

Comparison of nonhomogeneous regression models for probabilistic wind speed forecasting

Sebastian Lerch* and Thordis L. Thorarinsdottir†

May 10, 2013

In weather forecasting, nonhomogeneous regression is used to statistically postprocess forecast ensembles in order to obtain calibrated predictive distributions. For wind speed forecasts, the regression model is given by a truncated normal distribution where location and spread are derived from the ensemble. This paper proposes two alternative approaches which utilize the generalized extreme value (GEV) distribution. A direct alternative to the truncated normal regression is to apply a predictive distribution from the GEV family, while a regime switching approach based on the median of the forecast ensemble incorporates both distributions. In a case study on daily maximum wind speed over Germany with the forecast ensemble from the European Centre for Medium-Range Weather Forecasts, all three approaches provide calibrated and sharp predictive distributions with the regime switching approach showing the highest skill in the upper tail.

1 Introduction

Reliable forecasts of wind speed are a necessity in a diverse number of applications such as agriculture, most modern means of transportation and wind energy production. Wind power, as a renewable and emissions free alternative to fossil fuels, has been growing rapidly over the last decade. In Europe, the wind power's share of total installed power capacity amounted to about 11.4% at the end of 2012 and it has increased five-fold since 2000 (European Wind Energy Association, 2012). For wind energy production, accurate forecasts of wind speed at different lead times are required to regulate electricity markets, to schedule maintenance and, more generally, to improve the competitiveness of wind power compared to sources of electricity which allow for dispatchable generation (Genton and Hering, 2007; Pinson et al., 2007; Lei et al., 2009). In many of these applications and for weather warnings, high wind speeds are of particular importance.

The focus of this paper are daily forecasts with medium-range lead times of 1-3 days. In this setting, forecasts are usually based on outputs from numerical weather prediction

*Institute of Applied Mathematics, Heidelberg University, Im Neuenheimer Feld 294, 69120 Heidelberg, Germany

†Norwegian Computing Center, P.O. Box 114, Blindern, 0314 Oslo, Norway

(NWP) models which use physical descriptions of the atmosphere and oceans to propagate the state of the atmosphere forward in time based on the current weather conditions. Moreover, to account for uncertainties in the knowledge of the initial state of the atmosphere and the numerical model, the NWP models are often run several times with different initial conditions and/or numerical representations of the atmosphere resulting in an ensemble of forecasts (Gneiting and Raftery, 2005; Leutbecher and Palmer, 2008). The development of ensemble prediction systems plays a key role in the transition from deterministic to probabilistic forecasting and has become an established part of weather and climate prediction (Palmer, 2002).

Probabilistic forecasts are essential in many applications in that they allow for quantification of the associated prediction uncertainty. Further, optimal decision-making requires probabilistic forecasts, particularly for rapidly fluctuating resources such as wind energy where the optimal point forecast depends on permanently changing market features (Pinson et al., 2007; Thorarinsdottir and Gneiting, 2010). See Gneiting (2011) for a detailed discussion of optimal deterministic forecasts based on probabilistic forecasts. While ensemble systems are valuable in this context, they are finite and do not provide full predictive distributions. Also, they tend to be underdispersive and subject to a systematic bias, and thus require some form of statistical postprocessing (Hamill and Colucci, 1997; Gneiting and Raftery, 2005; Gneiting et al., 2007).

Statistical postprocessing methods for ensembles of wind speed forecasts include the ensemble Bayesian model averaging (BMA) method of Slughter et al. (2010) and the nonhomogeneous regression (NR) or ensemble model output statistics (EMOS) approach developed by Thorarinsdottir and Gneiting (2010). BMA predictions are given by weighted mixtures of parametric densities or kernels each of which depends on a single ensemble member, with the mixture weights determined based on the performance of the ensemble members in the training period. For wind speed, Slughter et al. (2010) apply a mixture of gamma distributions, see also Courtney et al. (2013). The NR method of Thorarinsdottir and Gneiting (2010), on the other hand, applies a single normal distribution truncated at zero, where the location parameter is an affine function of the ensemble members and the scale parameter is an affine function of their variance. In a comparison study, the two methods show very similar predictive performance (Thorarinsdottir and Gneiting, 2010). The NR method has been extended to wind gusts (Thorarinsdottir and Johnson, 2012) and a BMA approach for wind direction is proposed in Bao et al. (2010). Pinson (2012), Schuhen et al. (2012) and Slughter et al. (2013) study statistical postprocessing of bivariate wind vector ensembles.

Hourly average wind speeds are usually modeled using lognormal, gamma, Rayleigh or Weibull densities, with the Weibull model showing the best performance in many case studies, see for example Garcia et al. (1998) and Celik (2004). Here, we consider forecasts of daily maximum wind speed and the predictive distributions are conditioned on the ensemble forecast, the situation for which the postprocessing approaches of Slughter et al. (2010) and Thorarinsdottir and Gneiting (2010) were developed. Since daily maximum wind speeds are block maxima, results from extreme value theory imply that the generalized extreme value (GEV) distribution provides a suitable model (Coles, 2001). GEV distributions have especially received attention in modeling maxima of wind and gust speed observations over long return periods, typically 50 years, see Palutikof et al. (1999) and references therein. Friederichs and Thorarinsdottir (2012) apply a GEV model for probabilistic predictions of daily peak wind speed.

We propose to combine the NR ensemble postprocessing framework originally proposed by Gneiting et al. (2005) and later extended by Thorarinsdottir and Gneiting (2010) with

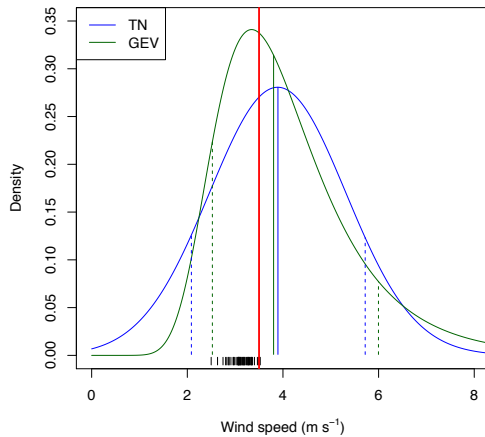


Figure 1: One day ahead forecasts for daily maximum wind speed at Frankfurt Airport valid on 19 March 2011. The ECMWF ensemble forecasts are indicated in black, the observation in red. The solid blue and green lines indicate the median of the truncated normal (TN) and the GEV predictive distribution, respectively. The dashed lines indicate the corresponding central 80% prediction intervals.

results from extreme value theory. That is, we apply a predictive distribution from the GEV family, where the location and scale parameters depend on the ensemble member forecasts. To illustrate the difference between the two NR approaches, Figure 1 shows the predictive distributions for Frankfurt Airport on 19 March 2011. Both the truncated normal (TN) and the GEV model correct the negative bias and the underdispersion of the ensemble, while the GEV density is less symmetric and exhibits a heavy right tail. We further investigate a regime switching approach which issues a TN predictive density when the ensemble median forecast takes a low value and a GEV predictive density for high values of the ensemble median.

The remainder of the paper is organized as follows. In Section 2, the ensemble forecasts and the observational data are introduced. In Section 3, we review the NR technique and describe our extensions using GEV distributions and a regime switching combination model. Section 4 summarizes the probabilistic scores used for estimating the model parameters and evaluating the competing forecasting procedures. In particular, we discuss how appropriately weighted proper scoring rules recently proposed in the economic literature (Gneiting and Ranjan, 2011; Diks et al., 2011) can be used to assess the predictive performance for high wind speeds. In Section 5, we report the results of a case study on daily maximum wind speed over Germany for lead times of 1-3 days with the ensemble issued by the European Centre for Medium-Range Forecasts from May 2010 to April 2011. We close with a discussion in Section 6.

2 Data

We consider 50 ensemble member forecasts of near-surface (10-meter) wind speed obtained from the global ensemble prediction system of the European Centre for Medium-Range Weather Forecasts (ECMWF). Ensemble forecasts for lead times up to 10 days ahead are issued twice a day at 00 UTC and 12 UTC, with a horizontal resolution of about 33 km and a temporal resolution of 3-6 hours. To account for uncertainties in the initial conditions and the numerical model, the ensemble members are generated from random




Figure 2: Map of Germany showing the locations of the 228 synoptic observation stations used in this study. The station at Frankfurt Airport is indicated with +.

perturbations in initial conditions and stochastic physics parametrization (Molteni et al., 1996; Leutbecher and Palmer, 2008; Pinson and Hagedorn, 2012). The ensemble members are thus statistically indistinguishable and can be treated as exchangeable (Fraley et al., 2010). We restrict attention to the ECMWF ensemble run initialized at 00 UTC and lead times of 1-3 days. To obtain predictions of daily maximum wind speed, we take the daily maximum of each ensemble member at each grip point location. For instance, one day ahead forecasts are given by the maximum over lead times of 3, 6, \dots , 24 hours.

The forecasts are verified over a set of 228 synoptic observation stations over Germany, see Figure 2. The observations are hourly observations of 10-minute average wind speed which is measured over the 10 minutes before the hour. To obtain daily maximum wind speed, we take the maximum over the 24 hours corresponding to the time frame of the ensemble forecast. In the following, the terms wind speed and daily maximum wind speed are used synonymously. Ensemble forecasts at individual stations are obtained by bilinear interpolation of the gridded model output. The results presented below are based on a verification period from 1 May 2010 to 30 April 2011, consisting of 83 220 individual forecast cases. Additionally, we use data from 1 February 2010 to 30 April 2011 to obtain training periods of equal lengths for all days in the verification period and for model selection purposes.

3 Nonhomogeneous regression prediction models

The nonhomogeneous regression (NR) methodology was originally developed for sea level pressure and surface temperature under a normal predictive distribution (Gneiting et al., 2005), see also Hagedorn et al. (2008) and Kann et al. (2009) for further applications. Thorarinsdottir and Gneiting (2010) extend the framework to wind speed using a normal distribution truncated in zero while Thorarinsdottir and Johnson (2012) apply the same set-up to predict gust speeds based on NWP forecasts of wind speed and gust factors. A bivariate normal model for wind vectors is discussed in Schuhen et al. (2012). An NR framework for quantitative precipitation has recently been proposed by Scheuerer (2012) using a GEV model for the precipitation accumulations.

3.1 Truncated normal model

Let Y denote wind speed and X_1, \dots, X_k the corresponding ensemble member forecasts. The predictive distribution for Y is given by a truncated normal (TN) distribution with a cutoff at 0,

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

where the location parameter $\mu = a + b_1 X_1 + \dots + b_k X_k$ is an affine function of the ensemble forecasts and the variance $\sigma^2 = c + dS^2$ is an affine function of the ensemble variance $S^2 = \frac{1}{k} \sum_{i=1}^k (X_i - \bar{X})^2$ with $\bar{X} = \sum_{i=1}^k X_i$. As the ECMWF ensemble members are exchangeable, we assume that $b_1 = \dots = b_k$, or $\mu = a + b\bar{X}$ (Fraley et al., 2010). The distribution function of the TN distribution is given by

$$F(z) = \Phi\left(\frac{\mu}{\sigma}\right)^{-1} \Phi\left(\frac{z - \mu}{\sigma}\right)$$

for $z > 0$ and 0 otherwise, where Φ denotes the cumulative distribution function of the standard normal distribution.

3.2 Generalized extreme value model

As an alternative to the TN model in (1), we consider a model based on extreme value theory. The cumulative distribution function of the GEV distribution with location parameter μ , scale parameter σ and shape parameter ξ is given by

$$G(z) = \begin{cases} \exp\left\{-\left[1 + \xi\left(\frac{z - \mu}{\sigma}\right)\right]^{-1/\xi}\right\} & \xi \neq 0 \\ \exp\left\{-\exp\left[-\left(\frac{z - \mu}{\sigma}\right)\right]\right\} & \xi = 0. \end{cases} \quad (2)$$

This distribution is defined on the set $\{z \in \mathbb{R} : 1 + \xi(z - \mu)/\sigma > 0\}$, where the parameters satisfy $\mu, \xi \in \mathbb{R}$ and $\sigma > 0$. For $\xi > 0$, G is of Fréchet type with a heavy right tail and it holds that $z \in [\mu - \sigma/\xi, \infty)$. We obtain the Fréchet type in approximately 99.5% of our forecast cases. We estimate the parameters of the model in (2) without any constraints on the parameter values. It is thus possible to obtain non-zero probabilities of negative wind speed. However, we find that this rarely happens in practice. The probability of negative wind speed is larger than 1% in about 0.1% of the forecast cases and it never exceeds 5%.

To link the parameters of the predictive GEV distribution to the ensemble, we apply the Bayesian covariate selection algorithm described in Friederichs and Thorarinsdottir (2012). We assume a constant shape parameter ξ while the location μ and the scale σ may depend on the ensemble mean and variance. Based on an analysis of the out-of-sample data from 1 February to 30 April 2010, we let the parameters depend on the ensemble mean only and set $\mu = \mu_0 + \mu_1 \bar{X}$, $\sigma = \sigma_0 + \sigma_1 \bar{X}$. While the average posterior inclusion probability for the ensemble mean is very high for both location and scale, the average posterior inclusion probability for the ensemble variance amounts to less than 0.1% for both parameters. The results were obtained for 100 000 iterations of the Metropolis within Gibbs algorithm with a burn-in period of 20 000 iterations.

3.3 Regime switching combination model

The third model we consider is a regime switching method which combines the TN approach in (1) and the GEV approach in (2). Conditional on the median of the ensemble predictions,

$$X^{med} = \text{median}(X_1, \dots, X_k),$$

we either issue a TN or a GEV predictive distribution. That is, for a model threshold $\theta \in \mathbb{R}_+$, we define the predictive distribution by

$$H = \begin{cases} \mathcal{N}_{[0,\infty)}(\mu^{\mathcal{N}}, \sigma^{2\mathcal{N}}), & \text{if } X^{med} < \theta \\ G(\mu^G, \sigma^G, \xi^G), & \text{if } X^{med} \geq \theta. \end{cases} \quad (3)$$

Here, the parameters of the TN and GEV model depend on the ensemble forecast as described above. However, we train the TN model only on training data for which it holds that $X^{med} < \theta$. Similarly, the parameters of the GEV distribution are learned from data where $X^{med} \geq \theta$. The model threshold θ is selected by comparing predictive performance over a range of possible thresholds based on the out-of-sample data from 1 February to 30 April 2010. Generally, thresholds between 7 and 8 m s⁻¹ prove optimal which approximately corresponds to the 75th and 85th percentile of the median ensemble predictions over the verification period. These results are discussed in detail below. Under this model, the probability of negative wind speed is less than 1.4×10^{-5} for all forecast cases.

4 Parameter estimation and prediction verification

The aim of the prediction is to “maximize the sharpness of the predictive distribution subject to calibration” (Gneiting et al., 2007). Calibration is a joint property of the predictive distribution and the associated observation. It essentially requires that the observation is indistinguishable from a random draw from the predictive distribution. Sharpness refers to the concentration of the predictive distribution and is a property of the forecasts only.

Anderson (1996) and Hamill and Colucci (1997) propose verification rank (VR) histograms as a graphical tool to assess the calibration of ensemble predictions. VR histograms show the distribution of the ranks of the observations when pooled within the ordered ensemble predictions. For a calibrated ensemble, the observations and the ensemble predictions should be exchangeable, resulting in a uniform VR histogram. The continuous analog of the VR histogram is the probability integral transform (PIT) histogram (Dawid, 1984; Gneiting et al., 2007). The PIT is the value of the predictive cumulative distribution function at the corresponding realizing observation. Again, for calibrated forecasts, the PIT values should follow a uniform distribution.

4.1 Proper scoring rules

Scoring rules assign numerical values to forecast-observation pairs and provide summary measures of predictive performance. Forecasting methods can be compared in this manner by averaging their scores over a test set. If the scoring rule evaluates the full predictive distribution, it can simultaneously address calibration and sharpness. A scoring rule is proper if the expected score is minimized when the true distribution of the observation is issued as the forecast (Bröcker and Smith, 2007; Gneiting and Raftery, 2007). Proper scores thus prevent hedging strategies.

Popular examples of proper scoring rules are the logarithmic or ignorance score (Good, 1952),

$$\text{LogS}(F, y) = -\log(f(y)), \quad (4)$$

where f denotes the density of F and y denotes the corresponding observation, and the continuous ranked probability score (CRPS) (Hersbach, 2000; Gneiting and Raftery, 2007),

$$\text{CRPS}(F, y) = \int_{-\infty}^{\infty} (F(z) - \mathbb{1}\{y \leq z\})^2 dz, \quad (5)$$

where the distribution F is assumed to have a finite first moment. Again, y denotes the corresponding observation. We furthermore use the absolute error $|x - y|$ for the point forecast x given by the median of the predictive distribution as a deterministic measure of accuracy. The median of the predictive distribution is the Bayes predictor under the absolute error loss function (Gneiting, 2011). All these scoring rules are negatively oriented, that is, the smaller, the better.

4.2 Evaluation of forecasts for high wind speeds

Despite the variety of theoretically justifiable methods to evaluate probabilistic forecasts, it is not obvious how to assess the predictive performance in the tails of the distribution, for example in the case of extreme wind speed observations. A natural approach is to select extreme events while discarding non-extreme events, and to proceed using standard evaluation procedures. However, it can be shown that restricting proper scoring rules to subsets of events results in improper scoring rules. This approach is thus bound to discredit even the most skillful forecasters (Gneiting and Ranjan, 2011). Instead, weighted scoring rules that emphasize specific regions of interest can be constructed.

Gneiting and Ranjan (2011) propose the threshold-weighted continuous ranked probability score (twCRPS),

$$\text{twCRPS}(F, y) = \int_{-\infty}^{\infty} (F(z) - \mathbb{1}\{y \leq z\})^2 w(z) dz, \quad (6)$$

where $w(z)$ is a non-negative weight function on the real line. For $w(z) \equiv 1$, the twCRPS reduces to the original CRPS in (5). If the interest lies in the right tail of the distribution, we may for example set $w(z) = \mathbb{1}\{z \geq r\}$. Note that the CRPS in (5) represents an integral of the Brier score (Brier, 1950) over all possible thresholds. The twCRPS in (6) with $w(z) = \mathbb{1}\{z \geq r\}$ thus allows us to simultaneously assess the exceedance probabilities for all thresholds greater or equal to r . Similarly, Diks et al. (2011) propose proper weighted versions of the logarithmic score in (4).

4.3 Optimum score estimation

The framework of proper scoring rules may also be applied to parameter estimation. Following the general optimum score estimation approach of Gneiting and Raftery (2007), the parameters of a distribution are determined by optimizing the average value of a proper scoring rule as a function of the parameters over a training set. Optimum score estimation based on the logarithmic score in (4) thus corresponds to maximum likelihood (ML) estimation. Minimum CRPS estimation, that is, optimum score estimation based on the CRPS in (5), provides a robust alternative to ML estimation if closed form expressions for the CRPS of the distribution family of interest are available.

Following Thorarinsdottir and Gneiting (2010), we estimate the parameters of the TN model using minimum CRPS estimation. Friederichs and Thorarinsdottir (2012) derive a closed form expression of the CRPS for the GEV distribution and compare ML and minimum CRPS estimation for their analysis of peak wind speed. For our data set, ML

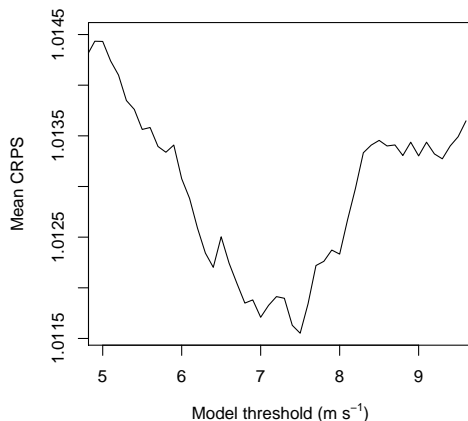


Figure 3: Mean continuous ranked probability score (CRPS) for the regime switching model as a function of the model threshold θ . The results are based on a rolling training period of 30 days during the out-of-sample time period from 1 February 2010 to 30 April 2010.

estimation proved to be more parsimonious and numerically stable. There is no analytical solution of the corresponding ML minimization problem (Coles, 2001). However, numerical approximations can be obtained using standard algorithms for any given dataset (Prescott and Walden, 1980). For the regime switching combination model in (3), minimum CRPS estimation is applied for the parameters of the TN distribution, and ML estimation for the parameters of the GEV distribution. For all three methods, the parameters are estimated over a rolling training period consisting of the forecast-observation pairs of the last m days. The parameters are estimated regionally in that training data from all stations are pooled together.

5 Results

Here, we present the results for 1-3 day ahead probabilistic forecasts of daily maximum wind speed over Germany produced by the three different postprocessing methods presented in Section 3. The verification period covers one year, from 1 May 2010 to 30 April 2011.

5.1 Selection of training period and regime switching threshold

The results presented here are based on a rolling training period of length $m = 30$ days for all methods. We have also performed the same analysis for training periods of length $m = 20, 25, \dots, 50$ days. In general, shorter training periods allow for a rapid adaption to changes in environmental conditions while longer training periods reduce the statistical variability in the parameter estimation (Gneiting et al., 2005). We found that the performance scores reported in Table 1 change by less than 1% for the different values of m , and in accordance with the results of Thorarinsdottir and Gneiting (2010) and Thorarinsdottir and Johnson (2012), we conclude that the methods are robust against changes in m .

The model threshold θ for the regime switching combination model in (3) is determined by computing the mean CRPS for a range of threshold values over an out-of-sample

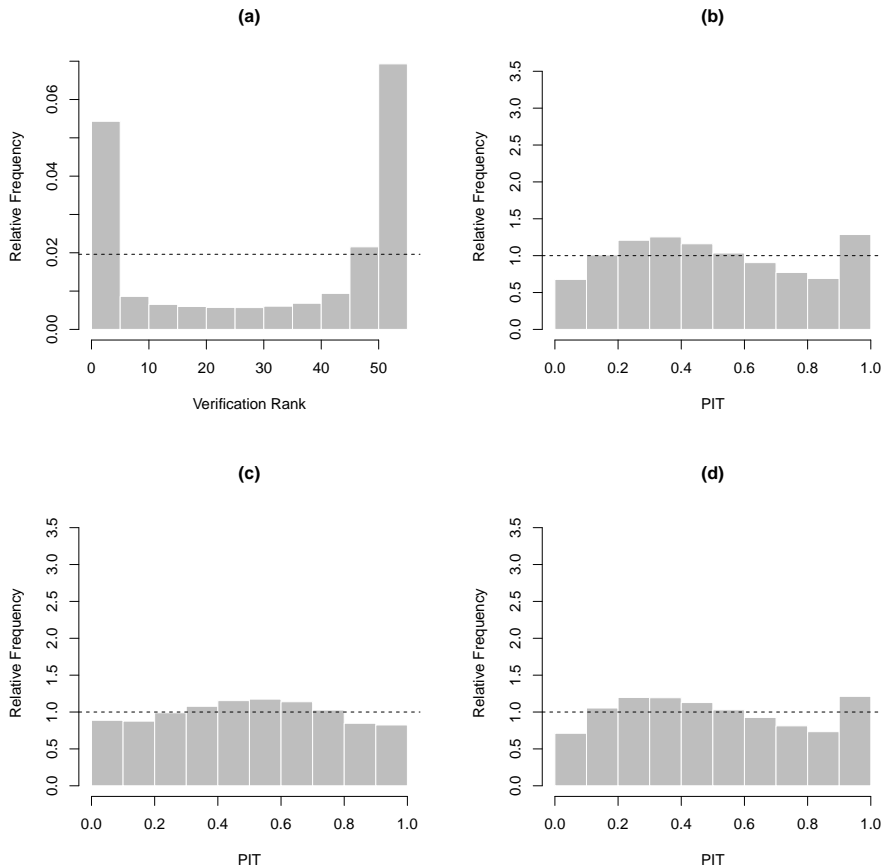


Figure 4: Calibration checks for probabilistic one day ahead forecasts of wind speed over Germany aggregated over 1 May 2010 to 30 April 2011 and the 228 stations: (a) verification rank histogram for the ECMWF ensemble forecasts; (b) PIT histogram for the TN model; (c) PIT histogram for the GEV model; (d) PIT histogram for the regime switching combination technique.

training period from 1 February 2010 to 30 April 2010. Using a rolling training period of $m = 30$ days, we obtain the optimal score for $\theta = 7.5 \text{ m s}^{-1}$, see Figure 3. A sensitivity analysis shows that the results in Table 1 are nearly constant for values of θ between 7 and 8 m s^{-1} and various choices of m . The threshold of 7.5 m s^{-1} corresponds approximately to the 80th percentile of the median ensemble predictions in the verification set. Over the verification period, a GEV distribution is used in around 18% of the forecast cases.

5.2 One day ahead predictive performance

We compare the three ensemble postprocessing methods discussed above to the raw, unprocessed ECMWF ensemble and a climatological reference forecast. For each day, the climatological reference forecast is obtained from the observed wind speeds in the 30 day training period used for the parameter estimation of the postprocessing methods. Verification rank and PIT histograms for the ensemble and the three postprocessing methods are shown in Figure 4. The ECMWF forecasts are underdispersive, with too many observations falling outside the ensemble range. All postprocessing methods significantly improve the calibration of the ensemble. While the GEV forecasts are slightly overdispersive and

Table 1: Mean continuous ranked probability score (CRPS), mean absolute error (MAE), average coverage and width of 80% prediction intervals of probabilistic one day ahead forecasts of daily maximum wind speed at 228 synoptic stations in Germany from 1 May 2010 to 30 April 2011.

Forecast	CRPS (m s^{-1})	MAE (m s^{-1})	Coverage (%)	Width (m s^{-1})
Climatology	1.54	2.13	64.4	6.6
Ensemble	1.26	1.44	26.6	1.0
TN	1.05	1.39	80.4	4.0
GEV	1.04	1.39	82.9	4.6
Combination	1.03	1.38	80.8	4.1

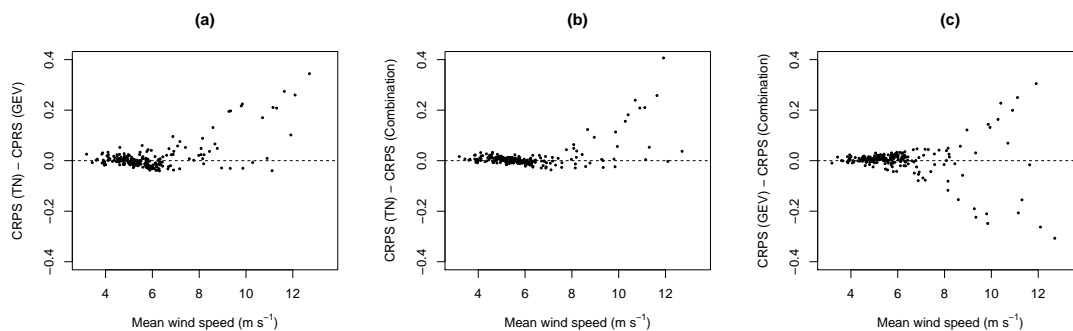


Figure 5: Station-specific comparisons of the continuous ranked probability score (CRPS) for the three postprocessing methods as a function of the average observed daily maximum wind speed at the station. The plots compare (a) the TN and the GEV model, (b) the TN and the regime switching combination model, and (c) the GEV and the regime switching combination model. The horizontal dashed lines indicate equal predictive performance.

show smaller deviations from uniformity than the TN forecasts, the PIT histogram of the combination model resembles the PIT histogram of the TN technique, with minor improvements for large PIT values.

Table 1 shows the mean CRPS, the MAE and average coverage and width of 80% prediction intervals for the competing forecasts. For calibrated forecasts, the average coverage of 80% prediction intervals should be close to 80% and narrower average prediction intervals indicate sharper forecasts. For discrete distributions such as the ECMWF ensemble and the climatology forecast, the CRPS can be calculated explicitly, see e.g. Berrocal et al. (2008). The CRPS for the TN model and the GEV model is calculated as described in Thorarinsdottir and Gneiting (2010) and Friederichs and Thorarinsdottir (2012), respectively. The ECMWF ensemble predictions outperform the climatological reference forecast and provide sharp prediction intervals at the cost of being uncalibrated. All post-processing methods outperform the ensemble predictions, with the GEV method showing small improvements in mean CRPS compared to the TN method. The regime switching combination method performs best in terms of both mean CRPS and MAE, slightly improving the results of the GEV method. Note that due to the heavier tails, the GEV model generally results in wider prediction intervals compared to the TN model.

Table 2: Mean threshold weighted continuous ranked probability score (twCRPS) for one day ahead forecasts of daily maximum wind speed at 228 synoptic stations in Germany from 1 May 2010 to 30 April 2011 using an indicator weight function $w_r(z) = \mathbb{1}\{z \geq r\}$ for different values of r .

Forecast	$r = 10$	$r = 12$	$r = 15$
Climatology	0.250	0.128	0.045
Ensemble	0.211	0.113	0.043
TN	0.200	0.111	0.042
GEV	0.195	0.107	0.041
Combination	0.191	0.103	0.039

Figure 5 compares the station-specific predictive performance of the individual postprocessing models as a function of the site-specific average observed wind speed. Figure 5(a) and (b) indicate that the overall improvements of the GEV and the regime switching combination model over the TN model are mainly due to improvements at stations with high average observed wind speeds. However, there appears to be no obvious pattern for stations with high average wind speeds when comparing the GEV model and the regime switching combination model in Figure 5(c). The combination model, on the other hand, outperforms the GEV model for most stations with average observed wind speeds below 7 m s^{-1} .

5.3 One day ahead performance in the upper tail

With a focus on the performance in the upper tail, Table 2 shows values of the mean twCRPS for the competing forecasts where we have employed the indicator weight function $w_r(z) = \mathbb{1}\{z \geq r\}$ for $r = 10, 12$ and 15 m s^{-1} . The threshold values approximately correspond to the 90th, 95th and 98th percentile of the marginal distribution of the wind speed observations. The twCRPS is here calculated using numerical integration methods. All three postprocessing methods improve the ECMWF ensemble predictions, and the GEV approach outperforms the TN method. The regime switching combination of the two models further improves the performance. Note that the relative improvement over the TN method for the upper tail is comparatively larger than the improvement under the unweighted CRPS in Table 1. Similar rankings hold for any value of r between 10 and 20 m s^{-1} .

We further consider the threshold-weighted continuous ranked probability skill score (twCRPSS) given by

$$\text{twCRPSS}(F, y) = 1 - \frac{\text{twCRPS}(F, y)}{\text{twCRPS}(F_{\text{ref}}, y)}, \quad (7)$$

where F_{ref} denotes the predictive cumulative distribution function of a reference forecast, in our case the TN method. The twCRPSS is positively oriented and can be interpreted as improvement over the reference forecast. Figure 6 shows the twCRPSS for the GEV and the regime switching combination method as a function of the threshold r for the indicator weight function, using the TN method as a reference forecast. For all thresholds and both models, the twCRPSS is strictly positive, indicating improved predictive performance compared to the TN model with the regime switching combination method showing greater improvement. In general, the score values increase for larger threshold values, with the largest differences obtained for threshold values around 14 m s^{-1} .




Figure 6: Threshold-weighted continuous ranked probability skill score of probabilistic one day ahead forecasts of daily maximum wind speed at 228 synoptic stations in Germany from 1 May 2010 to 30 April 2011 as a function of the threshold r in the indicator weight function $w_r(z) = \mathbb{1}\{z \geq r\}$, using the forecasts produced by the TN method as reference. The gray dashed vertical lines indicate the 50th, 90th, 95th and 99th percentile of the marginal distribution of the observations.

5.4 Performance for longer lead times

For lead times of two and three days, we obtain similar results as above. Table 3 shows the mean CRPS, MAE and average coverage and width of 80% prediction intervals for those lead times. Compared to the one day ahead forecasts, forecasts for longer lead times result in slightly less accurate predictions and wider prediction intervals. The ECMWF ensemble predictions exhibit wider prediction intervals compared to the one day ahead forecasts resulting in small improvements in calibration. However, the ensemble predictions are still underdispersive and the three postprocessing methods significantly improve the predictive skill of the ensemble. The differences among the three postprocessing models are less pronounced than for one day ahead forecasts. Results for the mean twCRPS and lead times of two and three days are shown in Table 4. We obtain the same ranking as for the one day ahead forecasts while the differences in the scores for the three models are small. Note that while the postprocessing methods show a slight decline in accuracy from two to three days ahead, the scores for the upper tail of the ensemble are identical for these lead times.

6 Discussion

We propose two extensions to the nonhomogeneous regression ensemble postprocessing approach of Thorarinsdottir and Gneiting (2010) employing generalized extreme value predictive distributions for daily maximum wind speed. In a case study over Germany using the exchangeable ECMWF ensemble, all three nonhomogeneous regression methods significantly improve the skill of the ECMWF ensemble and provide calibrated and accurate probabilistic forecasts. The best method according to our results is a regime

Table 3: Mean continuous ranked probability score (CRPS), mean absolute error (MAE), average coverage and width of 80% prediction intervals of daily maximum wind speed forecasts at 228 synoptic stations in Germany from 1 May 2010 to 30 April 2011. The upper half shows results for two day ahead forecasts, the lower half results for three day ahead forecasts.

Forecast	CRPS (m s ⁻¹)	MAE (m s ⁻¹)	Coverage (%)	Width (m s ⁻¹)
Two days ahead				
Climatology	1.55	2.14	64.4	6.6
Ensemble	1.22	1.47	38.8	1.6
TN	1.07	1.43	80.5	4.1
GEV	1.06	1.43	82.5	4.7
Combination	1.06	1.43	80.6	4.2
Three days ahead				
Climatology	1.55	2.14	64.4	6.6
Ensemble	1.22	1.52	48.0	2.1
TN	1.10	1.47	80.4	4.3
GEV	1.09	1.47	82.3	4.8
Combination	1.09	1.47	80.7	4.3

Table 4: Mean threshold weighted continuous ranked probability score (twCRPS) of daily maximum wind speed forecasts at 228 synoptic stations in Germany from 1 May 2010 to 30 April 2011 using an indicator weight function $w_r(z) = \mathbb{1}\{z \geq r\}$ for different values of r . The upper half shows results for two day ahead forecasts, the lower half shows results for three day ahead forecasts.

Forecast	$r = 10$	$r = 12$	$r = 15$
Two days ahead			
Climatology	0.250	0.128	0.045
Ensemble	0.209	0.113	0.043
TN	0.202	0.111	0.043
GEV	0.198	0.108	0.041
Combination	0.196	0.106	0.040
Three days ahead			
Climatology	0.250	0.128	0.045
Ensemble	0.209	0.113	0.043
TN	0.204	0.112	0.043
GEV	0.200	0.109	0.041
Combination	0.199	0.107	0.041

switching method, where a truncated normal model is applied when we expect low winds, while a GEV framework is used when high winds are expected.

In our GEV approach, we have not accounted for the possibility of the method predicting negative wind speeds. While this rarely happened in our case study, different results are likely to be obtained for less extreme wind variables. In an application to quantitative precipitation, Scheuerer (2012) considers the GEV to be left-censored at zero assigning all mass below zero to exactly zero. This approach seems very appropriate for precipitation, where there is often high probability of zero precipitation. However, it seems less appropriate for wind variables. Instead, one might consider a truncation of the GEV distribution similar to the truncated normal distribution in (1).

The regime-switching combination of the TN and the GEV model offers several starting points for further extensions and potential improvements. For each forecast case, either the TN or the GEV model is selected based on the median ensemble prediction falling below or exceeding a fixed threshold θ . Instead of assuming a fixed threshold value, it might be interesting to develop an adaptive method to automatically estimate θ , for example based on station-specific information, or covariate based on other weather variables during the rolling training period. Alternatively, improvements of the predictive performance for extreme events might be achieved by considering a mixture model using a truncated normal distribution for the bulk of the distribution, and an adaptive generalized Pareto distribution (GPD) for the tail. Bentzien and Friederichs (2012) propose such a mixture model for precipitation using lognormal and gamma mixtures for the bulk of the distribution and an adaptive GPD tail which is able to significantly improve the predictive performance for extreme quantiles. Similarly, all three methods could be extended by allowing for local adaption of the parameter estimation. This might be obtained with a geostatistical approach analogous to Kleiber et al. (2011b,a), where the spatially varying parameters are estimated locally and interpolated to locations without available observations.

When assessing the predictive performance in the upper tail, we have focused on the twCRPS with a simple indicator weight functions. Under this weight function, it is not possible to distinguish between forecasts with the same behavior on $[r, \infty)$, but different behavior on $(-\infty, r)$. Following Gneiting and Ranjan (2011), we have also performed our analysis using weight functions of the form $w_r(z) = \Phi(z|\mu_r, \sigma_r^2)$, where $\Phi(\cdot|\mu_r, \sigma_r^2)$ denotes the cumulative distribution function of a normal distribution with mean μ_r and variance σ_r^2 . For example, we might set $\mu_r = r$ and $\sigma_r^2 = 1$, or $\sigma_r^2 = S^2$, where S^2 is the sample variance of the wind speed observations during the out-of-sample training period from 1 February to 30 April 2010. The average values of the twCRPS using these weight functions result in the same ranking and in comparable relative improvements as reported above for the indicator weight function. We have therefore focused on the indicator weight function which is computationally less demanding.

7 Acknowledgments

We thank Tilmann Gneiting and Michael Scheuerer for helpful discussions and for providing the observation data. The support of the Volkswagen Foundation through the project “Mesoscale Weather Extremes–Theory, Spatial Modeling and Prediction (WEX-MOP)” and Statistics for Innovation *sfi*² in Oslo is gratefully acknowledged.

References

- J. L. Anderson. A method for producing and evaluating probabilistic forecasts from ensemble model integrations. *J. Climate*, 9:1518–1530, 1996.
- L. Bao, T. Gneiting, E. P. Grit, P. Guttorp, and A. E. Raftery. Bias correction and Bayesian model averaging for ensemble forecasts of surface wind direction. *Mon. Weather Rev.*, 138:1811–1821, 2010.
- S. Bentzien and P. Friederichs. Generating and Calibrating Probabilistic Quantitative Precipitation Forecasts from the High-Resolution NWP Model COSMO-DE. *Weather Forecast.*, 27:988–1002, 2012.
- V. J. Berrocal, T. Gneiting, and A. E. Raftery. Probabilistic quantitative precipitation field forecasting using a two-stage spatial model. *Ann. Appl. Stat.*, 2:1170–1193, 2008.
- G. W. Brier. Verification of forecasts expressed in terms of probability. *Mon. Weather Rev.*, 78:1–3, 1950.
- J. Bröcker and L. A. Smith. Scoring Probabilistic Forecasts: The Importance of Being Proper. *Weather Forecast.*, 22:382–388, 2007.
- A. N. Celik. A statistical analysis of wind power density based on the Weibull and Rayleigh models at the southern region of Turkey. *Renew. Energ.*, 29:93–604, 2004.
- S. Coles. *An Introduction to Statistical Modeling of Extreme Values*. Springer, London, 2001.
- J. F. Courtney, P. Lynch, and C. Sweeney. High resolution forecasting for wind energy applications using Bayesian model averaging. *Tellus A*, 65:19669, 2013.
- A. P. Dawid. Statistical theory: The prequential approach. *J. Roy. Stat. Soc. A*, 147:278–292, 1984.
- C. Diks, V. Panchenko, and D. van Dijk. Likelihood-based scoring rules for comparing density forecasts in tails. *J. Econometrics*, 163:215–230, 2011.
- European Wind Energy Association. Wind in power: 2012 European statistics, 2012. Available online at http://www.ewea.org/fileadmin/files/library/publications/statistics/Wind_in_power_annual_statistics_2012.pdf.
- C. Fraley, A. E. Raftery, and T. Gneiting. Calibrating multimodel forecast ensembles with exchangeable and missing members using Bayesian model averaging. *Mon. Weather Rev.*, 138:190–202, 2010.
- P. Friederichs and T. L. Thorarinsdottir. Forecast verification for extreme value distributions with an application to probabilistic peak wind prediction. *Environmetrics*, 7:579–594, 2012.
- A. Garcia, J. L. Torres, E. Prieto, and A. De Francisco. Fitting wind speed distributions: A case study. *Sol. Energy*, 62:139–144, 1998.
- M. Genton and A. Hering. Blowing in the wind. *Significance*, 4:11–14, 2007.
- T. Gneiting. Making and evaluating point forecasts. *J. Am. Stat. Assoc.*, 106:746–762, 2011.

- T. Gneiting and A. E. Raftery. Weather forecasting with ensemble methods. *Science*, 310:248–249, 2005.
- T. Gneiting and A. E. Raftery. Strictly proper scoring rules, prediction, and estimation. *J. Am. Stat. Assoc.*, 102:359–378, 2007.
- T. Gneiting and R. Ranjan. Comparing density forecasts using threshold-and quantile-weighted scoring rules. *J. Bus. Econ. Stat.*, 29:411–422, 2011.
- T. Gneiting, A. E. Raftery, A. H. Westveld, and T. Goldman. Calibrated probabilistic forecasting using ensemble model output statistics and minimum CRPS estimation. *Mon. Weather Rev.*, 133:1098–1118, 2005.
- T. Gneiting, F. Balabdaoui, and A. E. Raftery. Probabilistic forecasts, calibration and sharpness. *J. Roy. Stat. Soc. B*, 69:243–268, 2007.
- I. J. Good. Rational decisions. *J. Roy. Stat. Soc. B*, 14:107–114, 1952.
- R. Hagedorn, T. M. Hamill, and J. S. Whitaker. Probabilistic forecast calibration using ECMWF and GFS ensemble reforecasts. Part I: Two-meter temperatures. 136:2608–2619, 2008.
- T. M. Hamill and S. J. Colucci. Verification of Eta-RSM short-range ensemble forecasts. *Mon. Weather Rev.*, 125:1312–1327, 1997.
- H. Hersbach. Decomposition of the continuous ranked probability score for ensemble prediction systems. *Weather Forecast.*, 15:559–570, 2000.
- A. Kann, C. Wittmann, Y. Wang, and X. Ma. Calibrating 2-m temperature of limited-area ensemble forecasts using high-resolution analysis. *Mon. Weather Rev.*, 137:3373–3387, 2009.
- W. Kleiber, A. E. Raftery, J. Baars, T. Gneiting, C. F. Mass, and E. P. Gritmit. Locally calibrated probabilistic temperature forecasting using geostatistical model averaging and local Bayesian model averaging. *Mon. Weather Rev.*, 139:2630–2649, 2011a.
- W. Kleiber, A. E. Raftery, and T. Gneiting. Geostatistical model averaging for locally calibrated probabilistic quantitative precipitation forecasting. *J. Amer. Stat. Assoc.*, 106:1291–1303, 2011b.
- M. Lei, L. Shiyang, J. Chuanwen, L. Hongling, and Z. Yan. A review on the forecasting of wind speed and generated power. *Renew. Sust. Energ. Rev.*, 13:915–920, 2009.
- M. Leutbecher and T. N. Palmer. Ensemble forecasting. *J. Comput. Phys.*, 227:3515–3539, 2008.
- F. Molteni, R. Buizza, and T. N. Palmer. The ECMWF Ensemble Prediction System: Methodology and validation. *Q. J. Roy. Meteor. Soc.*, 122:73–119, 1996.
- T. N. Palmer. The Economic Value of Ensemble Forecasts as a Tool for Risk Assessment: From Days to Decades. *Q. J. Roy. Meteor. Soc.*, 128:747–774, 2002.
- J. P. Palutikof, B. B. Brabson, D. H. Lister, and S. T. Adcock. A review of methods to calculate extreme wind speeds. *Meteorol. Appl.*, 6:119–132, 1999.

- P. Pinson. Adaptive calibration of (u, v) -wind ensemble forecasts. *Q. J. Roy. Meteor. Soc.*, 138:1273–1284, 2012.
- P. Pinson and R. Hagedorn. Verification of the ECMWF ensemble forecasts of wind speed against analyses and observations. *Meteorol. Appl.*, 19:484–500, 2012.
- P. Pinson, C. Chevallier, and G. N. Kariniotakis. Trading wind generation from short-term probabilistic forecasts of wind power. *IEEE T. Power Syst.*, 22:1148–1156, 2007.
- P. Prescott and A. T. Walden. Maximum likelihood estimation of the parameters of the generalized extreme-value distribution. *Biometrika*, 67:723–724, 1980.
- M. Scheuerer. Probabilistic quantitative precipitation forecasting using ensemble model output statistics, 2012. arXiv preprint arXiv:1302.0893.
- N. Schuhen, T. L. Thorarinsdottir, and T. Gneiting. Ensemble model output statistics for wind vectors. *Mon. Weather Rev.*, 140:3204–3219, 2012.
- J. M. Sloughter, T. Gneiting, and A. E. Raftery. Probabilistic wind speed forecasting using ensembles and bayesian model averaging. *J. Amer. Stat. Assoc.*, 105:25–35, 2010.
- J. M. Sloughter, T. Gneiting, and A. E. Raftery. Probabilistic wind vector forecasting using ensembles and Bayesian model averaging. *Mon. Weather Rev.*, in press, 2013.
- T. L. Thorarinsdottir and T. Gneiting. Probabilistic forecasts of wind speed: Ensemble model output statistics by using heteroscedastic censored regression. *J. Roy. Stat. Soc. A*, 173:371–388, 2010.
- T. L. Thorarinsdottir and M. S. Johnson. Probabilistic wind gust forecasting using non-homogeneous Gaussian regression. *Mon. Weather Rev.*, 140:889–897, 2012.