottir and Gneiting (2010) showed that the local approach works well for probabilistic wind speed forecasts, it is difficult to obtain good results for gust speed because the observations are sparser in time and the marginal distribution can have very heavy tails.

A more flexible framework would allow the parameters to vary between locations while imposing constraints on the spatial structure; gust speed models with spatially varying parameters have been proposed by Etienne et al. (2010) and Li et al. (2010). Potential extension of the NGR method in this direction would be to model the parameters as a realization of a spatial random field, and similar extensions have been proposed for Bayesian model averaging applied to temperature and precipitation forecasts (Kleiber et al. 2011a,b). The assumption of a multiplicative relationship between wind speed and gust speed can furthermore be relaxed to include linear relationships, as a linear function of a random variable that follows a truncated Gaussian distribution will also follow a truncated Gaussian distribution. That is, a positive constant might be added to the right-hand side of (3). The requirement that the constant be positive is needed in order to ensure a positive support for the gust speed distribution function, as the addition of a constant will change the cutoff of the distribution.

Gust factors have also been shown to depend on altitude, terrain, and synoptic weather type (Ágústsson and Ólafsson 2004; Jungo et al. 2002). Using out-of-sample

data, we have investigated the observed gust factors in our dataset as a function of the altitude and the land use category of the observation stations. However, this has not led to conclusive results that could improve the predictive performance of the NGR method. While the observation data from the ASOS stations are generally of very high quality and thus desired for such a study, the stations are usually located at airports or, more precisely, at airport runway touchdown zones (National Weather Service 1998). Wind will therefore often have an unobstructed field close to the station, and the local topography is likely to be fairly flat. Instead, we have chosen to estimate daily gust factors, thereby allowing the gust factor estimates to change by seasons and/or prevailing weather patterns.

We have focused on gust speed forecasting in everyday situations, and the results indicate that our forecasting framework works very well in such situations. Furthermore, the marginal predictive density for gust speed under the NGR model offers a good fit to the marginal density of gust speed. When the aim of the statistical analysis is the modeling or forecasting of extreme wind gusts only, extreme value statistics are often employed [see, e.g., Walshaw and Anderson (2000) and Friederichs et al. (2009)]. An interesting further development would be to investigate the manner in which our framework can be applied to the forecasting of extreme events and what role extreme value statistics would play in this context.

Other approaches to gust speed forecasting include that of Brasseur (2001), who proposes a gust speed prediction method that is fully based on physical considerations. The approach assumes that the surface gusts result from the deflection of air parcels flowing higher in the boundary layer that are brought down by turbulent eddies, and it takes the sustained wind speed into account. The approach has been applied with some success using both NWP models (Ágústsson and Ólafsson 2009) and regional climate models (Goyette et al. 2003). However, the parameterization of the physical model depends on parameters that are usually not included in the output from an NWP model. For a successful application of the method, it thus needs to be implemented as a part of the NWP model. The NGR method can readily be applied as a statistical postprocessing method for an ensemble of gust speed forecasts obtained from an NWP model. Here, the gust factor estimates will not be needed, and, instead, a model similar to that in (1) but with a cutoff at 14 kn can be directly employed to post-process the gust speed forecast ensemble.

Acknowledgments. The authors are grateful to Tilmann Gneiting, Sebastian Lerch, Cliff Mass, Adrian Raftery,

Chris Fraley, and J. McLean Sloughter for helpful discussion and useful comments. The authors are also grateful to Jeff Baars for providing data and consultation during the production of this manuscript. This work was supported by National Science Foundation Award DMS-0706745.

REFERENCES

- Ágústsson, H., and H. Ólafsson, 2004: Mean gust factor in complex terrain. *Meteor. Z.*, **13**, 149–155.
- , and —, 2009: Forecasting wind gusts in complex terrain. *Meteor. Atmos. Phys.*, **103**, 173–185.
- Anderson, J. L., 1996: A method for producing and evaluating probabilistic forecasts from ensemble model integrations. *J. Climate*, **9**, 1518–1530.
- Baars, J., cited 2005: Observations QC summary page. [Available online at http://www.atmos.washington.edu/mm5rt/qc_obs/qc_obs_stats.html.]
- Bertsekas, D. P., 1999: *Nonlinear Programming*. 2nd ed. Athena Scientific, 802 pp.
- Brasseur, O., 2001: Development and application of a physical approach to estimating wind gusts. *Mon. Wea. Rev.*, **129**, 5–25.
- Brier, G. W., 1950: Verification of forecasts expressed in terms of probability. *Mon. Wea. Rev.*, **78**, 1–3.
- Dawid, A. P., 1984: Statistical theory: The prequential approach (with discussion and rejoinder). *J. Roy. Stat. Soc.*, **147A**, 278–292.
- Durst, C. S., 1960: Wind speeds over short period of time. *Meteor. Mag.*, **89**, 181–187.
- Eckel, A. F., and C. F. Mass, 2005: Aspects of effective mesoscale, short-range ensemble forecasting. *Wea. Forecasting*, **20**, 328–350.
- Etienne, C., A. Lehmann, S. Goyette, J. Lopez-Moreno, and M. Beniston, 2010: Spatial predictions of extreme wind speeds over Switzerland using generalized additive models. *J. Appl. Meteor. Climatol.*, **49**, 1956–1970.
- Friederichs, P., M. Göber, S. Bentzien, A. Lenz, and R. Krampitz, 2009: A probabilistic analysis of wind gusts using extreme value statistics. *Meteor. Z.*, **18**, 615–629.
- Gneiting, T., A. E. Raftery, A. H. Westveld, and T. Goldman, 2005: Calibrated probabilistic forecasting using ensemble model output statistics and minimum CRPS estimation. *Mon. Wea. Rev.*, **133**, 1098–1118.
- , F. Balabdaoui, and A. E. Raftery, 2007: Probabilistic forecasts, calibration and sharpness. *J. Roy. Stat. Soc.*, **69B**, 243–268.
- Goyette, S., O. Brasseur, and M. Beniston, 2003: Application of a new wind gust parameterization: Multiscale case studies performed with the Canadian regional climate model. *J. Geophys. Res.*, **108**, 4371–4389.
- Hamill, T. M., and S. J. Colucci, 1997: Verification of Eta-RSM short-range ensemble forecasts. *Mon. Wea. Rev.*, **125**, 1312–1327.
- He, Y., A. H. Monahan, C. G. Jones, A. Dai, S. Biner, D. Caya, and K. Winger, 2010: Probability distributions of land surface wind speeds over North America. *J. Geophys. Res.*, **115**, D04103, doi:10.1029/2008JD010708.
- Hersbach, H., 2000: Decomposition of the continuous ranked probability score for ensemble prediction systems. *Wea. Forecasting*, **15**, 559–570.
- Jungo, P., S. Goyette, and M. Beniston, 2002: Daily wind gust speed probabilities over Switzerland and according to three types of synoptic circulation. *Int. J. Climatol.*, **22**, 485–499.
- Kleiber, W., A. E. Raftery, J. Baars, T. Gneiting, C. Mass, and E. P. Grimit, 2011a: Locally calibrated probabilistic temperature forecasting using geostatistical model averaging and local Bayesian model averaging. *Mon. Wea. Rev.*, **139**, 2630–2649.
- , —, and T. Gneiting, 2011b: Geostatistical model averaging for locally calibrated probabilistic quantitative precipitation forecasting. *J. Amer. Stat. Assoc.*, in press.
- Li, H., F. Liu, J. R. Hosking, and Y. Amemiya, 2010: Gust speed forecasting using weather model outputs and meteorological observations. IBM Research Rep. RC25087 (W1012-060), 18 pp.
- Matheson, J. E., and R. L. Winkler, 1976: Scoring rules for continuous probability distributions. *Manage. Sci.*, **22**, 1087–1096.
- National Weather Service, cited 1998: Automated Surface Observing System (ASOS) user's guide. [Available online at <http://www.weather.gov/asos/aum-toc.pdf>.]
- R Development Core Team, cited 2010: R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria. [Available online at <http://www.R-project.org/>.]
- Sloughter, J. M., T. Gneiting, and A. E. Raftery, 2010: Probabilistic wind speed forecasting using ensembles and Bayesian model averaging. *J. Amer. Stat. Assoc.*, **105**, 25–35.
- Thorarinsdottir, T. L., and T. Gneiting, 2010: Probabilistic forecasts of wind speed: Ensemble model output statistics using heteroskedastic censored regression. *J. Roy. Stat. Soc.*, **173A**, 371–388.
- Walshaw, D., and C. W. Anderson, 2000: A model for extreme wind gusts. *J. Roy. Stat. Soc.*, **49C**, 499–508.
- Weggel, J. R., 1999: Maximum daily wind gusts related to mean daily wind speed. *J. Struct. Eng.*, **125**, 465–468.
- Wilks, D. S., 2006: *Statistical Methods in the Atmospheric Sciences*. 2nd ed. Academic Press, 648 pp.