

STOCHASTIC SIMULATION OF PREDICTIVE SPACE–TIME SCENARIOS OF WIND SPEED USING OBSERVATIONS AND PHYSICAL MODEL OUTPUTS¹

BY JULIE BESSAC*, EMIL CONSTANTINESCU*,[†] AND MIHAI ANITESCU*,[†]

Argonne National Laboratory and the University of Chicago[†]*

We propose a statistical space–time model for predicting atmospheric wind speed based on deterministic numerical weather predictions and historical measurements. We consider a Gaussian multivariate space–time framework that combines multiple sources of past physical model outputs and measurements in order to produce a probabilistic wind speed forecast within the prediction window. We illustrate this strategy on wind speed forecasts during several months in 2012 for a region near the Great Lakes in the United States. The results show that the prediction is improved in the mean-squared sense relative to the numerical forecasts as well as in probabilistic scores. Moreover, the samples are shown to produce realistic wind scenarios based on sample spectra and space–time correlation structure.

1. Introduction. In this study, we propose a statistical space–time model for predicting atmospheric wind speed based on numerical weather predictions and historical measurements. The wind speed predictions are based on deterministic numerical weather prediction (NWP) model outputs in a framework that integrates past dependence between observational measurements and the NWP model outputs. The past dependence between these two datasets is modeled linearly into a Gaussian process (GP) hierarchical framework. The aim of this work is to improve the wind speed forecasts and to produce samples (referred to as scenarios) from the predictive distribution of wind speeds. This is achieved by using the GP framework in conjunction with NWP forecast output.

Atmospheric near-surface wind conditions are important for numerous sectors of human activities, and the topic has received considerable attention in the past several years, for instance, in the study of crop models [Brisson et al. (2003)], object drift into the ocean [Ailliot, Frénod and Monbet (2006)], and severe weather forecasting [Thorarinsdottir and Johnson (2012)]. Arguably, one of the largest applications is in wind energy, because energy management with renewables relies heavily on an accurate description of the forecast uncertainty [Constantinescu et al. (2011), Li et al. (2015), Papavasiliou, Oren and Rountree (2015), Pinson (2013), Pinson et al. (2009)].

Received October 2016; revised September 2017.

¹Supported by the U.S. Department of Energy Office of Science through contract no. DE-AC02-06CH11357.

Key words and phrases. Hierarchical Gaussian model, multiple data sources, predictive scenarios, spatio-temporal wind speed.

1.1. *Existing context on statistical wind speed prediction.* Several components of the wind field can be predicted separately or jointly: the zonal and meridional components [Hering and Genton (2010), Sloughter, Gneiting and Raftery (2013)], wind speed [Brown, Katz and Murphy (1984), Gneiting et al. (2006), Sloughter, Gneiting and Raftery (2010)], and wind direction [Bao et al. (2010)]. Prediction of wind conditions can be generated by statistical models built for predicting observed wind [Brown, Katz and Murphy (1984), Gneiting et al. (2006), Hering and Genton (2010)], or they can be generated from the statistical post-processing of NWP model forecasts; these latter fall into the domain of model output statistics (MOS) methods [Glahn and Lowry (1972), Gneiting et al. (2005), Raftery et al. (2005)]. From a statistical point of view, the prediction error can be accounted for through the use of a predictive distribution. Ensemble forecasts aim at assessing the uncertainty associated with the numerical model; however, this strategy is known to be often uncalibrated and underdispersive [Gneiting et al. (2005)]. More recently, the generation of predictive scenarios has gained significant momentum. These scenarios enable accounting for the uncertainty of the forecasts for various locations and/or time-ahead lags [Pinson and Girard (2012), Pinson et al. (2008)]. Predictive scenarios fall into the class of multivariate probabilistic distributions, which can be handled, for instance, with copula-based methods [Clark et al. (2004), Pinson (2013), Wilks (2015)] or with parametric methods [Hering and Genton (2010), Kazor and Hering (2015)].

In the context of improving numerical forecasts, MOS methods provide probabilistic forecasts by post-processing the single or ensemble forecasts and tend to address the issue of bias and dispersion. Introduced at first for single-trajectory forecast [Glahn and Lowry (1972)], MOS methods have been extended to ensemble model outputs statistics (EMOS) post-processing methods for ensemble forecasts [Gneiting et al. (2005)]. Most of the MOS methods for ensembles are variations and extensions of the Bayesian model averaging (BMA) initially proposed in Raftery et al. (2005) and the nonhomogeneous regression (NGR) model proposed in Gneiting et al. (2005). Several variants of both models have been proposed with different distributions [Baran (2014), Baran and Lerch (2015), Lerch and Thorarinsdottir (2013), Sloughter, Gneiting and Raftery (2010)], a BMA extension to spatial interpolation [see Scheuerer and Möller (2015) and references therein], and an NGR with regime-switching [Baran and Lerch (2016)]. Lately, multivariate models have been introduced, namely, spatial models [Gel, Raftery and Gneiting (2004), Thorarinsdottir and Gneiting (2010)] and bivariate frameworks [Schuhen, Thorarinsdottir and Gneiting (2012), Sloughter, Gneiting and Raftery (2010)]. In Schefzik, Thorarinsdottir and Gneiting (2013), a multivariate tool based on the use of copulas is presented that allows one to account for time-ahead dependence, spatial dependence, and dependence among variables; see also Wilks (2015) for a discussion of multivariate MOS methods using copulas.

MOS and regression methods treat the NWP model outputs as covariates. Following the arguments in Schefzik, Thorarinsdottir and Gneiting (2013), these models neglect the dependence between variables, between the different time-ahead

lags, and between spatial locations. Notable exceptions are the use of copulas to model the variable dependence [Schefzik, Thorarinsdottir and Gneiting (2013)] and the use of a parametric model for spatial locations [Feldmann, Scheuerer and Thorarinsdottir (2015)]. Moreover, the use of covariates creates difficulties such as addressing misalignment in space and/or time with the response variables and possible nonexhaustive sampling of these covariates. Multivariate modeling alleviates these problems by jointly modeling several variables as a random process. Multivariate space–time modeling has been an area of intense research in the past two decades; see Fanshawe and Diggle (2012) for a review of bivariate geostatistical modeling, and see Berliner (2000), for a discussion of hierarchical Bayesian modeling for multiple dependent datasets.

In the multivariate modeling context, various statistical approaches have been proposed for hybrid NWP—physical observations usage. For example, in Fuentes et al. (2005) a Bayesian hierarchical model is presented that combines NWP outputs and observed measurements to provide spatial prediction for chemical species. A hidden process is used to represent the unobserved “true” concentration of sulfur dioxide, and the sources of data are affine transformations of this “true” process. A similar approach was used in a space–time context for multiple measurements of snow water equivalent data in Cowles et al. (2002). Several outputs of regional climate models are combined in a spatial framework by using a hierarchical model based on a spatial random effects model in Kang, Cressie and Sain (2012). Berrocal, Gelfand and Holland (2012) proposed a space–time hierarchical Bayesian model to fuse measurements and numerical model outputs of air-quality data, with an extension of a downscaling model introduced some years ago. Royle and Berliner (1999) introduced a conditional hierarchical model that combines two heterogeneous spatial datasets of ozone and temperature in the objective to predict these two variables on gridded points while they are recorded at two different irregular networks. A joint distribution is used to cope with the nonalignment of the two datasets and of the predicted values at gridded points. This distribution is specified in a hierarchical conditional way in order to avoid the direct specification of the joint distribution. Indeed, the modeling of multivariate covariance structure is challenging and is still an on-going research area; see Apanasovich and Genton (2010), Bourotte, Allard and Porcu (2016), Genton and Kleiber (2015).

1.2. *Proposed modeling framework and position within the existing literature.*

In this study, we fuse two heterogeneous datasets: numerical forecasts and physical observations. Our endeavor stems from the observation that physical observations alone cannot be used to deliver accurate forecasts 24–48 hours ahead, whereas NWP uncertainty analysis may be (as pointed out above) uncalibrated and underdispersive. By heterogeneity, we mean that the physical observations are not necessarily on the NWP output grid. We note that refining the grid would be a costly computational expense and would not guarantee that the sites of the physical

observations were exactly on the grid or that the forecast were improved. Moreover, NWP physics constrains assumptions that make refining below a grid size limit inconsistent, if not even erroneous [Palmer (2014)]. We propose a bivariate space–time Gaussian process to improve forecasts from an NWP model, where the physics model outputs and the measurements are modeled as two random space–time processes. GPs allow us to describe the conditional forecast distribution as a Gaussian distribution as well as facilitate robust computational algorithms [Anitescu, Chen and Wang (2012), Stein, Chen and Anitescu (2012)]. Moreover, we use a marginal Box–Cox transformation to address the typical positivity and skewness of wind speed distribution. Our model is specified in a hierarchical way in order to avoid characterizing the full space–time bivariate covariance. We extend the specification, which was initially proposed in Royle and Berliner (1999), Royle et al. (1999) in a spatial context, to space–time modeling.

The numerical forecasts of wind speed are combined with historical measurements data to provide scenarios of prediction. The modeling of the NWP forecasts as a random process with spatial or space–time structure, also carried out in Berrocal, Gelfand and Holland (2012), Fuentes et al. (2005), Royle et al. (1999), differentiates this approach from MOS methods, where dependencies are most of the time ignored. The NWP random process modeling allows for a consistent statistical framework for the discrepancy between the NWP output and physical measurement locations. A particularly important aspect of our model is that the proposed prediction framework accounts for the space–time dependence between these two heterogeneous datasets. Furthermore, in this work we consider a strategy where NWP future and past temporal information is used to provide a temporal prediction that begins at the current time. In our approach, we set a 24-hour ahead forecast window and use this entire window of NWP outputs to predict wind speed at each time hour during this window. In contrast, MOS methods commonly work with current information only. In other words, MOS methods adjust the forecast point-wise in space or time sequentially one step at a time, whereas our approach accounts for future and past information within the forecast window itself. The same argument applies to the strategies proposed in Berrocal, Gelfand and Holland (2012), Fuentes and Raftery (2005).

The paper is organized as follows. In Section 2, we introduce the modeling context and the statistical formulation. In Section 3, we describe the two sources of data that are used and combined. In Section 4, the model is validated on different months of the year, and the quality of the time prediction at one station is assessed. We highlight the improvements in terms of the forecasting accuracy of the proposed model with respect to the NWP forecasts. We conclude in Section 5 by presenting general improvements made by the model with respect to the NWP data.

2. Statistical model for NWP outputs. In this section, we introduce a Gaussian modeling framework that embeds the space–time dependence between measured observations and NWP model forecasts. We model two heterogeneous

spatio-temporal datasets as jointly distributed variables. We extend the hierarchical GP ideas of [Royle and Berliner \(1999\)](#) to a space–time context; the joint process is a space–time Gaussian process specified conditionally. We provide temporal prediction of observations given past and future NWP data while accounting for past space–time dependence between the two datasets; see equation (2.3) below.

2.1. *Overview of the proposed method.* We consider measured observations, Y_{Obs} , and NWP forecasts, Y_{NWP} . Both are available in the past; however, only future Y_{NWP} is available at the current time. We are therefore interested in generating samples from the distribution of future observations conditional on the past observations and NWP simulations and on future NWP simulations. This distribution is represented as a hierarchical Gaussian process and is calibrated by maximizing its likelihood. In particular, the ingredients of our proposed approach are as follows:

- We aim to construct a probabilistic model for future observations based on the current available data:

$$(2.1) \quad p(y_{\text{Obs}}^u(t_{\text{future}})|\theta, y_{\text{NWP}}^a(t_{\text{future}}), y_{\text{Obs}}^a(t_{\text{past}}), y_{\text{NWP}}^a(t_{\text{past}})),$$

where the superscript “ a ” stands for available and “ u ” for unavailable quantities and θ denotes parameters of a statistical model (see Section 2.2).

- To specify the probability in (2.1), we use a joint model for $(Y_{\text{Obs}}, Y_{\text{NWP}})$ represented by a GP. As indicated in Section 2.3, the full joint distribution can be expressed as

$$(2.2) \quad \begin{pmatrix} Y_{\text{Obs}} \\ Y_{\text{NWP}} \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} \mu_{\text{Obs}}(\theta) \\ \mu_{\text{NWP}}(\theta) \end{pmatrix}, \begin{pmatrix} \Sigma_{\text{Obs}}(\theta) & \Sigma_{\text{Obs,NWP}}(\theta) \\ \Sigma_{\text{Obs,NWP}}^T(\theta) & \Sigma_{\text{NWP}}(\theta) \end{pmatrix} \right),$$

where θ is the set of parameters that describe the parametric shapes of the means and covariances. One of the main objectives is to specify an appropriate structure for the mean and covariance of (2.2).

- The joint distribution of $(Y_{\text{Obs}}, Y_{\text{NWP}})$ in (2.2) is specified in a conditional hierarchical way where the distributions of $(Y_{\text{Obs}}|Y_{\text{NWP}})$ and (Y_{NWP}) are described separately in terms of parametric space–time mean and covariance. This is discussed in detail in Section 2.4.
- The parameters θ in (2.2) are estimated by maximizing the likelihood, as discussed in Section 2.5, so that $\theta^* = \arg \max_{\theta} \mathcal{L}(\theta; y_{\text{Obs}}^a(t_{\text{past}}), y_{\text{NWP}}^a(t_{\text{past}}))$. This aspect is also discussed at the end of Section 2.2 in the more abstract context of equation (2.1), where we also specify the probabilistic modeling assumptions.
- The predictive scenarios sampled from $p(y_{\text{Obs}}^u(t_{\text{future}})|\dots)$ are obtained via kriging equations and are detailed in Section 2.6.

In other words, we represent the joint distribution of the numerical simulations and observations through a Gaussian model that is calibrated by maximizing the likelihood of the model parameters. This distribution is then used to condition on the numerical forecasts at the current time in order to forecast observations. In the following sections, we give details of each step in constructing this framework.

2.2. *Prediction setting.* Let us assume that both measured observations Y_{Obs} and NWP forecasts Y_{NWP} are available from time t_1 to time t_{k_K} . In the following, the term “observations” refers to the observational measurements. Observations are available at J_0 locations $S = \{s_1, \dots, s_{J_0}\}$, and NWP forecasts are available over a grid that covers these stations. Each day, the Weather Research and Forecasting (WRF) model is run for a period of h hours independently from the previous day, because WRF is initialized from a reanalysis or assimilated dataset; time can then be written in terms of blocks of length h . Henceforth, we consider a time window of $h = 24$ hours. We denote by b_i the i th time block of length h , $b_i = \{t_{k_i}, \dots, t_{k_i+h-1}\}$. The objective here is to predict the measurements Y_{Obs} between time $t_{k_{K+1}}$ and $t_{k_{K+1}+h-1}$ at stations $S = \{s_1, \dots, s_{J_0}\}$, and possibly at locations $S_0 = \{s_{J_0+1}, \dots, s_J\}$ where no historical measurements are recorded, from NWP forecasts that are available between $t_{k_{K+1}}$ and $t_{k_{K+1}+h-1}$. S denotes the sets of measurements stations, Y_{NWP} are taken at the closest grid point. S_0 denotes a set of unobserved stations. This can be summarized by

$$(2.3) \quad \left(\begin{matrix} y_{\text{Obs}}^a(b_{1:K}; S) \\ y_{\text{Obs}}^u(b_{K+1}; S, S_0) \end{matrix} \right) \quad \text{and} \quad \left(\begin{matrix} y_{\text{NWP}}^a(b_{1:K}; S) \\ y_{\text{NWP}}^a(b_{K+1}; S, S_0) \end{matrix} \right),$$

with $b_{1:K} = \{b_1, b_2, \dots, b_K\}$.

In this context, the model is trained on the following available pairs:

$$(y_{\text{Obs}}^a(b_{1:K}; S), y_{\text{NWP}}^a(b_{1:K}; S)),$$

and the prediction is made from $y_{\text{NWP}}^a(b_{K+1}; S, S_0)$ to estimate $y_{\text{Obs}}^u(b_{K+1}; S, S_0)$, where $b_{K+1} = \{t_{k_{K+1}}, \dots, t_{k_{K+1}+h-1}\}$. In a probabilistic sense, we aim to compute

$$(2.4) \quad p(y_{\text{Obs}}^u(b_{K+1}) | y_{\text{NWP}}^a(b_{K+1}), y_{\text{Obs}}^a(b_{1:K}), y_{\text{NWP}}^a(b_{1:K}))$$

$$= \int p(y_{\text{Obs}}^u(b_{K+1}), \theta | y_{\text{NWP}}^a(b_{K+1}), y_{\text{Obs}}^a(b_{1:K}), y_{\text{NWP}}^a(b_{1:K})) d\theta,$$

where θ is a random set of model parameters, blocks $b_{1:K}$ are available, and b_{K+1} is a predicted block; the spatial components are suppressed for brevity. Note that b_{K+1} is not necessarily a block coming right after b_K , but rather a day that is not observed. To simplify the computation of (2.4), we now make several assumptions. First, we assume that we have approximate independence of $y_{\text{Obs}}^u(b_{K+1})$ and $y_{\text{Obs}}^a(b_{1:K}), y_{\text{NWP}}^a(b_{1:K})$ conditional on $y_{\text{NWP}}^a(b_{K+1})$. In hierarchical models such as ours, which has NWP predictions as its first layer and the observation sites as the second layer, one commonly assumes that random variables on the second layer are independent conditional on the realizations of the ones in the first layer; see [Cressie and Wikle \(2011\)](#). We assume as well that θ^* can be obtained by maximizing the likelihood

$$(2.5) \quad \theta^* = \underset{\theta}{\text{argmax}} \mathcal{L}(\theta; y_{\text{Obs}}^a(b_{1:K}), y_{\text{NWP}}^a(b_{1:K}))$$

$$= \underset{\theta}{\text{argmax}} p(\theta | y_{\text{Obs}}^a(b_{1:K}), y_{\text{NWP}}^a(b_{1:K})).$$

The probability (2.4) that we aim to compute is then approximated by the following probability obtained by conditioning only on the NWP predictions for the same temporal block and plugging-in the maximum likelihood estimate θ^* :

$$(2.6) \quad \begin{aligned} p(y_{\text{Obs}}^u(b_{K+1}) | y_{\text{NWP}}^a(b_{K+1}), y_{\text{Obs}}^a(b_{1:K}), y_{\text{NWP}}^a(b_{1:K})) \\ \approx p(y_{\text{Obs}}^u(b_{K+1}) | \theta^*, y_{\text{NWP}}^a(b_{K+1})). \end{aligned}$$

In what follows, we consider multivariate normal distributions to describe the involved probabilities of (2.6). We have found that a sensible approach is to model statistically the output of NWP itself. Indeed, a probabilistic structure for NWP outputs enables us to address misalignment issues between the two datasets and express a kriging distribution (2.6) in that case. Moreover, modeling the joint distribution allows us to embed statistically the space–time structure of the NWP outputs. Efforts will be directed to specify the full multivariate distribution for the two involved datasets, namely, the mean and covariance structures. In Section 2.3, we review a hierarchical approach for Gaussian processes, and in Section 2.4 we present the model used for the mean and covariance functions that introduce the parametrization θ .

2.3. Hierarchical bivariate framework. Gaussian processes are chosen for their convenience in expressing conditional distributions and in a multivariate and space–time context. Power transformations are commonly used to approximate Gaussian margins. To address the typical skewness of the wind speed distribution, we apply the Box–Cox transformation; see, for instance, [Brown, Katz and Murphy (1984)]. A specific transformation is used for each dataset (NWP and measurements) to account for the heterogeneity between the two datasets. Within each dataset, the same power transformation is applied to each spatial location in order to preserve the variance structure. The model is fitted on the transformed data. We write the joint distribution of the process $(Y_{\text{Obs}}, Y_{\text{NWP}})$ as

$$(2.7) \quad \begin{pmatrix} Y_{\text{Obs}} \\ Y_{\text{NWP}} \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} \mu_{\text{Obs}} \\ \mu_{\text{NWP}} \end{pmatrix}, \begin{pmatrix} \Sigma_{\text{Obs}} & \Sigma_{\text{Obs}, \text{NWP}} \\ \Sigma_{\text{Obs}, \text{NWP}}^T & \Sigma_{\text{NWP}} \end{pmatrix} \right).$$

The goal here is to find accurate representations of the data with this distribution in terms of mean and variance. The positive-definiteness of block matrices is generally difficult to ensure when specifying the three blocks in (2.7) independently. Therefore, to avoid specifying of the full covariance in (2.7), we follow the hierarchical conditional modeling proposed by [Royle and Berliner (1999), Royle et al. (1999)], and we model $(Y_{\text{Obs}} | Y_{\text{NWP}})$ and (Y_{NWP}) , where $(Y_{\text{Obs}} | Y_{\text{NWP}})$ stands for the conditional distribution of Y_{Obs} given Y_{NWP} . When $(Y_{\text{Obs}}, Y_{\text{NWP}})$ is a Gaussian process, $(Y_{\text{Obs}} | Y_{\text{NWP}})$ and (Y_{NWP}) follow a Gaussian distribution; then only first- and second-order structures are to be specified. Consequently, the model is described by the following distributions:

$$(2.8) \quad (Y_{\text{Obs}} | Y_{\text{NWP}}) \sim \mathcal{N}(\mu_{\text{Obs} | \text{NWP}}, \Sigma_{\text{Obs} | \text{NWP}}).$$

The Gaussian joint distribution of (Y_{Obs}, Y_{NWP}) implies a conditional linear dependence between Y_{Obs} and Y_{NWP} , which agrees reasonably with the data analysis:

$$(2.9) \quad \mu_{Obs|NWP} = E(Y_{Obs}|Y_{NWP}) = \mu + \Lambda Y_{NWP},$$

where μ and Λ , which is called the transition matrix, will be parameterized below, and

$$(2.10) \quad Y_{NWP} \sim \mathcal{N}(\mu_{NWP}, \Sigma_{NWP}).$$

From these equations, we express the full joint distribution given by (2.7) as

$$(2.11) \quad \begin{pmatrix} Y_{Obs} \\ Y_{NWP} \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} \mu + \Lambda \mu_{NWP} \\ \mu_{NWP} \end{pmatrix}, \begin{pmatrix} \Sigma_{Obs|NWP} + \Lambda \Sigma_{NWP} \Lambda^T & \Lambda \Sigma_{NWP} \\ (\Lambda \Sigma_{NWP})^T & \Sigma_{NWP} \end{pmatrix} \right).$$

The derivation of the latter joint mean and covariance do not necessitate the assumption of normality, the covariance can be derived from the law of total variance. The normality of the joint distribution is derived from the Bayes' rule.

2.4. *Statistical model.* To provide time prediction and to ensure model parsimony, we propose a parameterization in space and time of the involved first- and second-order structures of the conditional and the marginal distributions defined by (2.8) and (2.10) are specified following an exploratory analysis of the datasets.

2.4.1. *Marginal mean structure of (Y_{NWP}) .* The empirical mean function of Y_{NWP} exhibits spatial patterns associated with the geographical coordinates but also with several parameters of the NWP model. Indeed the studied area is in the Great Lakes region with the large water mass of Lake Michigan (the "land use," LU, is used, which is a categorical variable that represents the type of land used in the parameterization of the NWP model). Time-periodic effects are present in the first-order structure of Y_{NWP} and are accounted for through harmonics of different frequencies. In Figure 2, these spatial and temporal patterns are plotted. We write

$$(2.12) \quad \begin{aligned} \mu_{NWP}(t, s) &= E(Y_{NWP}(t, s)) \\ &= \left(\beta_0 + \beta_1 \cos\left(\frac{2\pi t}{24}\right) + \beta_2 \sin\left(\frac{2\pi t}{24}\right) + \beta_3 \cos\left(\frac{2\pi t}{12}\right) \right. \\ &\quad \left. + \beta_4 \sin\left(\frac{2\pi t}{12}\right) + \beta_5 \cos\left(\frac{2\pi t}{8}\right) + \beta_6 \sin\left(\frac{2\pi t}{8}\right) \right) \\ &\quad \times (\alpha_0(LU(s)) + \alpha_1 Lat(s) + \alpha_2 Long(s)), \end{aligned}$$

where t is measured in hours, $LU(s)$ is a categorical variable that represents the land use associated with station s used in the model; Lat and $Long$ are the latitude and longitude coordinates; $(\alpha_0(l))_{l=1, \dots, n}$, with n the number of possible land uses, α_1 , α_2 , and $(\beta_k)_{k=0, \dots, 6}$ are real numbers to be estimated.

2.4.2. *Marginal covariance structure of (Y_{NWP}) .* The block structure of the space–time covariance of the data suggests expressing wind speed at each station as a linear transformation of an unobserved common signal Y_0 with added noise; see Figure 3. Intuitively, we can think of this common signal as an average flow over the studied region. The wind speed at each site is a linear transformation of this average flow. The temporal dynamics of the unobserved signal is modeled with a squared exponential covariance. The following structure is used:

$$Y(b_i, s_j) = \Psi_{s_j} Y_0(b_i) + \varepsilon_{s_j}(b_i),$$

where Y represents Y_{NWP} in this paragraph and will represent $(Y_{\text{Obs}}|Y_{\text{NWP}})$ in Section 2.4.4, b_i is a temporal window of $h = 24$ lags, s_j is the spatial location, and Ψ_{s_j} is an $h \times h$ -matrix. The various ε_{s_j} are assumed independent from each other and from Y_0 . This model is inspired in part by an earlier study [Constantinescu and Anitescu (2013)] where the Ψ operators were used to represent a known functional relation. In our case, Ψ_{s_j} is a parameterized matrix that is inferred from the data.

The overall space–time covariance of Y_{NWP} has the following structure:

$$\begin{aligned} \Sigma_{\text{NWP}}(\cdot, s_i; \cdot, s_j) &= \text{cov}(Y_{\text{NWP}}(\cdot, s_i), Y_{\text{NWP}}(\cdot, s_j)) \\ (2.13) \qquad \qquad \qquad &= (\Psi_{s_i} \Gamma_0 \Psi_{s_j}^T) + \delta_{i-j} \Gamma_{s_i}, \end{aligned}$$

for $j \in \{1, \dots, J_0\}$ and where δ stands for the Kronecker symbol. The $h \times h$ -matrices Γ_{s_j} are written as

$$\Gamma_{s_j}[l, k] = \sigma_{s_j} \exp(-\lambda_{s_j}(|t_k - t_l|)^2) + \delta_{k-l} \gamma_{s_j},$$

with

$$\sigma_{s_j} = \sigma_1 + \sigma_2 \text{Lat}(s_j) + \sigma_3 \text{Long}(s_j),$$

$$\lambda_{s_j} = \lambda_1 + \lambda_2 \text{Lat}(s_j) + \lambda_3 \text{Long}(s_j),$$

$$\gamma_{s_j} = \gamma_1 + \gamma_2 \text{Lat}(s_j) + \gamma_3 \text{Long}(s_j),$$

and

$$\Gamma_0[l, k] = \sigma_0 \exp(-\lambda_0(|t_k - t_l|)^2) + \delta_{k-l} \gamma_0,$$

where $\sigma_1, \sigma_2, \sigma_3, \lambda_1, \lambda_2, \lambda_3, \gamma_1, \gamma_2, \gamma_3, \sigma_0, \lambda_0,$ and γ_0 are positive real numbers to be estimated.

Following the data analysis, the $h \times h$ -matrices Ψ_{s_j} are parameterized as tridiagonal matrices. Given the study of the variance in space and time, the diagonal and off-diagonal quantities are modeled with a quadratic dependence in time and spatially dependent coefficients. The diagonal, subdiagonal, and superdiagonal of

the matrix Ψ_{s_j} are written respectively as

$$\begin{aligned} \Psi_{s_j}[i, i] &= (1 + v_1\text{Lat}(s_j) + v_2\text{Long}(s_j)) \\ &\quad + (1 + v_3\text{Lat}(s_j) + v_4\text{Long}(s_j)) \times i \\ &\quad + (1 + v_5\text{Lat}(s_j) + v_6\text{Long}(s_j)) \times i^2, \\ \Psi_{s_j}[i, i - 1] &= (1 + v_7\text{Lat}(s_j) + v_8\text{Long}(s_j)) \\ &\quad + (1 + v_9\text{Lat}(s_j) + v_{10}\text{Long}(s_j)) \times i \\ &\quad + (1 + v_{11}\text{Lat}(s_j) + v_{12}\text{Long}(s_j)) \times i^2, \\ \Psi_{s_j}[i, i + 1] &= (1 + v_{13}\text{Lat}(s_j) + v_{14}\text{Long}(s_j)) \\ &\quad + (1 + v_{15}\text{Lat}(s_j) + v_{16}\text{Long}(s_j)) \times i \\ &\quad + (1 + v_{17}\text{Lat}(s_j) + v_{18}\text{Long}(s_j)) \times i^2, \end{aligned}$$

for $i \in \{1, \dots, h\}$; v_1, \dots, v_{18} are real numbers to be estimated. We work in relatively small areas and use distances in latitude and longitude here and for the rest of this work.

2.4.3. *Conditional mean structure of $(Y_{\text{Obs}}|Y_{\text{NWP}})$.* In Royle and Berliner (1999), several configurations of the transition matrix Λ are proposed depending on its use. For instance, a transition matrix from atmospheric pressure to wind speed is derived from geostrophic equations in Royle et al. (1999). The observations exhibit daily and half-daily periodicity (with various intensities depending on the month of the year) and spatial patterns; see Figure 2. However, the relation between the two datasets does not exhibit significant time dependence that requires a time-varying dependence. We use spatial and temporal neighbors to explain the observed wind speed. The land use LU is included in the transition matrix because it defines different behaviors in the NWP model data. We choose the following transition between the two datasets:

$$\mu_{\text{Obs}|\text{NWP}}(t, s) = E(Y_{\text{Obs}}(t, s)|Y_{\text{NWP}}) = \mu(t, s) + (\Lambda Y_{\text{NWP}})(t, s),$$

with

$$\begin{aligned} \mu(t, s) &= \left(\beta_7 + \beta_8 \cos\left(\frac{2\pi t}{24}\right) + \beta_9 \sin\left(\frac{2\pi t}{24}\right) \right. \\ &\quad \left. + \beta_{10} \cos\left(\frac{2\pi t}{12}\right) + \beta_{11} \sin\left(\frac{2\pi t}{12}\right) \right) \\ &\quad \times (1 + \alpha_3\text{Lat}(s) + \alpha_4\text{Long}(s)), \\ (\Lambda Y_{\text{NWP}})(t, s) &= \sum_{i=1}^h \rho(\text{LU}(s), |t - t_i|) \\ &\quad \times \sum_{k=1}^3 \Phi_k(\Delta\text{Lat}, \Delta\text{Long})(s, s_k) Y_{\text{NWP}}(t_i, s_k), \quad t_1 \leq t \leq t_h, \end{aligned}$$

where

- $\rho(\cdot, \cdot)$ are temporal weights, parameterized according to $\rho(l, \Delta t) = \rho_0(l) \exp(-\rho_1(l)|\Delta t|) + \rho_2(l)$, for the time difference Δt in $\{0, \dots, h-1\}$; the integer $l \in \{1, \dots, n\}$ is the land use value of the closest grid point of s ; $\rho(l, 0) = 1$ for identifiability purposes;
- $\Phi(\Delta \text{Lat}, \Delta \text{Long}) = \phi_{0,\cdot} + \phi_{1,\cdot} \Delta \text{Lat} + \phi_{2,\cdot} \Delta \text{Long}$, with $\Delta \text{Lat}(s_i, s_j) = |\text{Lat}(s_i) - \text{Lat}(s_j)|$ and $\Delta \text{Long}(s_i, s_j) = |\text{Long}(s_i) - \text{Long}(s_j)|$;
- s_1, s_2, s_3 are nearest spatial neighbor grid points of s selected according to the radial distance, but other distances are possible. Moreover, for simplicity we consider here nearest neighbors, but other choices of predictors can be made, such as upwind stations; and
- all the parameters $\beta_7, \dots, \beta_{11}, \alpha_3, \alpha_4, (\rho_0(l))_{l=1, \dots, n}, (\rho_1(l))_{l=1, \dots, n}, (\rho_2(l))_{l=1, \dots, n}, (\phi_{0,k})_{k=1, \dots, 3}, (\phi_{1,k})_{k=1, \dots, 3}, (\phi_{2,k})_{k=1, \dots, 3}$ are real numbers to be estimated.

2.4.4. *Conditional covariance structure of $(Y_{\text{Obs}}|Y_{\text{NWP}})$.* Analysis of the empirical conditional covariance suggests the use of the parametric shape proposed in (2.13), with a different set of parameters.

2.5. *Estimation of the parameters.* Maximum likelihood is chosen for estimating the parameters. The likelihood of the model for the observed dataset $(y_{\text{Obs}}(t_1, \dots, t_{k_K}; s_1, \dots, s_{J_0}), y_{\text{NWP}}(t_1, \dots, t_{k_K}; s_1, \dots, s_{J_0}))$ is written as

$$\begin{aligned} \mathcal{L}(\theta; y_{\text{Obs}}(t_1, \dots, t_{k_K}; s_1, \dots, s_{J_0}), y_{\text{NWP}}(t_1, \dots, t_{k_K}; s_1, \dots, s_{J_0})) \\ &= p_{\theta}(y_{\text{Obs}}(t_1, \dots, t_{k_K}; s_1, \dots, s_{J_0}), y_{\text{NWP}}(t_1, \dots, t_{k_K}; s_1, \dots, s_{J_0})) \\ &= p_{\theta}(y_{\text{NWP}}(t_1, \dots, t_{k_K}; s_1, \dots, s_{J_0})) \\ &\quad \times p_{\theta}(y_{\text{Obs}}(t_1, \dots, t_{k_K}; s_1, \dots, s_{J_0}) | y_{\text{NWP}}(t_1, \dots, t_{k_K}; s_1, \dots, s_{J_0})). \end{aligned}$$

This is the particular instantiation of (2.5). Each day, the WRF model is run independently from the previous day. Here, we have assumed short temporal error correlations (see Figure 3). Furthermore, the forecasts are restarted from reanalyses, and at least in the linear case the innovations are independent of observations [Shumway and Stoffer (2010), Section 6.3]. Therefore, we consider statistical independence between each day, which leads to the following product:

$$\begin{aligned} p_{\theta}(y_{\text{NWP}}(t_1, \dots, t_{k_K}; s_1, \dots, s_{J_0})) &= \prod_{i=1}^K p_{\theta}(y_{\text{NWP}}(t_{k_i}; \mathbf{S}), \dots, y_{\text{NWP}}(t_{k_i+23}; \mathbf{S})) \\ &= \prod_{i=1}^K p_{\theta}(y_{\text{NWP}}(b_i; \mathbf{S})), \end{aligned}$$

where $S = \{s_1, \dots, s_{J_0}\}$ and $\{b_1, \dots, b_K\} = \{t_1, \dots, t_{24}, t_{25}, \dots, t_{k_K}\}$ with $b_i = \{t_{k_i}, \dots, t_{k_i+23}\}$. For each $i \in \{1, \dots, K\}$ the associated log-likelihood is written as

$$\log(p_\theta(y_{NWP}(b_i; S))) = \log(f(y_{NWP}(b_i; S); \mu_{NWP}, \Sigma_{NWP})),$$

where $f(\cdot; \mu, \Sigma)$ denotes the multivariate normal density with mean μ and covariance Σ and μ_{NWP} and Σ_{NWP} are the parametric mean and covariance, respectively, expressed in (2.12) and (2.13). Similarly, the likelihood of conditional distribution is written and maximized. In practice, a preliminary least-squares estimation of the parameters is realized between the empirical and parametric first- and second-order structures of Y_{Obs} and Y_{NWP} . These estimates are used as initial conditions of the maximum likelihood procedure.

2.6. *Kriging.* Predictions of Y_{Obs} from Y_{NWP} are obtained from the kriging equations [Stein (1999)], with the mean and covariance defined by (2.11). For t_0 in $b_{K+1} = \{t_{k_{K+1}}, \dots, t_{k_{K+1}+h-1}\}$ and s_0 in $\{S, S_0\} = \{1, \dots, J_0, J_0 + 1, \dots, J\}$ defined in (2.3), we have

$$(2.14) \quad \begin{aligned} & (Y_{Obs}(t_0; s_0) | Y_{NWP}(b_{K+1}; 1, \dots, J_0, J_0 + 1, \dots, J)) \\ & \sim \mathcal{N}(\hat{\mu}_{Obs}(t_0; s_0), \hat{\Sigma}_{Obs}(t_0; s_0)) \end{aligned}$$

with

$$(2.15a) \quad \begin{aligned} \hat{\mu}_{Obs}(t_0; s_0) &= (\mu + \Lambda \mu_{NWP})(t_0; s_0) + c_0^T \Sigma_{NWP}^{-1}((b_{K+1}; S, S_0); \\ & (b_{K+1}; S, S_0))(Y_{NWP} - \mu_{NWP})(b_{K+1}; S, S_0), \end{aligned}$$

$$(2.15b) \quad \begin{aligned} \hat{\Sigma}_{Obs}(t_0; s_0) &= \Sigma_{Obs}((t_0; s_0); (t_0; s_0)) + c_0^T \Sigma_{NWP}^{-1}((b_{K+1}; S, S_0); \\ & (b_{K+1}; S, S_0))c_0, \end{aligned}$$

$$c_0 = \Sigma_{Obs, NWP}((t_0; s_0); (b_{K+1}; S, S_0)).$$

The distribution (2.14) is used to generate the scenarios of prediction of wind speed in Section 4. This is in fact the predictive distribution presented in (2.6). Note that the conditional distribution (2.8) could directly be used for predictions; however, this would not enable predictions with misaligned data. Nevertheless, in our results the predictive variance is on average 2.7% greater when using directly the conditional distribution instead of the kriging one (2.14).

3. **Wind data.** In order to improve forecasts from the considered numerical model, two sources of data are combined: ground measurements and WRF model outputs. The measurement data are recorded across an irregular network, and at each observational station we pick the closest gridded point of NWP outputs. As a result, the two datasets have the same number of spatial locations; however, the

proposed model is not restricted to this spatial layout and can handle datasets with different numbers of stations. In the following, the time series of the two datasets are filtered in time by a moving average process over a window of 1 hour to remove small-scale effects and focus on a larger temporal scale; they are picked every hour. We focus on a region around Lake Michigan in the United States; however, the framework proposed here is not specific to that region.

3.1. *Direct observations.* Observational data are extracted from the Automated Surface Observing System (ASOS) network, available at <ftp://ftp.ncdc.noaa.gov/pub/data/asos-onemin>. The network of collecting stations covers the U.S. territory. The studied data are 1-minute data selected from Wisconsin, Illinois, Indiana, and Michigan; see Figure 1. The measured wind speed is discretized in integer knots (one knot is about 0.5 m/s). We do not apply any additional treatment to account for this discretization because the data are filtered over a window of 1 hour; see [Slougher, Gneiting and Raftery \(2010\)](#) for a discussion of the discretization of the wind speed. The orography of this region is simple and flat; however, the presence of Lake Michigan has a strong impact on the wind conditions. Several months are investigated and reveal different behaviors; in particular, periodicities differ from winter to summer months. In the following, for homogeneity purposes the dataset of 31 stations is subdivided into three spatial clusters of respectively




FIG. 1. Map of the considered area (Midwest; visible are Lake Michigan, Michigan, Illinois, Indiana, and Wisconsin). Clusters are depicted with different colors: respectively in black dots are the 11 stations of subregion C_1 , in triangle the 12 stations of C_2 , and in diamond the 8 stations of C_3 .

11, 12, and 8 stations, depicted in Figure 1. A spatial clustering is performed on wind speed in order to distinguish among different average regional weather conditions. This is a proxy for different NWP forecast behaviors. These three clusters are treated independently hereafter.

3.2. Numerical weather prediction data. State-of-the-art NWP forecasts are generated by using WRF v3.6 [Skamarock et al. (2008)], which is a state-of-the-art numerical weather prediction system designed to serve both operational forecasting and atmospheric research needs. WRF has a comprehensive description of the atmospheric physics that includes cloud parameterization, land-surface models, atmosphere-ocean coupling, and broad radiation models. The terrain resolution can go up to 30 seconds of a degree (less than 1 km^2). The NWP forecasts are initialized by using the North American Regional Reanalysis fields dataset that covers the North American continent (160W-20W; 10N-80N) with a resolution of 10 minutes of a degree, 29 pressure levels (1000–100 hPa, excluding the surface), every 3 hours from the year 1979 until the present. Simulations are started every day during January and August 2012 and cover the continental United States on a grid of $25 \times 25 \text{ km}$ with a time resolution of 10 minutes.

4. Results. In this section, we first analyze the estimated parameters and then explore qualitatively and quantitatively the ability of the model to provide accurate forecasts. Two months of the year (January and August) are considered and are studied independently in order to investigate the model performance under different conditions. Moreover, the model is compared with two embedded models: one model with only temporal dependencies but without spatial interactions and one model without temporal or spatial dependencies. For each month, the model is trained on contiguous two-thirds of the month and predicted on the remaining third. The training periods are rolled over the three possible permutations of one-third to fill in the entire month.

4.1. Analysis of the estimated parameters. In this section, we investigate the maximum likelihood estimation of the mean and covariance of the process. First, the empirical mean and covariance are compared with the fitted parametric ones proposed in Section 2. The mean of the process ($Y_{\text{Obs}}, Y_{\text{NWP}}$) is depicted in Figure 2; for each station, the mean at each hour of the day is plotted. The structure of the estimated mean of the two processes is accurately reproduced in terms of temporal and spatial patterns, the empirical and fitted spatial trends are not pronounced for this studied month. In Figure 3, the empirical and fitted space–time correlations are plotted. A great part of the structure is captured by the proposed parametric shapes; however, the global shapes tend to be smoothed by the parametric models. The nonseparability between space and time that is visible on the empirical off-diagonal blocks is not entirely captured by the parametric model and tends to be overestimated. Analysis of the matrices Ψ_s that are involved in the




FIG. 2. Empirical and fitted parametric mean of wind speed at each hour of a day and at each station in the subregion C_1 in August. Vertical lines separate each station. Within each of these windows, each hour of the day is considered. Top panel: mean of Y_{NWP} ; bottom panel: mean of Y_{Obs} .

covariance model (2.13) reveals different configurations given the subregion and the period of the year. These can be expected because these operators can be interpreted as a linear projector of a process that is common to all the stations. Average air flows differ according to the season and the spatial location; the dependence

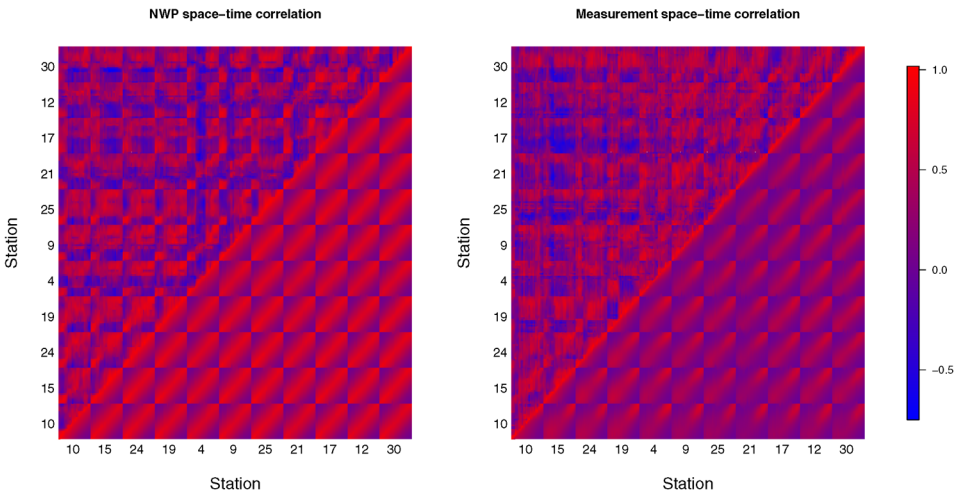


FIG. 3. Empirical and fitted parametric space–time correlation estimated in August 2012 in the subregion C_1 . Left: NWP space–time correlation; right: measurement correlation. The upper left part of the matrix corresponds to the empirical correlation; the right bottom blocks display the fitted parametric correlation. Diagonal blocks correspond to temporal correlations at each station, and off-diagonal blocks correspond to temporal cross-correlation between stations.




FIG. 4. Maximum likelihood estimation of the parameters of the model and associated standard deviation of estimation. Vertical lines separate the different sets of parameters. From left to right are the parameters of $\mu_{\text{Obs|NWP}}$, $\Sigma_{\text{Obs|NWP}}$, μ_{NWP} , and Σ_{NWP} . Parameters are plotted following this order: $\mu_{\text{Obs|NWP}}$: $\beta_7, \dots, \beta_{11}, \alpha_3, \alpha_4, (\rho_0^{(l)})_{l=1, \dots, n}, (\rho_1^{(l)})_{l=1, \dots, n}, (\rho_2^{(l)})_{l=1, \dots, n}, (\phi_{0,k})_{k=1, \dots, 3}, (\phi_{1,k})_{k=1, \dots, 3}, (\phi_{2,k})_{k=1, \dots, 3}, \Sigma_{\text{Obs|NWP}}$: $\nu_1, \dots, \nu_{18}, \sigma_0, \lambda_0, \gamma_0, \sigma_1, \sigma_2, \sigma_3, \lambda_1, \lambda_2, \lambda_3, \gamma_1, \gamma_2, \gamma_3, \mu_{\text{NWP}}$: $(\beta_k)_{k=0, \dots, 6}, (\alpha_0^{(l)})_{l=1, \dots, n}, \alpha_1, \alpha_2, \Sigma_{\text{NWP}}$: same indexation as $\Sigma_{\text{Obs|NWP}}$.

from a common process that would contain this information is likely to differ in space and in time across the year.

The matrix Λ , which appears in both the mean and covariance components, is important because it links the NWP forecasts to the objective predictive quantities. The analysis of Λ reveals that the intensity of temporal dependence varies with the land use; however, the temporal persistence is curtailed to a few hours across the different land use.

In the second step, the uncertainty associated with the estimation of the parameters is accounted for. Figure 4 plots the maximum likelihood estimation of the parameters and the associated standard deviation of estimation, given by the inverse of the Hessian of the log-likelihood calculated at the maximum. Note that the uncertainty is relatively narrow. For both distributions, the estimation of the covariance parameters tends to show more variance than the one of the mean structure parameters. The greatest estimation variance is presented by several parameters ν_i that appear in the matrices Ψ of both covariances in Section 2.4.2. In the parameters of $\mu_{\text{Obs|NWP}}$, parameters with a high estimation variance are the ones associated with ϕ defined in Section 2.4.3. In these cases, a lack of data in the estimation of these specific parameters may cause this high estimation variance.

4.2. Assessment of the quality of the predictive model. In this part, samples (or scenarios) are generated from the predictive distribution defined by equation (2.14). The mean of these samples can be used as a pointwise prediction, but the objective here is to embed the uncertainty associated with the prediction by working with samples from the predictive distribution. The presented predictive scenarios are back-transformed by using the inverse Box–Cox transformation.

4.2.1. Description of the embedded models. In the current study, the model proposed in Section 2 is compared to two embedded versions of it. These two ver-

sions of the model use the same NWP-predictors from Section 2.4.3; however, the covariance structures are less sophisticated. First, we consider a model where only temporal dependence is accounted for and where spatial interactions are ignored in the covariance structures. This model will be called temporal model in the following. This version of the model has all the characteristics of the full model of Section 2, except that both covariance structures do not handle the spatial cross-covariance of equation (2.13) between two stations s_i and s_j , and the matrices Ψ_{s_i} are diagonal. Second, another reduction of the model is implemented, where only the bias in the space–time mean is corrected, and neither temporal nor spatial dependencies are accounted for in the covariance structures. We call this reduced version bias-correction model. In this reduced version, only characteristics described by equation (2.12) and Section 2.4.3 hold true, and the covariance is reduced to a site-specific variance without temporal features.

The proposed model and its embedded versions can be linked to the NGR model proposed by Gneiting et al. (2005) through the conditional distribution of equation (2.8). Indeed, the ensemble members of the NGR model could be replaced here by the NWP-predictors in space and time described in Section 2.4.3 with some weight constraints given by μ and Λ . The full space–time model can be seen as a multidimensional NGR in space and time with a constrained covariance structure given in Section 2.4.4. Similarly, the temporal model version appears as a multidimensional NGR in time with a constrained squared-exponential temporal covariance structure. However, the second reduction, the bias-corrected model, relates to an univariate NGR model where the dispersion of the ensemble would not be corrected. Further correspondences are harder to draw between the two models because predictions with the present model are made through a kriging procedure that is derived from the joint distribution of Y_{Obs} and Y_{NWP} , whereas in the NGR model the prediction is made from an analogue of the conditional distribution ($Y_{\text{Obs}}|Y_{\text{NWP}}$). Both the proposed and the NGR models have their respective merits, the former has more constraints but handles nonaligned data and accounts for the statistical space–time structure of Y_{NWP} , whereas the latter is less parametrically constrained but is more difficult to use in a multidimensional context.

4.2.2. Qualitative exploration of the predictions. We first investigate the time series of the predictions by a visual assessment in Figure 5. We next explore how the forecast scenarios are representative in terms of calibration through rank histograms (Figure 6) and in terms of temporal and spatio-temporal structures through spectral and correlation investigations (Figures 7 and 8).

Diagnostic tools and scores to evaluate vector-valued prediction have been proposed recently for single and ensemble forecasts [Gneiting et al. (2008), Pinson and Girard (2012), Scheuerer and Hamill (2015), Smith and Hansen (2004), Thorarindottir, Scheuerer and Heinz (2016)]. Some of these criteria evaluate simultaneously univariate predictive skills and multivariate dependencies. Criteria for ensemble forecasts can be used for predictive scenarios where each sample of

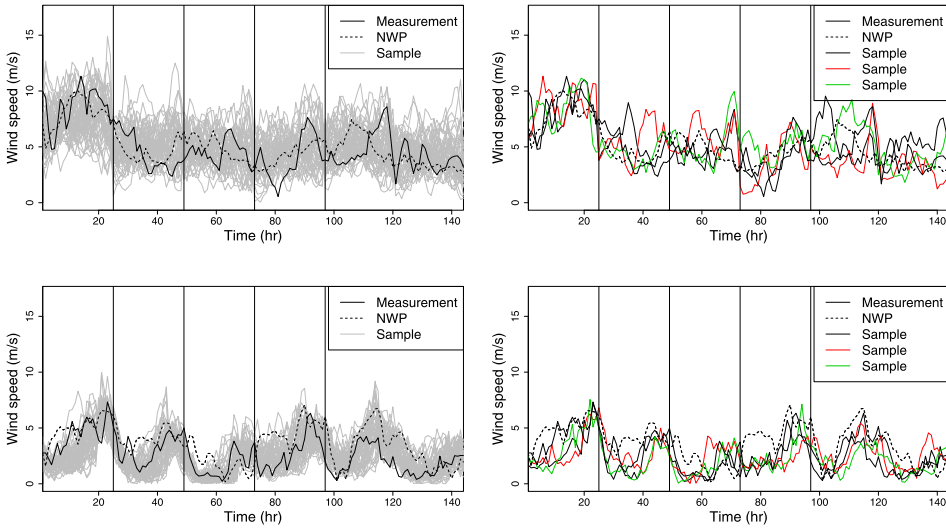


FIG. 5. Time series of wind speed at the station with the median RMSE in subregion C_1 . January 2012 (top) and August 2012 (bottom) for six days that are separated by vertical lines. Left panels: 50 predictive samples are plotted; right panel: 3 samples are plotted.

the predictive distribution is treated as a member of the ensemble [Pinson and Girard (2012)]. Here, we focus on univariate criteria for the calibration separately from the temporal and spatio-temporal dependencies; however, diagnostic tools such as multivariate rank histograms [Gneiting et al. (2008), Thorarinsdottir, Scheuerer and Heinz (2016)] can be used to assess univariate predictive skills and multivariate dynamics. Multivariate scores are discussed and compared in Section 4.2.3.




FIG. 6. Univariate rank histograms at the station with the median RMSE in subregion C_1 . Left: January 2012; right: August 2012. The red error bars are 95%-confidence intervals (confidence intervals are computed with a binomial distribution); the horizontal red line represents the density of a uniform distribution.

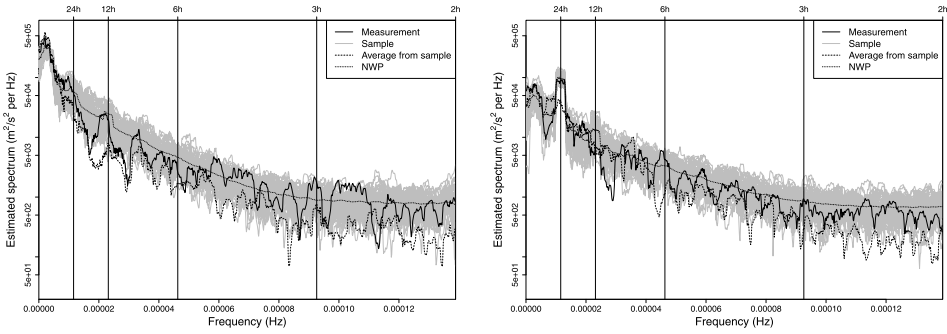


FIG. 7. *Estimated spectrum in January (left) and August (right) for the station with the median RMSE in subregion C_1 .*

Visual assessment of time series. We investigate observed time series and generated predictive scenarios for part of the months of January and August; see Figure 5. Measured wind speed, which is to be predicted, is plotted as a reference in order to evaluate the accuracy of the prediction. NWP wind forecasts are also plotted because they are predictors and a target to be improved with respect to the measurements. For both months, the global trend of the measured time series is well captured by the predictive mean and by the scenarios. The predictive samples cover the measurements that are to be predicted (see left panels); and the predictive mean realizes, most of the time, an improvement with respect to the NWP forecasts. Moreover, each sample has a temporal dynamics consistent with the observed temporal behavior (right panels). The scenarios take negative values; however, such values happen only 0.3% of the time in January and 0.2% in August. The improvement of the proposed prediction is more visible in August (bottom panels), likely because of the periodic components that are stronger




FIG. 8. *Space–time correlation of wind speed in January in subregion C_1 . The upper left part of the matrix corresponds to the empirical correlation; the right bottom blocks display the fitted correlation. Diagonal block correspond to temporal correlations at each station, and off-diagonal blocks correspond to temporal cross-correlation between stations. From left to right: full space–time model, temporal model, bias-correction model.*

TABLE 1

Statistics and metrics for the station, representing the median RMSE in each cluster denoted as C_i , for $i = 1, 2, 3$. ES is the energy score and M/S is measurements or samples. They are evaluated on the concerned month for time prediction. Associated with the model RMSE is the percentage of improvement of the model with respect to the NWP data. The bottom block corresponds to metrics computed on the three proposed reductions of the model

Model	RMSE (m/s)	ES (m/s)	M/S	Mean(Y_{Obs}) (m/s)	Var(Y_{Obs}) (m^2/s^2)
NWP (Jan. 2012, C_1)	1.85	48	M	4.7	5.6
Model (Jan. 2012, C_1)	1.64 (11%)	29	S	4.68	5.08
NWP (Aug. 2012, C_1)	1.7	44	M	2.5	1.99
Model (Aug. 2012, C_1)	1.15 (34%)	21	S	2.5	2
NWP (Jan. 2012, C_2)	2.01	52	M	3.74	5.09
Model (Jan. 2012, C_2)	1.62 (19%)	43	S	3.78	4.96
NWP (Aug. 2012, C_2)	1.62	41	M	2.76	1.97
Model (Aug. 2012, C_2)	1.04 (36%)	19	S	2.86	2.23
NWP (Jan. 2012, C_3)	2.31	60	M	4.72	6.54
Model (Jan. 2012, C_3)	1.72 (25%)	32	S	4.71	6.09
NWP (Aug. 2012, C_3)	1.72	44	M	2.31	2.23
Model (Aug. 2012, C_3)	1.02 (41%)	18	S	2.33	2.04
NWP (Jan. 2012, C_1)	1.85	48	M	4.7	5.6
Model (Jan. 2012, C_1)	1.64 (11%)	29	S	4.68	5.08
Model Temp. (Jan. 2012, C_2)	1.63 (12%)	29	S	4.45	4.37
Model Bias (Jan. 2012, C_2)	1.81 (2.1%)	32	S	5.23	4.93

in this period of the year are well captured by the model, and also because the global variance of the wind speed is smaller in August than in January; see also Figure 7. Furthermore, the spread of the scenarios is more important in January than in August, likely because the wind speed has more variability in winter, as illustrated in the observed variances in Table 1, which may make it less predictable. We note that the scenarios are not spreading at the end of each prediction window, as observed in the literature. The reason is that the NWP predictors are available over the entire prediction window and such spread increase is not obvious in the model–measurement discrepancy.

Calibration assessment. Rank histograms are commonly used to assess the calibration of predictive ensembles; they can be seen as analogous to ensembles of the probability integral transform (PIT) that evaluates the calibration of single forecasts [Anderson (1996), Hamill (2001)]. In Figure 6, univariate rank histograms are plotted for one station, and each predictive scenario is treated as an ensemble member in the rank histogram. No strong signs of trend or over- and under-dispersion are seen in these histograms, except the upper tail of distribution for August. The two panels indicate a well-calibrated ensemble of samples; the hori-

zontal line representing the uniform distribution is covered most of the time by the confidence intervals associated with the proportions of the histogram. The predictions from January tend to show better calibration than the August ones do.

Temporal spectral analysis. The spectral content of the scenarios and of the observations is estimated and depicted in Figure 7; the average spectrum of the estimated spectrum on each sample is also plotted. We note that a similar analysis was carried out in Bouallegue et al. (2016). The estimated spectra of the scenarios cover most of the spectrum of the observations. The overall shape of the estimated spectra and of the average spectrum indicates a robust agreement, especially in August where small frequencies are accurately captured. In this and other spectral estimates, the spectral content at high frequency is sometimes slightly overpredicted; we believe the reason is that the time series has discontinuities at the boundaries between temporal blocks due to overlapping prediction windows, which in turn are a result of assimilating data and restarting the NWP forecast. Nevertheless, the features of the spectrum of the measurements appear well captured by our model. Notice that the spectrum of NWP data in August provides a poor description of the observations; however, the model is able to correct this. Therefore, our model appears to be a suitable and realistic wind scenario generator.

Space–time correlation structure. In Figure 8, empirical space–time correlations of measurements are compared with those of scenarios from the model proposed in Section 2 and with those of scenarios from the simplified models. The space–time model enables of the space–time correlation structure to be captured, whereas the embedded models capture significantly less information than does the full model. These two embedded models miss most of the spatial cross-correlation between stations. From this figure, we see the importance of space–time information in the structure of the wind speed prediction.

4.2.3. *Quantitative assessment of the quality of the predictions.* As our second step, we assess quantitatively the overall improvement of the model in comparison with the WRF model outputs. We study general metrics because we would like to preserve a general application scope. See Pinson (2013) for reflections on links between improvement of these general metrics and user-specific metrics and also for general challenges associated with forecast verification.

Univariate predictive skills. In Table 1, the root mean square error (RMSE) is computed for the predictive mean of the proposed distribution and for the NWP forecasts. We consider also the energy score (ES), which represents a generalization of the continuous ranked probability score (CRPS) for ensemble predictions [see Gneiting et al. (2008), Pinson and Girard (2012)]. This metric is an omnibus metric that enables comparison of ensemble forecasts and scenarios with point-wise prediction; it is computed on predictive samples and on NWP forecasts. The energy score is a proper scoring rule, the lower the energy score, the better the

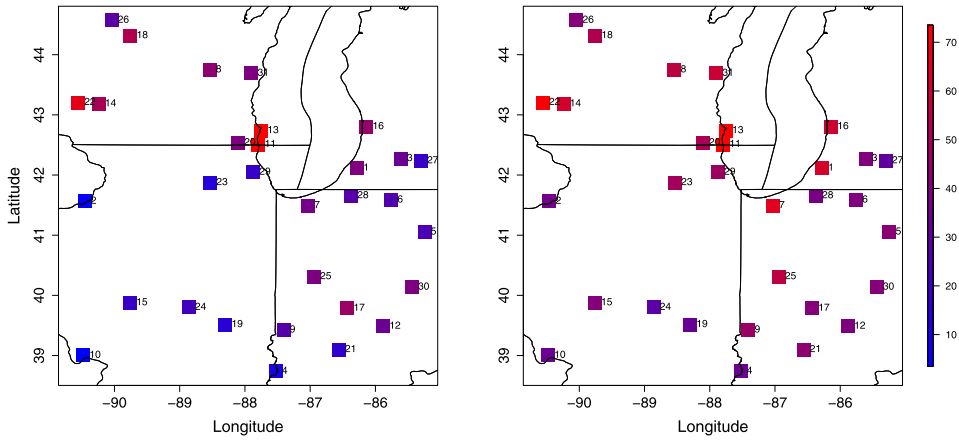


FIG. 9. Maps of percentage of improvement of RMSE of the proposed prediction with respect to RMSE of NWP prediction. Left: January; right: August.

proposed forecast. In subregion C_3 , the model shows the greatest improvement in terms of RMSE and energy score, likely because this subregion is the most homogeneous one and the model fits the best for this subregion. In subregion C_2 , the model also presents good improvements in terms of RMSE and energy score, due mostly to the influence of Lake Michigan on the results. Indeed around Lake Michigan, NWP forecasts overestimate wind speed values. Indeed, the NWP embeds this presence through the lake mask and land use, but this may be overestimated in comparison with the behaviors of the observations. The improvement in terms of RMSE is more significant in August, likely because the periodic components are well captured by the model and the variability of wind speed is smaller. The energy score clearly favors the proposed model in comparison with the WRF outputs. Most of the means and variances of the observations are well captured by the prediction made with the model. In addition, in Table 1, the metrics are computed for the full space–time model and for its two simplifications (temporal model and bias-correction model). The proposed space–time model reveals better results than do the embedded models, as expected. Results are presented only for January in subregion C_1 ; however, similar conclusions can be drawn from the other months and subregions. In Figure 9, maps of the percentage of improvement of RMSE are shown for each station and for both months. The improvement is greater in August and also around Lake Michigan in subregion C_2 .

Multivariate predictive skills. In Figure 10, we compute three multivariate proper scores for the proposed model and its two reductions: the energy score; the Dawid–Sebastiani score [Dawid and Sebastiani (1999)], which is equivalent to the log-score for a Gaussian multidimensional predictive distribution; and the variogram-based score [Scheuerer and Hamill (2015)], which measures the dissimilarity between variograms of the observation and of the forecasts. The distributions of the

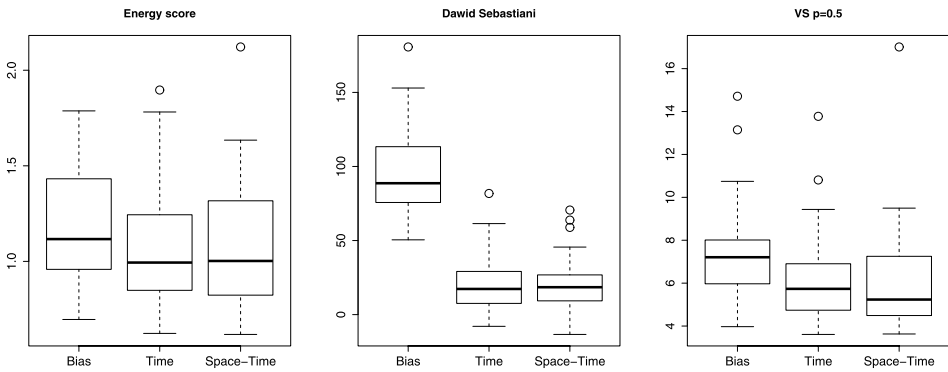


FIG. 10. *Energy score, Dawid–Sebastiani score, and Variogram-based score computed for the station with median RMSE in subregion C_1 and its closest neighbor for 24 hours, for January 2012. The variogram-based score is computed with equal weights and for an order $p = 0.5$, which shows here the most discrimination between the models. Scores are computed for the proposed model and its two reductions: a temporal model and a bias-correction model. Each point of the boxplot corresponds to the score computed for one day of the month.*

scores computed for every day of the month of January are displayed in the boxplots of Figure 10. The energy score may not be discriminative in the case of misspecified correlations; see Pinson and Tastu (2013), Scheuerer and Hamill (2015) and Figure 10. However, it helps discriminate accurate intensity in forecasts; see Table 1. Note that in Figure 10 the values of the energy score differ from those in Table 1 because different windows are under consideration: a space–time window is used in Figure 10 and a temporal one in Table 1. The Dawid–Sebastiani score enables the model that corrects only the bias to be strictly distinguished from the models that embed also the temporal and space–time structures. The variogram-based score enables the three models to be distinguished, the separation between the temporal and space–time models is not as clear as expected. These multivariate proper scores assess different properties of the predictions and exhibit different results. A sensible approach, therefore, is to apply multiple scores in order to select appropriate predictive multivariate models and assess their skills.

A potential limitation of our approach stems from the assumption that an inherent stationarity exists across the calibration and forecast windows. In our case, all days of a third of each month are modeled with the same model. We currently are exploring introducing nonstationarity between each temporal block (day) in order to account for such potential shortcomings.

5. Conclusions. We have introduced a statistical space–time modeling framework for predicting atmospheric wind speed based on deterministic numerical weather predictions and historical measurements. We have used a Gaussian multivariate space–time process that combines multiple sources of past physical model outputs and measurements along with model predictions to forecast wind speed at

observation sites. We applied this strategy to surface wind-speed forecasts for a region near the U.S. Great Lakes. The results show that the prediction is improved in the mean-squared sense as well as in probabilistic scores. Moreover, the samples are shown to produce realistic wind scenarios based on the sample spectrum. Using the proposed model, one can correct the first- and second-order space–time structure of the numerical forecasts in order to match the structure of the measurements.

Acknowledgments. We thank Professor Michael Stein for comments on multiple versions of this manuscript. We thank Michael Scheuerer for providing us codes on multivariate scores and his helpful comments.

REFERENCES

- AILLIOT, P., FRÉNOT, E. and MONBET, V. (2006). Long term object drift forecast in the ocean with tide and wind. *Multiscale Model. Simul.* **5** 514–531. [MR2247761](#)
- ANDERSON, J. L. (1996). A method for producing and evaluating probabilistic forecasts from ensemble model integrations. *J. Climate* **9** 1518–1530.
- ANITESCU, M., CHEN, J. and WANG, L. (2012). A matrix-free approach for solving the parametric Gaussian process maximum likelihood problem. *SIAM J. Sci. Comput.* **34** A240–A262.
- APANASOVICH, T. and GENTON, M. G. (2010). Cross-covariance functions for multivariate random fields based on latent dimensions. *Biometrika* **97** 15–30. [MR2594414](#)
- BAO, L., GNEITING, T., GRIMIT, E. P., GUTTORG, P. and RAFTERY, A. E. (2010). Bias correction and Bayesian model averaging for ensemble forecasts of surface wind direction. *Mon. Weather Rev.* **138** 1811–1821.
- BARAN, S. (2014). Probabilistic wind speed forecasting using Bayesian model averaging with truncated normal components. *Comput. Statist. Data Anal.* **75** 227–238.
- BARAN, S. and LERCH, S. (2015). Log-normal distribution based Ensemble Model Output Statistics models for probabilistic wind-speed forecasting. *Q. J. R. Meteorol. Soc.* **141** 2289–2299.
- BARAN, S. and LERCH, S. (2016). Mixture EMOS model for calibrating ensemble forecasts of wind speed. *Environmetrics* **27** 116–130. [MR3481324](#)
- BERLINER, M. (2000). Hierarchical Bayesian modeling in the environmental sciences. *AStA Adv. Stat. Anal.* **84** 141–153.
- BERROCAL, V. J., GELFAND, A. E. and HOLLAND, D. M. (2012). Space–time data fusion under error in computer model output: An application to modeling air quality. *Biometrics* **68** 837–848.
- BOUALLEGUE, Z. B., HEPELMANN, T., THEIS, S. E. and PINSON, P. (2016). Generation of scenarios from calibrated ensemble forecasts with a dual-ensemble copula-coupling approach. *Mon. Weather Rev.* **144** 4737–4750.
- BOUROTTE, M., ALLARD, D. and PORCU, E. (2016). A flexible class of non-separable cross-covariance functions for multivariate space–time data. *Spat. Stat.* **18** 125–146. [MR3573273](#)
- BRISSON, N., GARY, C., JUSTES, E., ROCHE, R., MARY, B., RIPOCHE, D., ZIMMER, D., SIERRA, J., BERTUZZI, P. and BURGER, P. (2003). An overview of the crop model STICS. *Eur. J. Agron.* **18** 309–332.
- BROWN, B. G., KATZ, R. W. and MURPHY, A. H. (1984). Time series models to simulate and forecast wind speed and wind power. *J. Clim. Appl. Meteorol.* **23** 1184–1195.
- CLARK, M., GANGOPADHYAY, S., HAY, L., RAJAGOPALAN, B. and WILBY, R. (2004). The Schaake shuffle: A method for reconstructing space–time variability in forecasted precipitation and temperature fields. *J. Hydrometeorol.* **5** 243–262.

- CONSTANTINESCU, E. M. and ANITESCU, M. (2013). Physics-based covariance models for Gaussian processes with multiple outputs. *Int. J. Uncertain. Quantif.* **3** 47–71. [MR3044910](#)
- CONSTANTINESCU, E., ZAVALA, V., ROCKLIN, M., LEE, S. and ANITESCU, M. (2011). A computational framework for uncertainty quantification and stochastic optimization in unit commitment with wind power generation. *IEEE Trans. Power Syst.* **26** 431–441.
- COWLES, M. K., ZIMMERMAN, D. L., CHRIST, A. and MCGINNIS, D. L. (2002). Combining snow water equivalent data from multiple sources to estimate spatio-temporal trends and compare measurement systems. *J. Agric. Biol. Environ. Stat.* **7** 536–557.
- CRESSIE, N. and WIKLE, C. K. (2011). *Statistics for Spatio-Temporal Data*. Wiley, Hoboken, NJ. [MR2848400](#)
- DAWID, A. P. and SEBASTIANI, P. (1999). Coherent dispersion criteria for optimal experimental design. *Ann. Statist.* **27** 65–81.
- FANSHAWE, T. R. and DIGGLE, P. J. (2012). Bivariate geostatistical modelling: A review and an application to spatial variation in Radon concentrations. *Environ. Ecol. Stat.* **19** 139–160. [MR2946168](#)
- FELDMANN, K., SCHEUERER, M. and THORARINSDOTTIR, T. L. (2015). Spatial postprocessing of ensemble forecasts for temperature using nonhomogeneous Gaussian regression. *Monthly Weather Rev.* **143** 955–971.
- FUENTES, M. and RAFTERY, A. E. (2005). Model evaluation and spatial interpolation by Bayesian combination of observations with outputs from numerical models. *Biometrics* **61** 36–45.
- FUENTES, M., CHEN, L., DAVIS, J. M. and LACKMANN, G. M. (2005). Modeling and predicting complex space–time structures and patterns of coastal wind fields. *Environmetrics* **16** 449–464.
- GEL, Y., RAFTERY, A. E. and GNEITING, T. (2004). Calibrated probabilistic mesoscale weather field forecasting: The geostatistical output perturbation method. *J. Amer. Statist. Assoc.* **99** 575–583.
- GENTON, M. G. and KLEIBER, W. (2015). Cross-covariance functions for multivariate geostatistics. *Statist. Sci.* **30** 147–163. [MR3353096](#)
- GLAHN, H. R. and LOWRY, D. A. (1972). The use of model output statistics (MOS) in objective weather forecasting. *J. Appl. Meteorol.* **11** 1203–1211.
- GNEITING, T., RAFTERY, A. E., WESTVELD, A. H. III and GOLDMAN, T. (2005). Calibrated probabilistic forecasting using ensemble model output statistics and minimum CRPS estimation. *Mon. Weather Rev.* **133** 1098–1118.
- GNEITING, T., LARSON, K., WESTRICK, K., GENTON, M. G. and ALDRICH, E. (2006). Calibrated probabilistic forecasting at the stateline wind energy center: The regime-switching space–time method. *J. Amer. Statist. Assoc.* **101** 968–979.
- GNEITING, T., STANBERRY, L. I., GRIMIT, E. P., HELD, L. and JOHNSON, N. A. (2008). Assessing probabilistic forecasts of multivariate quantities, with an application to ensemble predictions of surface winds. *TEST* **17** 211–235.
- HAMILL, T. M. (2001). Interpretation of rank histograms for verifying ensemble forecasts. *Mon. Weather Rev.* **129** 550–560.
- HERING, A. S. and GENTON, M. G. (2010). Powering up with space–time wind forecasting. *J. Amer. Statist. Assoc.* **105** 92–104.
- KANG, E. L., CRESSIE, N. and SAIN, S. R. (2012). Combining outputs from the North American regional climate change assessment program by using a Bayesian hierarchical model. *J. R. Stat. Soc. Ser. C. Appl. Stat.* **61** 291–313.
- KAZOR, K. and HERING, A. S. (2015). The role of regimes in short-term wind speed forecasting at multiple wind farms. *Stat* **4** 271–290. [MR3429323](#)
- LERCH, S. and THORARINSDOTTIR, T. L. (2013). Comparison of non-homogeneous regression models for probabilistic wind speed forecasting. *Tellus A* **65** 21206.

- LI, N., UCKUN, C., CONSTANTINESCU, E., BIRGE, J., HEDMAN, K. and BOTTERUD, A. (2015). Flexible operation of batteries in power system scheduling with renewable energy. *IEEE Trans. Sustain. Energy* **7** 685–696.
- PALMER, T. (2014). More reliable forecasts with less precise computations: A fast-track route to cloud-resolved weather and climate simulators? *Philos. Trans. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci.* **372** 20130391.
- PAPAVASILIOU, A., OREN, S. S. and ROUNTREE, B. (2015). Applying high performance computing to transmission-constrained stochastic unit commitment for renewable energy integration. *IEEE Trans. Power Syst.* **30** 1109–1120.
- PINSON, P. (2013). Wind energy: Forecasting challenges for its operational management. *Statist. Sci.* **28** 564–585.
- PINSON, P. and GIRARD, R. (2012). Evaluating the quality of scenarios of short-term wind power generation. *Appl. Energy* **96** 12–20.
- PINSON, P. and TASTU, J. (2013). Discrimination ability of the energy score. Technical University of Denmark. Tech. rep.
- PINSON, P., CHRISTENSEN, L. E. A., MADSEN, H., SORENSEN, P. E., DONOVAN, M. H. and JENSEN, L. E. (2008). Regime-switching modelling of the fluctuations of offshore wind generation. *J. Wind Eng. Ind. Aerodyn.* **96** 2327–2347.
- PINSON, P., MADSEN, H., NIELSEN, H., PAPAETHYMIU, G. and KLÖCKL, B. (2009). From probabilistic forecasts to statistical scenarios of short-term wind power production. *Wind Energy* **12** 51–62.
- RAFTERY, A. E., GNEITING, T., BALABDAOUI, F. and POLAKOWSKI, M. (2005). Using Bayesian model averaging to calibrate forecast ensembles. *Mon. Weather Rev.* **133** 1155–1174.
- ROYLE, J. A. and BERLINER, L. M. (1999). A hierarchical approach to multivariate spatial modeling and prediction. *J. Agric. Biol. Environ. Stat.* **4** 29–56. [MR1812239](#)
- ROYLE, J., BERLINER, L., WIKLE, C. and MILLIFF, R. (1999). A hierarchical spatial model for constructing wind fields from scatterometer data in the Labrador Sea. In *Case Studies in Bayesian Statistics* 367–382. Springer, Berlin.
- SCHEFZIK, R., THORARINSDOTTIR, T. L. and GNEITING, T. (2013). Uncertainty quantification in complex simulation models using ensemble copula coupling. *Statist. Sci.* **28** 616–640.
- SCHUEERER, M. and HAMILL, T. M. (2015). Variogram-based proper scoring rules for probabilistic forecasts of multivariate quantities. *Mon. Weather Rev.* **143** 1321–1334.
- SCHUEERER, M. and MÖLLER, D. (2015). Probabilistic wind speed forecasting on a grid based on ensemble model output statistics. *Ann. Appl. Stat.* **9** 1328–1349. [MR3418725](#)
- SCHUHEN, N., THORARINSDOTTIR, T. L. and GNEITING, T. (2012). Ensemble model output statistics for wind vectors. *Mon. Weather Rev.* **140** 3204–3219.
- SHUMWAY, R. and STOFFER, D. (2010). *Time Series Analysis and Its Applications: With R Examples*. Springer Science & Business Media.
- SKAMAROCK, W., KLEMP, J., DUDHIA, J., GILL, D., BARKER, D., DUDA, M., HUANG, X.-Y., WANG, W. and POWERS, J. (2008). A description of the Advanced Research WRF Version 3 Tech. Rep. Tech Notes-475+ STR, NCAR.
- SLOUGHTER, J. M. L., GNEITING, T. and RAFTERY, A. E. (2010). Probabilistic wind speed forecasting using ensembles and Bayesian model averaging. *J. Amer. Statist. Assoc.* **105** 25–35.
- SLOUGHTER, J. M. L., GNEITING, T. and RAFTERY, A. E. (2013). Probabilistic wind vector forecasting using ensembles and Bayesian model averaging. *Mon. Weather Rev.* **141** 2107–2119.
- SMITH, L. A. and HANSEN, J. A. (2004). Extending the limits of ensemble forecast verification with the minimum spanning tree. *Mon. Weather Rev.* **132** 1522–1528.
- STEIN, M. L. (1999). *Interpolation of Spatial Data: Some Theory for Kriging*. Springer, New York. [MR1697409](#)
- STEIN, M., CHEN, J. and ANITESCU, M. (2012). Difference filter preconditioning for large covariance matrices. *SIAM J. Matrix Anal. Appl.* **33** 52–72.

- THORARINSDOTTIR, T. L. and GNEITING, T. (2010). Probabilistic forecasts of wind speed: Ensemble model output statistics by using heteroscedastic censored regression. *J. Roy. Statist. Soc. Ser. A* **173** 371–388.
- THORARINSDOTTIR, T. L. and JOHNSON, M. S. (2012). Probabilistic wind gust forecasting using non-homogeneous Gaussian regression. *Mon. Weather Rev.* **140** 889–897.
- THORARINSDOTTIR, T. L., SCHEUERER, M. and HEINZ, C. (2016). Assessing the calibration of high-dimensional ensemble forecasts using rank histograms. *J. Comput. Graph. Statist.* **25** 105–122.
- WILKS, D. S. (2015). Multivariate ensemble Model Output Statistics using empirical copulas. *Q. J. R. Meteorol. Soc.* **141** 945–952.

J. BESSAC
ARGONNE NATIONAL LABORATORY
9700 SOUTH CASS AVENUE
LEMONT, ILLINOIS 60439-4844
USA
E-MAIL: jbessac@anl.gov
emconsta@mcs.anl.gov

E. CONSTANTINESCU
ARGONNE NATIONAL LABORATORY
9700 SOUTH CASS AVENUE
LEMONT, ILLINOIS 60439-4844
USA
AND
UNIVERSITY OF CHICAGO
COMPUTATION INSTITUTE
5735 SOUTH ELLIS AVENUE
CHICAGO, ILLINOIS 60637
USA
E-MAIL: emconsta@mcs.anl.gov

M. ANITESCU
ARGONNE NATIONAL LABORATORY
9700 SOUTH CASS AVENUE
LEMONT, ILLINOIS 60439-4844
USA
AND
UNIVERSITY OF CHICAGO
DEPARTMENT OF STATISTICS
GEORGE HERBERT JONES LABORATORY
5747 S. ELLIS AVENUE
CHICAGO, ILLINOIS 60637
USA
E-MAIL: anitescu@mcs.anl.gov