SPECIAL ISSUE

A NON-GAUSSIAN SPATIO-TEMPORAL MODEL FOR DAILY WIND SPEEDS BASED ON A MULTI-VARIATE SKEW-*T* DISTRIBUTION

FELIPE TAGLE,[a] STEFANO CASTRUCCIO,[a] PAOLA CRIPPA[b] AND MARC G. GENTON[c*]

[a] *Department of Applied and Computational Mathematics and Statistics, University of Notre Dame, Notre Dame, IN, USA*
[b] *Department of Civil and Environmental Engineering and Earth Sciences, University of Notre Dame, Notre Dame, IN, USA*
[c] *Statistics Program, King Abdullah University of Science and Technology, Thuwal, Saudi Arabia*

Facing increasing domestic energy consumption from population growth and industrialization, Saudi Arabia is aiming to reduce its reliance on fossil fuels and to broaden its energy mix by expanding investment in renewable energy sources, including wind energy. A preliminary task in the development of wind energy infrastructure is the assessment of wind energy potential, a key aspect of which is the characterization of its spatio-temporal behavior. In this study we examine the impact of internal climate variability on seasonal wind power density fluctuations over Saudi Arabia using 30 simulations from the Large Ensemble Project (LENS) developed at the National Center for Atmospheric Research. Furthermore, a spatio-temporal model for daily wind speed is proposed with neighbor-based cross-temporal dependence, and a multi-variate skew-*t* distribution to capture the spatial patterns of higher-order moments. The model can be used to generate synthetic time series over the entire spatial domain that adequately reproduce the internal variability of the LENS dataset.

1. INTRODUCTION

Wind energy has become an important component of energy portfolios for many developed and developing nations worldwide. This trend is largely driven by technological advances that enable more efficient harnessing of available wind energy and its integration into energy transmission networks, but is also due to growing recognition of the importance of renewable energy sources in climate change mitigation strategies (Bruckner *et al.*, 2014; Zhu and Genton, 2012, and references therein). For instance, Denmark relies on wind energy for nearly 40% of its domestic energy consumption, and this number is expected to rise to 50% by 2020 (DEA, 2016). In 2015, China added 33 GW of new installed capacity, to reach a total capacity of 148 GW, representing an increase of 29% (WWEA, 2016). These and other examples provide convincing evidence that wind energy offers a viable alternative to traditional fossil fuel-based energy sources, and are paving the way for other nations to exploit this renewable and clean energy resource.

Saudia Arabia, with its vast oil reserves, has a long tradition of relying solely on fossil fuels for its energy needs. However, faced with rising energy demands due to population growth and industrialization, it is actively seeking to diversify its energy mix by expanding its renewable energy portfolio to 54 GW by 2032, of which 9 GW will come from wind power (KA-CARE, 2012). Due to the inherent variability and limited predictability of the wind resource, assessing its spatial and temporal characteristics is a critical step in the development of

---

* Correspondence to: Marc G. Genton, Statistics Program, King Abdullah University of Science and Technology, Thuwal 23955-6900, Saudi Arabia. E-mail: marc.genton@kaust.edu.sa

wind energy infrastructure. Several studies have sought to quantify the wind potential over Saudi Arabia (e.g., Rehman and Ahmad, 2004; Rehman *et al.*, 2007; Shaahid *et al.*, 2014), however, these have mostly focused on just a few locations, largely due to the country's sparse observational network with varying record lengths. A recent study by Yip *et al.* (2016) used a gridded reanalysis dataset with a multi-decadal record period to provide a more comprehensive assessment, evaluating the abundance of the wind resource over the Arabian Peninsula. The authors also investigated its variability and intermittency, as these pose the greatest challenges to the integration of wind energy into existing power grid systems. In this study, as part of an ongoing collaborative effort with the National Center for Atmospheric Research (NCAR), we focus on the spatio-temporal characteristics of this resource over Saudi Arabia. Wind speeds fluctuate over a wide range of frequencies, with those of the order of days governed by general weather patterns, while shorter frequencies are driven by local effects and turbulence (Pinson, 2013). Our focus is on very low frequencies, where long-term wind trends and the effects of human activities may potentially play a role. In particular, this is the first study that investigates the sensitivity of wind energy potential to internal climate variability, that is, the variability that is intrinsic to the climate system due to the complex interactions between its components at various temporal and spatial scales. Only recently has this source of variability been properly identified in the geophysical community (i.e., Deser *et al.*, 2012, 2014), and its impact is just beginning to be understood. Because of the decadal timescales at which this source of variability operates, the characterization of its impact usually relies on simulations from Earth system models (ESMs), which, in addition to solving the atmospheric governing equations, include processes from other components of the climate system, namely, land, ocean and sea ice; and so are well-suited to represent the complex interactions that occur at such timescales. In this work we use a collection of simulations based on the community ESM (CESM), developed at NCAR under the Large Ensemble Project (LENS) that are specifically designed to isolate the effects of internal climate variability. These simulations and ESM simulations in general, however, are computationally expensive to run, and many simulations are necessary to properly quantify the effects of natural variability. An alternative to simulating atmospheric processes using these models is provided by stochastic weather generators (Wilks and Wilby, 1999; Ailliot *et al.*, 2015a), which are statistical models designed to generate synthetic sequences of meteorological variables whose statistical properties closely resemble those of the observational datasets on which they are trained. We develop one such model designed to capture the natural variability of daily wind speed in Saudi Arabia, that may be used, among other applications, to assess future wind energy potential under different climate scenarios.

Stochastic weather generators are widely used in applications ranging from agricultural models to climate impact studies. For instance, they are commonly used in the energy industry to generate short-temporal frequency simulations of wind data to assist utilities in grid integration studies, where these and other components of an electrical grid are simulated to assess general performance measures and determine best practices (Archer *et al.*, 2017). In Ailliot and Monbet (2012), a model of this kind was used to generate wind time series at different meteorological stations near potential wind farm sites to assess various quantities related to wind power production. Most of these stochastic wind generators have focused on single-site models, and typically use Markov chains to describe the temporal dependence on different wind regimes (Ailliot *et al.*, 2006). Recently, focus has shifted towards a multisite framework, which poses the challenge of having to explicitly account for the complex cross-correlation patterns among neighboring sites (Bessac *et al.*, 2015; Hering *et al.*, 2015; Bessac *et al.*, 2016). Thus, the number of sites under consideration is usually limited, because the quality of fit quickly deteriorates as the quantity increases (Hering *et al.*, 2015). The model presented here, however, with the approximately 1° horizontal resolution of the LENS dataset, is applied to a 149-point spatial domain which, together with a temporal coverage spanning more than 80 years, results in a spatio-temporal domain of nearly 5 million data points. To emphasize the different framework of this statistical model, consistently with Jeong *et al.* (2018, 2019), we denote this as a *stochastic generator*, and in particular in this application, a *stochastic wind generator*. While stochastic generators rely on a stochastic approximation of climate model output, they are also fundamentally different from an emulator in that they are not used to explore the input parameter space and perform sensitivity analysis.

Environmental data often exhibit departures from Gaussianity, such as skewness and heavy tails. Traditionally, to continue exploiting the appealing properties of the normal distribution and the well-developed theory of Gaussian processes, a transformation is usually applied to the data; for example, a square root transformation (Gneiting,

2002), a power transformation (Ailliot *et al.*, 2015b; Bessac *et al.*, 2016) or a Gaussian copula (Hering *et al.*, 2015). Only recently have there been studies that exploit the flexibility of skew-elliptical distributions, of which the skew-*t* (Azzalini and Capitanio, 2003) and the skew-normal (Azzalini, 2005) distributions are special cases that directly address the skewness and excess kurtosis that are commonly found in wind data. For example, Hering and Genton (2010) developed a short-term wind speed forecasting model using a bivariate skew-*t* distribution, while Flecher *et al.* (2010) used a closed skew-normal distribution (González-Farías *et al.*, 2004) as part of a multi-variate weather generator. This family of distributions represents an extension of the normal model and retains several of its convenient properties, such as closure under conditioning and marginalization (see Genton (2004) or Azzalini and Capitanio (2014) for an overview of the theory and applications). In this article we propose a stochastic wind generator that leverages the flexibility of the skew-*t* distribution to capture the non-Gaussian features that are common to wind-speed time series. In particular, we consider a vector-autoregressive process to capture the temporal auto- and cross-correlations, coupled with independent realizations from a multi-variate skew-*t* distribution as a model for the spatial residuals. Given that our interest lies in generating spatial replicates at locations and support that coincide with those of the training dataset, we consider the skew-*t* distribution only in its multi-variate form. An extension to a spatial process (e.g., Beranger *et al.*, 2017; Morris *et al.*, 2017) would have been necessary had the scope of this study involved an application of spatial interpolation, but it is not contemplated here.

The remainder of the article is organized as follows. Section 2 describes the daily wind speed data used for the training and validation of the spatio-temporal model described in Section 3. Section 4 provides an assessment of the temporal variability of the wind energy resource in Saudi Arabia, and Section 5 offers a discussion and conclusion.

2. WIND DATA

We consider daily wind speed in Saudi Arabia from the publicly available LENS dataset developed at NCAR (Kay *et al.*, 2015). The ensemble consists of 30 fully-coupled simulations, based on CESM version 1, with the community atmosphere model (CAM), version 5, run at approximately 1° horizontal resolution in all model components. The simulations span from the year 1920 to 2100, with radiative forcing following the CMIP5 protocol; namely, historical forcing from 1920 to 2005, and the representative concentration pathway 8.5 (RCP8.5) from 2006 to 2100. Each simulation represents a unique climate trajectory, generated by introducing small round-off level differences into their initial atmospheric conditions. For this study we use the historical segment (1920–2005) of 30 ensemble members over Saudi Arabia, bounded roughly by 15–32°N and 35–55°E, which at the noted horizontal resolution corresponds to $N = 149$ points, in the temporal dimension $T = 31{,}390$ points (the model does not account for leap years), resulting in nearly 5 million data points.

To illustrate the diverse wind regimes across Saudi Arabia, we examine the various shapes of the annual cycle, as represented by the LENS ensemble mean. Figure 1 displays the annual cycles for five locations that broadly capture the variety of regimes. The seasonal patterns can be interpreted with the aid of the mean monthly wind fields (Figure S1 in Appendix S1, Supporting Information), constructed from the lowest level zonal and meridional winds (UBOT and VBOT respectively) obtained from the LENS database. The Arabian Peninsula lies within the trade-wind belt (Hadley cell) of the Northern Hemisphere. During winter, westerlies from the Mediterranean travel southward towards the Persian Gulf (Shamal trade winds), and turn south and southwest through the Rub al-Khali Desert toward Yemen. However beginning in spring, the northerly winds intensify, and the onset of the monsoon circulation from the Indian Ocean induces a southeasterly flow along the southeastern part of the peninsula and the Arabian Sea, into the Rub al-Khali Desert. As a result, the mild easterly flow over Point 3, located east of the Asir Mountains, becomes strongly westerly during the summer, explaining the maximum daily wind speed during this season. Moreover, the low speeds in March and September correspond to the transitions between easterly and westerly flow, characterized by strong wind convergence. Despite their distance, Points 2 and 4 display a similar pattern, with dual peaks arising in early spring and the middle of summer, largely due to the intensification of the northerly winds during the monsoon season. The temporal pattern over Point 5, close to the eastern coast, most clearly reflects the second intensification of the Shamal trade winds during May and June. The mountainous region
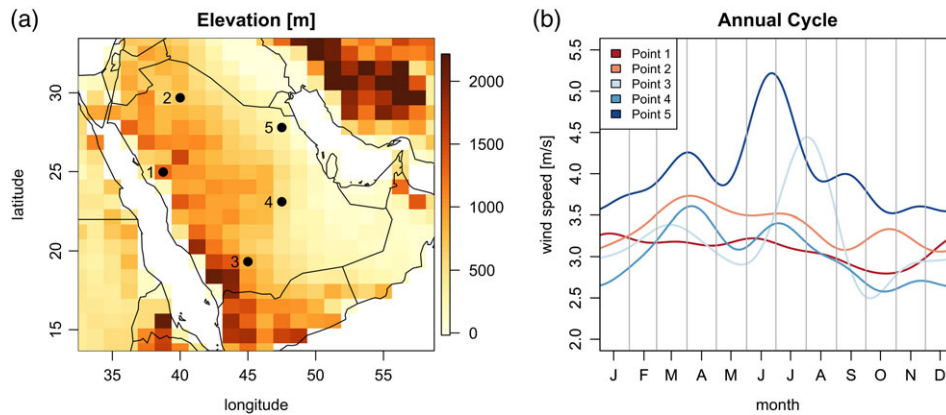
Figure 1. (a) Selection of points over Saudi Arabia with colors showing the elevation in meters, and their respective (b) daily wind speed annual cycles, computed as the mean annual cycle from the LENS. Wind speed data is based on the 1920–2005 output [Color figure can be viewed at wileyonlinelibrary.com]

on the eastern coast of the peninsula, as represented by Point 1, displays a dampened annual cycle, reflecting the effect of the complex orography.

3. MODEL DESCRIPTION

We describe the proposed spatio-temporal model, and discuss inferential aspects, highlighting several difficulties and the approach undertaken to overcome them. Finally, we validate the model along several metrics that focus on low-frequency aspects of the wind-speed time series.

3.1. Spatio-temporal Model

Let $\mathbf{W}_{t,r} = (W_{1,t,r}, \ldots, W_{N,t,r})^\top$ denote the N-vector of daily wind speeds over the domain at day t, in realization r. We first consider a standardization of the form $(W_{i,t,r} - \mu_{i,t})/\sigma_{i,t}$, where $\mu_{i,t}$ refers to a seasonal effect and $\sigma_{i,t}$ the seasonal fluctuation in the standard deviation of the residuals, both indexed by time t and location i. The former is estimated by regressing the mean of each calendar day of the time series of daily wind speed at each gridpoint on a small set of harmonics, ranging from frequencies of 1–5 cycles per year. Similarly, the latter is obtained by fitting the same range of harmonics to the standard deviation of each calendar day of the residual time series at each gridpoint. We make the assumption that both standardization terms are common to all ensemble members, to be consistent with the design of the LENS. Despite the realizations having a time span of over 80 years, we did not find evidence of temporal non-stationarity in the estimates (see the diagnostics in Figures S2–S5). For ease of notation, we assume henceforth that $\mathbf{W}_{t,r}$ and its component terms have undergone said standardization. For the residuals, we propose a model of the form

$$\mathbf{W}_{t,r} = \mathbf{A}_1 \mathbf{W}_{t-1,r} + \mathbf{A}_2 \mathbf{W}_{t-2,r} + \boldsymbol{\varepsilon}_{t,r} \quad t = 3, \ldots, T, \tag{1}$$

where $\mathbf{A}_k$, $k = 1, 2$, are $N \times N$ coefficient matrices and $\boldsymbol{\varepsilon}_{t,r}$ corresponds to the time t innovations of the rth realization, for which $\mathrm{E}(\boldsymbol{\varepsilon}_{t,r}) = \mathbf{0}$, and $\mathrm{E}(\boldsymbol{\varepsilon}_{t,r} \boldsymbol{\varepsilon}_{t,r}^\top) = \boldsymbol{\Sigma}_\varepsilon$. We assume that the coefficient matrices are identical across the ensemble, while the distinct trajectories of each realization are the result of i.i.d. innovations. This is in agreement with the construction of the LENS and atmospheric flow more generally; the initial condition memory is lost within a few weeks, after which each ensemble member evolves chaotically as if driven by random atmospheric fluctuations (Kay *et al.*, 2015). To assess whether fitting the model to a few realizations is sufficient to capture the

variability across the LENS, the model in (1) is fit to the time series belonging to three members under the above assumption of i.i.d. realizations. These three are chosen arbitrarily among the 30 realizations, replicating a scenario where only three members are available. However, the results are robust to the number of members of the training set (see Figure S6 and its associated section in Appendix S1).

The temporal dependence in wind time series at both hourly and daily time scales is usually well represented in the univariate case by an autoregressive model of order 2 (e.g., Brown *et al.*, 1984; Haslett and Raftery, 1989; Ailliot *et al.*, 2006) and in the multi-site setting, by its multi-variate counterpart, the vector-autoregressive (VAR) process of order 2 (e.g., Hering *et al.*, 2015; Bessac *et al.*, 2016). An obvious difficulty that emerges in the latter context is the estimation of the autoregressive matrices as the dimension of the multi-variate process grows. Parametric models for off-diagonal elements based on the distance between sites or shrinkage methods for reducing elements to zero have been proposed to address this issue (Monbet and Ailliot, 2017). Schweinberger *et al.* (2017) investigated the $p \gg N$ case, where p is the number of parameters and N is the number of observations, under the assumption that informative spatial structure is available. Here the gridded nature of the spatial domain naturally suggests a nearest-neighbor specification. Consider the following notation: $\mathbf{W} = (\mathbf{W}_3, \ldots, \mathbf{W}_T)$, $\mathbf{B} = (\mathbf{A}_1, \mathbf{A}_2)$, $\mathbf{Z}_t = \mathrm{vec}(\mathbf{W}_t, \mathbf{W}_{t-1})$, $\mathbf{Z} = (\mathbf{Z}_2, \ldots, \mathbf{Z}_{T-1})$ and $\mathbf{U} = (\boldsymbol{\varepsilon}_3, \ldots, \boldsymbol{\varepsilon}_T)$. Then the VAR(2) process in (1) can be expressed compactly as

$$\mathbf{W} = \mathbf{BZ} + \mathbf{U}$$

or equivalently as $\mathbf{W}^* = (\mathbf{Z}^\top \otimes \mathbf{I}_N)\boldsymbol{\beta} + \mathbf{U}^*$, which results from applying the vec operator to both sides of the equation, so that $\mathbf{W}^* = \mathrm{vec}(\mathbf{W})$, $\boldsymbol{\beta} = \mathrm{vec}(\mathbf{B})$ and $\mathbf{U}^* = \mathrm{vec}(\mathbf{U})$. The nearest-neighbor specification amounts to imposing zero restrictions on the non-neighbor coefficients, which correspond to linear constraints of the form $\boldsymbol{\beta} = \mathbf{R}\boldsymbol{\gamma}$, where $\mathbf{R}$ is a $(2N^2) \times M$ matrix of rank M that encodes the $A_{i,j,k} = 0$ restrictions, for $i, j = 1, \ldots, N$, $k = 1, 2$. Here, $\boldsymbol{\gamma}$ is an unrestricted $M \times 1$ vector of unknown parameters, representing the non-zero elements of $\mathbf{B}$ and M the number of such elements; see Chapter 5 in Lütkepohl (2005) for details. The generalized least squares (GLS) estimator of $\boldsymbol{\gamma}$ is given by

$$\hat{\boldsymbol{\gamma}} = [\mathbf{R}^\top(\mathbf{ZZ}^\top \otimes \boldsymbol{\Sigma}_{\boldsymbol{\varepsilon}}^{-1})\mathbf{R}]^{-1}\mathbf{R}^\top(\mathbf{Z} \otimes \boldsymbol{\Sigma}_{\boldsymbol{\varepsilon}}^{-1}), \tag{2}$$

where $\boldsymbol{\Sigma}_{\boldsymbol{\varepsilon}}$ denotes the covariance matrix of $\boldsymbol{\varepsilon}_t$. The residuals arising from the GLS estimation of $\boldsymbol{\gamma}$, among several nearest-neighbor configurations, did not adequately capture the autoregressive structure, in comparison with the traditional OLS approach that replaces $\boldsymbol{\Sigma}_{\mathbf{U}^*}$ in (2) with $\mathbf{I}_N$. The GLS approach attempts to account for the second-order structure in the residuals through the autoregressive parameters, which is undesirable in this case as we will model that structure in a subsequent stage. We therefore opted to proceed with the OLS approach. The stability condition of the VAR(2) process is checked in the usual manner, verifying that all of the eigenvalues of the autoregressive matrix of the compact form have modulus less than one.

We find that a first-order stencil neighborhood scheme (North, West, South, East) for $\mathbf{A}_1$ and a diagonal form for $\mathbf{A}_2$ adequately represent the temporal and cross-temporal structure (Figure S7). Figure 2 displays the estimates for the leading and diagonal terms in $\mathbf{A}_1$ and $\mathbf{A}_2$ respectively. Symbols in Figure 2(a) denote neighbor to which the value of the leading term corresponds. The estimates in $\mathbf{A}_1$ are consistently positive across the entire domain, with values exceeding 0.7 along the boundaries. The leading terms in the northern part of the country are generally from the West neighbor, whereas those along the center belong to the North neighbor, in agreement with the predominant wind patterns in Figure S1. Those of $\mathbf{A}_2$, however, are consistently negative, and even more so in the Eastern Province.

Given a domain of this size, with its wide range of wind regimes, the assumption of stationarity for a parametric form of the associated correlation matrix would be inappropriate. Several approaches have been advocated in the statistical literature to deal with non-stationary spatial covariance structures, beginning with the seminal work of Sampson and Guttorp (1992) and their spatial deformation approach. More recently, other approaches based on spatially weighted combinations of stationary spatial covariance functions (Fuentes, 2001; Fuentes and Smith,
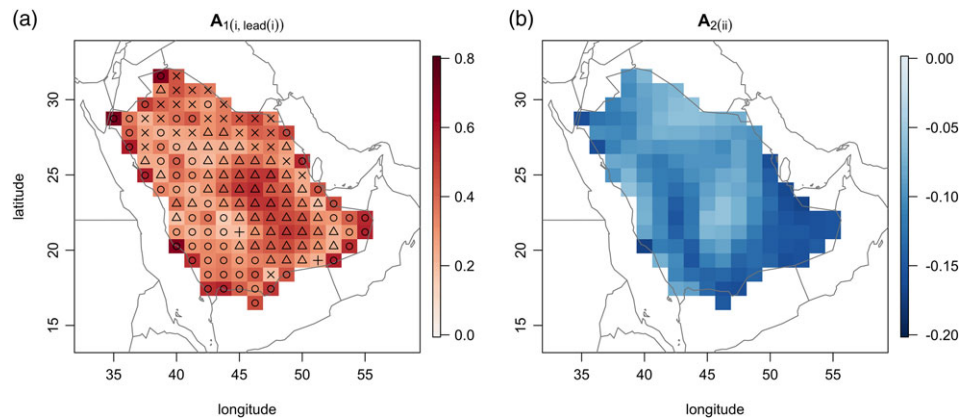
Figure 2. Estimates for the (a) leading terms in $\mathbf{A}_1$ and (b) the diagonal elements of $\mathbf{A}_2$; all statistically significant at 1%, withstanding an adjustment for multiplicity using the false discovery rate (Benjamini and Hochberg, 1995). Symbols in (a) indicate the neighbor corresponding to the leading term: circle, itself; triangle point-up, North; triangle point-down, South; slanted cross, West; cross, East [Color figure can be viewed at wileyonlinelibrary.com]
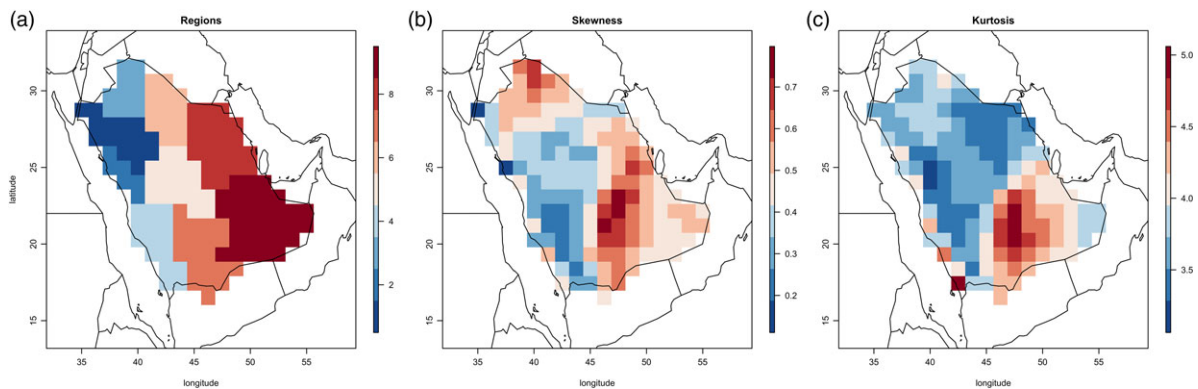

Figure 3. (a) Nine regions used in the fitting of the innovation $\boldsymbol{\varepsilon}_t$, defined using Ward's method, an agglomerative method that minimizes the sum of squared Euclidean distances between the points and the centroids of their respective cluster. Here each point corresponds to the N-dimensional vector of rescaled VAR(2) residuals. The (b) skewness and (c) kurtosis coefficients have been added as visual guides to help explain the regional shapes [Color figure can be viewed at wileyonlinelibrary.com]

2001) and process convolutions (Higdon, 1998; Paciorek and Schervish, 2006) have received particular attention (see Sampson (2010) for a review). Here we partition the spatial domain into regions where the assumption of stationarity is plausible; that is, we divide the vector $\boldsymbol{\varepsilon}_t$ (henceforth we drop the subscript r for convenience) into N_c subvectors, $\boldsymbol{\varepsilon}_t = (\boldsymbol{\varepsilon}_{1,t}^\top, \ldots, \boldsymbol{\varepsilon}_{N_c,t}^\top)^\top$, where each $\boldsymbol{\varepsilon}_{c,t}$, composed of d_c number of gridpoints, is designed to capture regional features. We use Ward's hierarchical clustering method to partition the domain, as it tends to produce clusters of approximately equal size (Everitt *et al.*, 2011). Nine regions (Figure 3) result in cluster sizes that are suitable for parameter estimation. The clustering of spatial locations has a long history in environmental applications (Wilks, 2011). Recently, Lorente-Plazas *et al.* (2015) used a combination of hierarchical and non-hierarchical techniques to cluster wind station data collected in the Iberian Peninsula. In the Markov regime-switching literature, clustering is often performed along the temporal dimension to identify different weather regimes (Kazor and Hering, 2015).

Each $\boldsymbol{\varepsilon}_{c,t}$ is assumed to follow an independent multi-variate skew-t distribution, based on the formulation of Azzalini and Capitanio (2003) which has an analogous derivation to the familiar Student-t distribution, as opposed
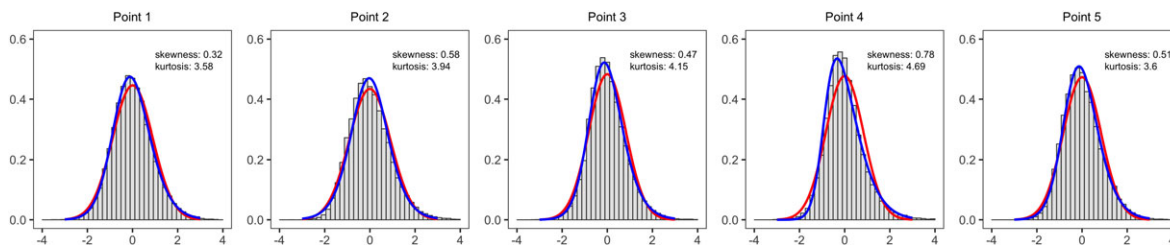
Figure 4. Histograms of VAR(2) time series of the five points denoted in Figure 1. The marginal skew-*t* distribution, obtained by marginalizing its multi-variate parent, is denoted in blue. A reference Gaussian density has been added and denoted in red. The sample skewness and kurtosis are noted in each panel [Color figure can be viewed at wileyonlinelibrary.com]

to others such as that of Branco and Dey (2001); see Azzalini *et al.* (2016) for a comparison of these formulations. A *d*-dimensional random variable **Z** follows a multi-variate skew-*t* distribution, denoted by $ST_d(\xi, \Omega, \alpha, v)$, with ξ being a location *d*-vector, Ω a $d \times d$ scale matrix, α a *d*-vector skewness parameter, and *v* being the degrees of freedom, if it has a probability density of the form

$$f(\mathbf{z}) = 2t_d(\mathbf{z} - \xi; \Omega, v)T\left(\alpha^\top \omega^{-1}(\mathbf{z} - \xi)\sqrt{\frac{v+d}{v+Q(\mathbf{z})}}; v+d\right), \tag{3}$$

with $Q(\mathbf{z}) = (\mathbf{z} - \xi)^\top \Omega(\mathbf{z} - \xi)$, $\omega = \mathrm{diag}(\omega_1, \dots, \omega_d) > 0$, $\Omega = \omega\bar{\Omega}\omega$, $\bar{\Omega}$ a $d \times d$ correlation matrix, $t_d(\cdot; \Omega, v)$ is the probability density of the *d*-dimensional Student-*t* distribution, and $T(\cdot; v)$ denotes the cumulative distribution function of a univariate Student-*t* distribution with *v* degrees of freedom.

Empirical justification for the model in (1) with non-Gaussian innovations is displayed in Figure 3(b, c), where the pointwise estimates of skewness and kurtosis for the VAR(2) residuals are presented. Clear departures from zero and three, arise respectively, in the southwest part of the country. Furthermore, in Figure 4 the marginal probability densities of the fitted multi-variate skew-*t* distributions for the five points denoted in Figure 1 are contrasted with their respective histograms and a reference Gaussian distribution. The sample skewness is consistently positive across the various points, as is the excess kurtosis, in agreement with the values shown in Figure 3(b, c).

3.2. Inference

One of the drawbacks of the parameterization used in (3), called the *direct parameterization* (DP), is that its parameters are not easy to interpret, as there is no simple relationship between them and the moments or cumulants of the distribution (Arellano-Valle and Azzalini, 2013). This is most readily apparent in the location parameter ξ, which bears no resemblance to typical centrality measures such as the mean or median. This is problematic also for inference, as the reductions in the parameter space that usually accompany assumptions such as a mean zero and a variance one do not materialize under this parameterization, as both the mean and the variance are functions of all of the above parameters (see Appendix). A centered parameterization (CP) based on the moments of the distribution was developed by Azzalini (1985) for the DP of the skew-normal distribution; partly to offer a more intuitive parameterization, but more importantly, to respond to the singularity of its associated information matrix at $\alpha = \mathbf{0}$, which does not apply to the skew-*t* distribution. Recently, Arellano-Valle and Azzalini (2013) extended this CP to the skew-*t* distribution. In the univariate setting, with the skew-*t* DP given by $(\xi, \omega^2, \alpha, v)$, the CP corresponds to the four centered moments $(\mu, \sigma^2, \gamma_1, \gamma_2)$, where γ_1 refers to the skewness coefficient and γ_2 is the coefficient of excess kurtosis. An implicit assumption in the CP is that these moments exist, requiring that $v > 4$, which may not hold in practice; and the exact value of *v* imposes constraints on the feasible space of γ_1 and γ_2 (see Figure 1 in Arellano-Valle and Azzalini (2013)). An examination of the marginals revealed that the values of the tuples (γ_1, γ_2) encountered here were well within this feasible set. In the multi-variate case, the first two components of the CP are the familiar mean vector and the covariance matrix, μ and Σ, while the skewness

term, γ_1, consists of the vector of marginal skewness coefficients. The kurtosis term corresponds to the Mardia index of multi-variate kurtosis (see Appendix), denoted by γ_2^M.

Because the singularity at $\alpha = 0$, which characterizes the skew-normal likelihood, is lacking here, maximization of the skew-t likelihood is often the method of choice to conduct parameter inference. However, we find that when maximizing the log-likelihood to each set of residuals, the maximum likelihood estimator (MLE) tends to favor a more precise representation of the correlation structure at the expense of the marginal structure. This is not acceptable for our current application, which places a strong emphasis on the accurate representation of the marginal structure given the sensitivity of wind power density to wind speed. We considered adding a penalization term to the likelihood to constrain the optimization towards an improved marginal fit, but we obtained better results with a method of moments approach based on the above CP, as it afforded us greater control in isolating the fitting of the marginal structure from that of the spatial dependence. More specifically, using the CP gave us the possibility to directly impose the condition that $\mu = 0$ and adopt a parametrized form for the correlation matrix $\hat{\Sigma}$ for each region, thereby considerably reducing the CP parameter space. $\hat{\Sigma}$ is parameterized by $\mathcal{M}(\phi, \kappa)$, a Matérn correlation function with range parameter ϕ and smoothness κ. The latter was fixed at 1.5, and ϕ was fit by OLS. Good agreement was found between the sample estimates of correlation of the VAR(2) residuals and the corresponding value from the OLS-estimated Matérn correlation function. Furthermore, errors from the OLS fit did not meaningfully contaminate the model-simulated seasonal variability in Section 4 (Figure S10). Although the mapping from DP to CP is available in closed-form, the inverse map is not, and so it must be computed numerically. In fact, only the mapping $(\gamma_1, \gamma_2^M) \mapsto (\delta, \nu)$ needs to be specified, since the rest of the DP parameterization follows directly from these parameters and those of the CP parameters (see sections 6.2.2 and 6.2.3 in Azzalini and Capitanio (2014) for details). The `sn` package (Azzalini, 2016) implements an algorithm which seeks the $(d + 1)$-vector $(\delta^\top, \nu)$ that minimizes the ℓ^2-distance between $(\gamma_1(\delta, \nu)^\top, \gamma_2^M(\delta, \nu))$ and its CP counterparts. As is common with many numerical optimization problems, finding a global maximum, or minimum in this case, is not guaranteed. In the few regions where such a situation arose, that is, the ℓ^2-distance does not vanish, the difference between the CP implied by the DP and the original CP was sufficiently small to be of no concern.

The definition of the γ_1 vector naturally suggests the use of sample estimates of the skewness coefficients at each gridpoint. For the multi-variate kurtosis coefficient, γ_2^M, we considered its sample estimate, however, in some cases it yielded unsatisfactory results in terms of the implied skewness of the marginals once mapped to ν. In such cases, we replaced the sample estimate with the value which, together with $\hat{\gamma}_1$, minimizes the ℓ^2-distance between $\gamma_1(\delta, \nu)$ and $\gamma_2(\delta, \nu)$ and its sample counterparts, where here the dependence on (δ, ν) indicates that these are the respective moment vectors after performing the mapping using the `sn` package.

As an indication of the sampling variability, we considered bootstrap standard errors based on 100 bootstrap samples for an arbitrarily chosen ensemble member and found that estimates of these for γ_1 across the domain were contained in $(0.014, 0.035)$ and γ_2 in $(0.035, 0.179)$. Regarding the mapping from CP to DP, we recorded for each cluster the mean-squared-error (MSE) between the sample γ_1 and γ_2 of their constituent gridpoints and the respective values implied by the DP. Note that the errors corresponding to γ_1 serve as a reliable indicator of the accuracy of the mapping between CP and DP, whereas those for γ_2 are contaminated by an inherent characteristic of the DP paramerization, whose value of ν is inherited by all of the marginal skew-t distributions. Therefore the latter errors also reflect the discrepancy between the actual marginal ν and this inherited value. The bootstrap estimates of the means and standard deviations for the MSE of γ_1, across the domain, were contained in $(0.002, 0.071)$ and $(0.01, 0.095)$ respectively. The analogous values for the MSE of γ_2 are $(0.032, 0.49)$ and $(0.02, 0.89)$.

3.3. Model Validation

To assess the performance of the model in simulating realistic realizations of the daily surface wind speed over Saudi Arabia, the model was fit to three ensemble members, and 30 realizations were generated. Several aspects of the simulated time series can be investigated as a basis for validation, such as the distribution of excursion durations above or below some fixed threshold as in Ailliot and Monbet (2012). Had wind direction been part of the analysis, for instance, the joint and marginal distribution of wind speed and direction, as well as their respective
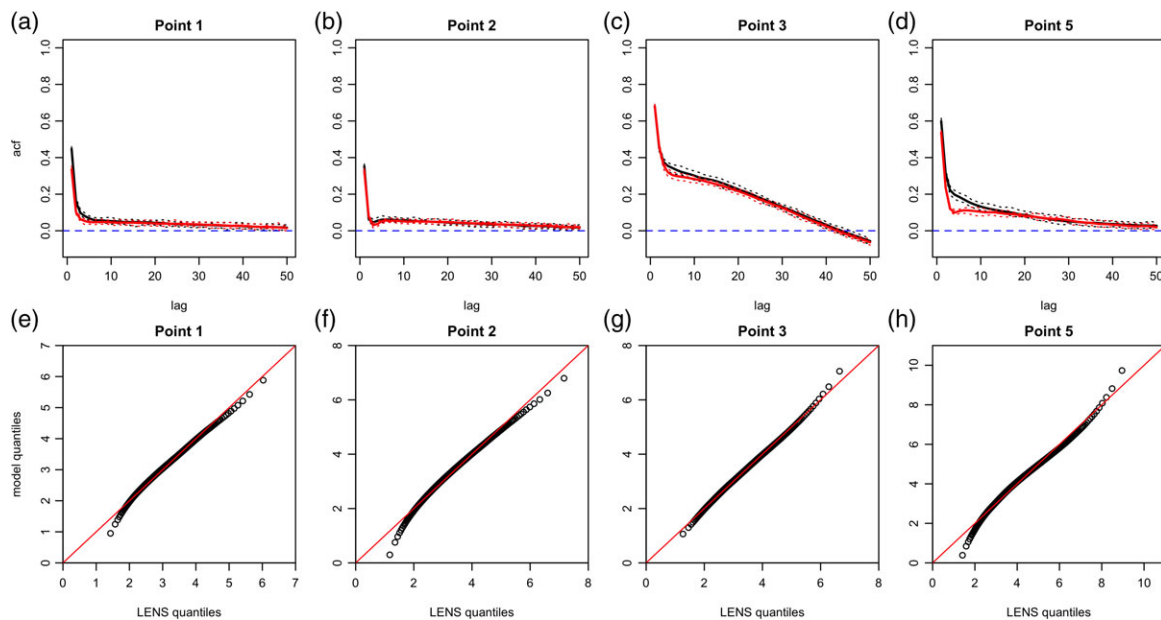
Figure 5. (a)–(d) Average autocorrelation functions of 30 realizations of daily wind speed of (red) the proposed model and (black) the LENS, for Points 1, 2, 3 and 5 indicated in Figure 1. Dotted red and black lines represent the maximums and minimums of the simulated and LENS values respectively, at each lag. A horizontal dashed blue line has been placed at zero. (e)–(h) Mean quantile-quantile plots of the daily surface wind speed for the LENS and the 30 model realizations, for the same points [Color figure can be viewed at wileyonlinelibrary.com]

autocorrelation functions may have been investigated (e.g., Hering *et al.*, 2015). Given our interest in capturing low frequency aspects of wind speed, we focus here on the marginal distributions of daily wind speed and its serial dependence. However, for illustrative purposes, we have included an examination of the distribution of excursion durations in Appendix S1 (Figure S8 and associated section).

Figure 5(a)–(d) shows the mean autocorrelation function (ACF) for the 30 realizations and 30 LENS members, for four of the five representative points displayed in Figure 1, as the ACF for Point 4 is quite similar to that of Point 2, and so is excluded. The realizations adequately reproduce this feature of the LENS data, despite the marked differences among the autocorrelation structures across the various points. In particular, the ACFs in Points 1 and 2 drop sharply beyond the first few lags, whereas Point 3, which is in a region where both the northerly Shamal trade winds and southerly summer monsoon winds interact, shows a temporal dependence that remains relevant beyond the 50-day lag.

A comparison of the agreement between wind speed distributions is displayed across Figure 5(e)–(h). The mean quantile-quantile (QQ) plots of the first three points show very good agreement between the simulated distributions and the LENS. Slight discrepancies emerge in the lower quantiles, most notably in Point 5 where even negative wind speeds are simulated, but this is a common drawback of models that do not have mechanisms in place to ensure non-negative values (Ailliot and Monbet, 2012). The discrepancies occur below the 2 m/s level, speeds at which common wind turbines are idle, therefore, these can be disregarded for the purposes of our application. Of greater importance is the upper tail of the QQ plots and the fit in all four cases is acceptable.

4. WIND ENERGY IN SAUDI ARABIA

The average wind speed or wind power density (WPD) at measurement height or adjusted-to-hub height is frequently used in assessments of wind energy potential. WPD refers to the wind power that is available per unit area, and is given by $\frac{1}{2}\rho w^3$, where ρ represents the air density and w is the wind speed. Several methods exist for
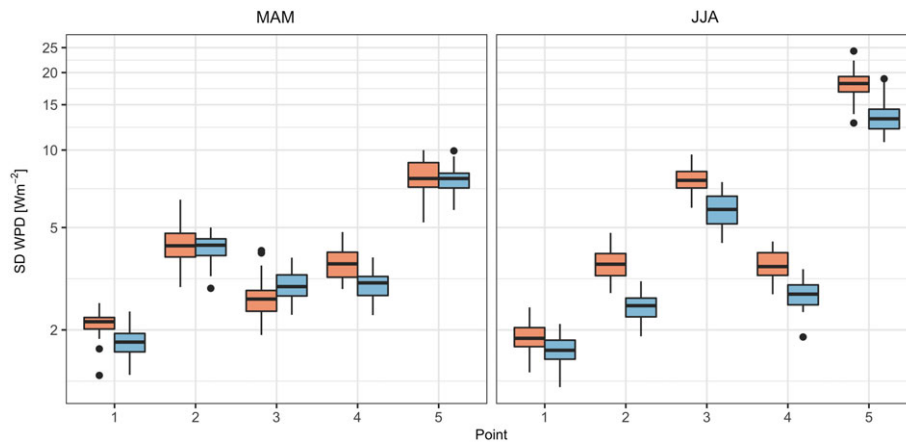
Figure 6. Standard deviation of wind power density (WPD) at 80 m for 1976–2005, in W/m^2 for the five points denoted in Figure 1, during the (a) May–April–May (MAM) season and the (b) June–July–August (JJA) season. At each point, the orange boxplot corresponds to the WPD estimates derived from the LENS; the blue boxplots correspond to 30 generated realizations of the proposed model. Log10-scale is used on the *y* axis [Color figure can be viewed at wileyonlinelibrary.com]

extrapolating near surface wind measurements to heights at which wind turbines operate (Emeis, 2012). Here we use the power law method, which assumes that the vertical wind profile at height *z* is given by $w(z) = w(z_r)(z/z_r)^a$, where z_r is a reference height and *a* is the power law exponent, which depends on the surface roughness and the thermal stability of the boundary layer. Typically, z_r represents the height at which observations are available. While the recommended value of *a* varies according to the surface type, it is frequently fixed at 1/7; this value has been found to be appropriate over open land surfaces, under neutral atmospheric stability. Since this value has been used in other wind studies in this region (e.g., Rehman *et al.*, 2007) and the coarse resolution of the LENS dataset already limits the representativeness of our estimates, we use $a = 1/7$ henceforth.

Our computation of WPD starts with surface wind speed, defined in the LENS dataset as the wind speed at the lowest model layer (variable ZBOT in the database), which is then extrapolated to an 80 m height using the above power law formula. The value of air density is assumed to be constant across space and time, and equal to 1.225 kg/m^3. Given the general increase in wind speeds during spring and summer, we focus on the March–April–May (MAM) and June–July–August (JJA) WPD seasonal averages, which are derived from daily WPD estimates; and we consider the last 30 years of the LENS dataset for comparability with recent studies. The mean MAM WPD averages approximately 36 W/m^2 across Saudi Arabia, and it increases to 39 W/m^2 during JJA, particularly along the northern coast of the Persian Gulf and west of the Asir Mountains, reflecting the strengthening of the northerly Shamal trade winds and the southerly monsoon winds respectively (Figure S9).

Figure 6 displays the standard deviation (SD) of the LENS and simulated seasonal WPD estimates for both seasons. Median differences across seasons and points tend to follow the fluctuations of the annual cycles denoted in Section 2, as higher wind speeds tend to be accompanied with greater variability. For instance, the median is nearly constant at Point 1, where the annual cycle experiences muted seasonal variation, but it more than doubles near the coastal areas of the Persian Gulf at Point 5, increasing from over 6 W/m^2 in MAM to nearly 15 W/m^2 in JJA. The fluctuations of the seasonal SD due to internal climate variability are shown to be considerable. At Point 5, the SD of JJA WPD varies between less than 15 and nearly 25 W/m^2 across the 30 LENS realizations. At other points, the variation is smaller yet still relevant, as in Point 2 or 3 for MAM, where the highest value is nearly double that of the lowest. Comparing the values from the LENS with those implied by the model realizations suggests that the model reproduces the median and variability in seasonal WPD reasonably well during MAM, and this applies to the rest of the domain (Figure 7(b)). However, negative biases arise in the median during JJA, especially at Point 5, where the bias reaches close to 4 W/m^2. Figure 7(d) shows how this bias also extends to the neighboring points along the Persian Gulf coastline.
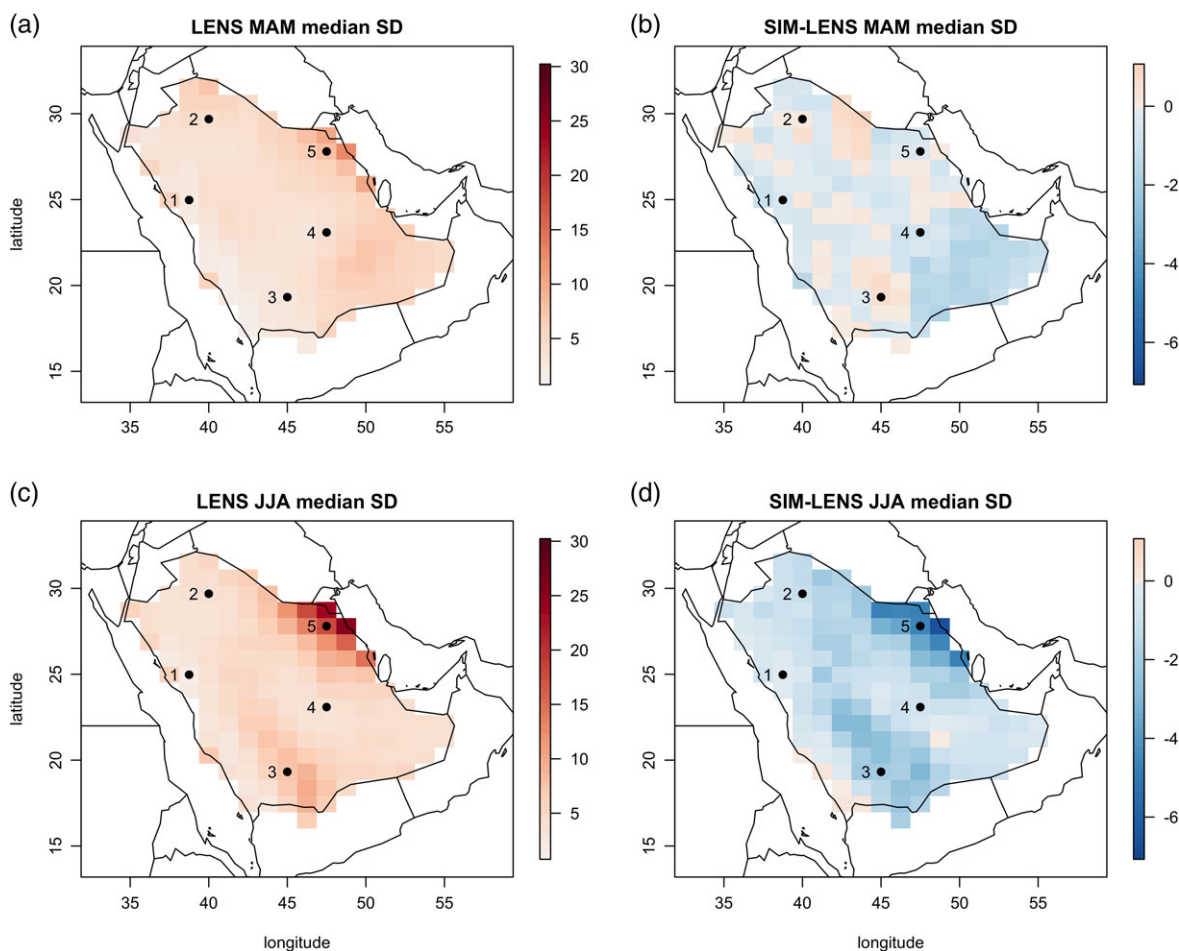
Figure 7. Median standard deviation (SD) of 80 m WPD, over 1976–2005, in W/m² for the (a, c) LENS dataset over May–April–May (MAM) and June–July–August (JJA) and the (b, d) differences with respect to the estimates derived from the 30 generated realizations. Selected points denoted in Figure 1 are also depicted [Color figure can be viewed at wileyonlinelibrary.com]

In JJA, the median SD (Figure 7(c)), as represented in the LENS, displays a spatial pattern closely resembling that of the mean WPD shown in Figure S9. However, the generated realizations have difficulty in capturing this rise in variability, as negative biases in excess of 4 and 2 W/m² can be seen over the areas surrounding Points 5 and 3 respectively. Analogous results apply to other quantiles, such as the 5% and 95% quantiles. To assess the impact of the errors of the OLS estimation on the WPD analysis, we fit the multi-variate skew-*t* distributions using the sample correlation matrices in place of the parameterized Matérn correlation matrices, and we reported the results in Figure S10.

5. DISCUSSION AND CONCLUSIONS

This study proposed a model for daily wind speeds at multiple sites based on a multi-variate skew-*t* distribution. Most recent models developed for wind speed and the two-dimensional wind field combine a Hidden Markov Model (HMM), to account for the influence of distinct weather regimes, with an autoregressive structure to capture the high degree of autocorrelation that is common in atmospheric flow (Ailliot *et al.*, 2006; Pinson and Madsen,

2012; Hering *et al.*, 2015; Bessac *et al.*, 2016). These Markov-switching autoregressive (MS-AR) models have been shown to adequately represent the marginal distribution of daily wind speed, as the Markovian structure provides the flexibility necessary to capture its higher-order moments. Recent efforts have attempted to extend this framework to the multisite setting; however, the number of sites is usually limited as inference quickly becomes problematic, in part due to the computationally demanding expectation–maximization (EM) algorithm that is used to deal with the latent Markov chain. Here we proposed a model that retains the vector-autoregressive structure, but that uses a neighbor-based scheme to reduce its dimensionality, and that replaces the latent weather regime process with a multi-variate skew-*t* distribution in the innovations to reproduce the observed skewness and excess kurtosis. The model was applied to a spatial domain of 149 points, which we partitioned into smaller regions of sizes ranging between 3 and 31 points where the assumption of stationarity is plausible, and we then performed the fit independently at each region. Our findings show that the model adequately matches the marginal distributions of daily wind speed, as well as the autocorrelation. However, the model shows some difficulty in matching the distribution of excursion durations above or below a high and low threshold respectively, especially when the degree of persistence between the two excursion types is markedly different. We conducted an analogous assessment replacing the innovations in (1) with multi-variate Gaussian distributions. As expected, the fit of marginal distributions showed a clear degradation, however, both the autocorrelation functions as well as the excursion distribution were nearly identical to those presented in Figure 5, see Figures S11 and S12. This highlights the relevance of the proposed mean structure to the performance in these two aspects, and suggests that adding seasonally varying vector-autoregressive coefficients may alleviate the shortcoming in the latter. Further work will involve incorporating a latent process to account for the large-scale variation of the wind field. Preliminary work reveals that such variation can be incorporated by exploiting the fact that the multi-variate skew-*t* distribution is constructed from a multi-variate skew-normal distribution, which has the convenient property that sums of it with a normal distribution remains within the skew-normal family.

The model was used to assess the sensitivity of wind resource potential in Saudi Arabia to the natural variability of the climate system. Studies devoted to the quantification of this resource have usually focused on the annual or seasonal mean wind speed, its temporal variability or the energy it can potentially provide as measured by WPD. A recent work by Yip *et al.* (2016) investigating the wind power potential of the Arabian Peninsula exceptionally considered other aspects, beyond its abundance, that are of particular relevance to the integration of wind energy into existing grid systems, such as the intermittency and persistence of the wind resource. To date, however, no study has addressed the impact of internal climate variability on wind energy related statistics, as its impact has only recently begun to be understood, thanks in part to projects such as the LENS that are designed to isolate its effects on the various earth system components. Our results, limited here to the variability of seasonal mean WPD based on 30 ensemble members from the LENS project, highlight the considerable effect of this source of variability on seasonal WPD. For instance, near the coastal areas of the Persian Gulf, the standard deviation of JJA WPD varies between 14 and 22 W/m^2. At other points further inland the variation is not as severe, but is still relevant. Further work will expand on this analysis to include other aspects related to WPD, along the lines of Yip *et al.* (2016). The accuracy of these WPD estimates are limited by the coarse horizontal resolution of the LENS dataset, of approximately 1°. In fact, we should expect to see departures from these as the resolution is refined, if only because of the better representation of the orography and the fine-scale processes that affect the boundary layer dynamics. A comparison of mean WPD at 80 m (Figure S9(c)) with the findings of Yip *et al.* (2016), based on a comparatively finer horizontal resolution of 0.5° (latitude) and 0.67° (longitude), hints at these potential discrepancies. Their work identifies the western mountains as offering more abundant wind resource than the Persian Gulf coastal areas, whereas here it is not the case. Ongoing work by the present authors with a high-resolution dataset are consistent with the findings of Yip *et al.* (2016); thus the estimates presented in this article should be perceived as references from which to guide further inquiries performed at finer resolution.

The lack of reliable observational data over Saudi Arabia makes validation of the proposed model against observations challenging. A more structural hurdle to such a comparison is the difference in support between the LENS dataset, whose wind measurements at each grid cell correspond to area averages, and wind speed observations, obtained from meteorological masts or remote sensing systems such as LiDARs, which may be regarded as point

measurements. Overcoming this difficulty would involve the application of some form of downscaling technique which exceeds the scope of this work.

Ideally, climate centers should make projects such as the LENS regularly available to the research community at increasingly higher resolution to promote further analyses. However, projects of this magnitude, requiring scores of ensemble members, are costly from both a computational and a storage perspective. A cheaper alternative to running these numerical models is provided by stochastic weather generators, which attempt to reproduce the statistics of specific climate processes. The proposed model was fit to only three ensemble members, to assess its ability to reproduce the variability of the entire ensemble. Although some degree of underestimation was revealed over areas with abundant wind resource that warrants further study, it reasonably reproduced the variability of seasonal WPD, thus justifying its use as a potential stochastic generator of wind speed time series for other datasets.

ACKNOWLEDGEMENTS

SUPPORTING INFORMATION

Additional Supporting Information may be found online in the supporting information tab for this article.

REFERENCES

Ailliot P, Allard D, Monbet V, Naveau P. 2015a. Stochastic weather generators: an overview of weather type models. *Journal de la Société Française de Statistique* **156**: 101–113.

Ailliot P, Bessac J, Monbet V, Pene F. 2015b. Non-homogeneous hidden Markov-switching models for wind time series. *Journal of Statistical Planning and Inference* **160**: 75–88.

Ailliot P, Monbet V. 2012. Markov-switching autoregressive models for wind time series. *Environmental Modelling & Software* **30**: 92–101.

Ailliot P, Monbet V, Prevosto M. 2006. An autoregressive model with time-varying coefficients for wind fields. *Environmetrics* **17**: 107–117.

Archer C, Simao H, Kempton W, Powell W, Dvorak M. 2017. The challenge of integrating offshore wind power in the US electric grid. Part I: Wind forecast error. *Renewable Energy* **103**: 346–360.

Arellano-Valle RB, Azzalini A. 2013. The centred parameterization and related quantities of the skew-*t* distribution. *Journal of Multivariate Analysis* **113**: 73–90.

Azzalini A. 1985. A class of distributions which includes the normal ones. *Scandinavian Journal of Statistics*: 171–178.

Azzalini A. 2005. The skew-normal distribution and related multivariate families. *Scandinavian Journal of Statistics* **32**: 159–188.

Azzalini A. 2016. *The R package sn: The Skew-Normal and Skew-t distributions (version 1.4-0)*. Italia: Università di Padova.

Azzalini A, Browne RP, Genton MG, McNicholas PD. 2016. On nomenclature for, and the relative merits of, two formulations of skew distributions. *Statistics & Probability Letters* **110**: 201–206.

Azzalini A, Capitanio A. 2003. Distributions generated by perturbation of symmetry with emphasis on a multivariate skew t-distribution. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **65**: 367–389.

Azzalini A, Capitanio A. 2014. *The Skew-Normal and Related Families. IMS Monograph Series*. Cambridge: Cambridge University Press.

Benjamini Y, Hochberg Y. 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*: 289–300.

Beranger B, Padoan SA, Sisson SA. 2017. Models for extremal dependence derived from skew-symmetric families. *Scandinavian Journal of Statistics* **44**: 21–45.

Bessac J, Ailliot P, Cattiaux J, Monbet V. 2016. Comparison of hidden and observed regime-switching autoregressive models for (*u*, *v*)-components of wind fields in the Northeast Atlantic. *Advances in Statistical Climatology, Meteorology and Oceanography* **1**: 1–16.

Bessac J, Ailliot P, Monbet V. 2015. Gaussian linear state-space model for wind fields in the Northeast Atlantic. *Environmetrics* **26**: 29–38.

Branco MD, Dey DK. 2001. A general class of multivariate skew-elliptical distributions. *Journal of Multivariate Analysis* **79**: 99–113.

Brown BG, Katz RW, Murphy AH. 1984. Time series models to simulate and forecast wind speed and wind power. *Journal of Climate and Applied Meteorology* **23**: 1184–1195.

Bruckner T, Bashmakov I, Mulugetta Y, Chum H, de la Vega Navarro A, Edmonds J, Faaij A, Fungtammasan B, Garg A, Hertwich E, Honnery D, Infield D, Kainuma M, Khennas S, Kim S, Nimir HB, Riahi K, Strachan N, Wiser R, Zhang X. 2014. Energy Systems. In *In Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, Edenhofer O, Pichs-Madruga R, Sokona Y, Farahani E, Kadner S, Seyboth K, Adler A, Baum I, Brunner S, Eickemeier P, Kriemann B, Savolainen J, Schlömer S, von Stechow C, Zwickel T, Minx JC (eds.) Cambridge, New York, NY: Cambridge University Press.

DEA. 2016. *The Danish Energy Model: Innovative, Efficient and Sustainable*. Kobenhavn, Denmark: Danish Energy Agency.

Deser C, Phillips A, Bourdette V, Teng H. 2012. Uncertainty in climate change projections: the role of internal variability. *Climate Dynamics* **38**: 527–546.

Deser C, Phillips AS, Alexander MA, Smoliak BV. 2014. Projecting North American climate over the next 50 years: uncertainty due to internal variability. *Journal of Climate* **27**: 2271–2296.

Emeis S. 2012. *Wind Energy Meteorology: Atmospheric Physics for Wind Power Generation*. Berlin: Springer Science & Business Media.

Everitt BS, Landau S, Leese M, Stahl D. 2011. *Cluster Analysis. Wiley Series in Probability and Statistics*. New York: Wiley.

Flecher C, Naveau P, Allard D, Brisson N. 2010. Model for daily wind speeds data. *Water Resources Research* **46**: W07519.

Fuentes M. 2001. A high frequency kriging approach for non-stationary environmental processes. *Environmetrics* **12**: 469–483.

Fuentes M, Smith RL. 2001. *A new class of nonstationary spatial models (Tech. rep.)* Raleigh, NC: North Carolina State University.

Genton MG. 2004. *Skew-Elliptical Distributions and their Applications: A Journey Beyond Normality*. Boca Raton, FL: CRC Press.

Gneiting T. 2002. Nonseparable, stationary covariance functions for space–time data. *Journal of the American Statistical Association* **97**: 590–600.

González-Farías G, Dominguez-Molina J, Gupta A. 2004. The closed skew-normal distribution. In *Skew-Elliptical Distributions and their Applications: A Journey Beyond Normality*, Genton MG (ed.) Boca Raton, FL: Chapman & Hall/CRC; 25–42.

Haslett J, Raftery AE. 1989. Space-time modelling with long-memory dependence: Assessing Ireland's wind power resource. *Applied Statistics* **38**: 1–50.

Hering AS, Genton MG. 2010. Powering up with space-time wind forecasting. *Journal of the American Statistical Association* **105**: 92–104.

Hering AS, Kazor K, Kleiber W. 2015. A Markov-switching vector autoregressive stochastic wind generator for multiple spatial and temporal scales. *Resources* **4**: 70–92.

Higdon D. 1998. A process-convolution approach to modelling temperatures in the North Atlantic Ocean. *Environmental and Ecological Statistics* **5**: 173–190.

Jeong J, Castruccio S, Crippa P, Genton MG. 2018. Reducing storage of global wind ensembles with stochastic generators. *The Annals of Applied Statistics* **12**: 490–509.

Jeong J, Yan Y, Castruccio S, Genton MG. 2019. A stochastic generator of global monthly wind energy with Tukey g-and-h autoregressive processes. *Statistica Sinica*. in press. http:/doi.org/10.5705/ss.202017.0474.

KA-CARE. 2012. King Abdullah City for Atomic and Renewable Energy: building the renewable energy sector in Saudi Arabia. Report.

Kay J, Deser C, Phillips A, Mai A, Hannay C, Strand G, Arblaster J, Bates S, Danabasoglu G, Edwards J, Holland M, Kushner P, Lamarque J-F, Lawrence D, Lindsay K, Middleton A, Munoz E, Neale R, Oleson K, Polvani L, Vertenstein M. 2015. The Community Earth System Model (CESM) Large Ensemble Project: A community resource for studying climate change in the presence of internal climate variability. *Bulletin of the American Meteorological Society* **96**: 1333–1349.

Kazor K, Hering AS. 2015. Assessing the performance of model-based clustering methods in multivariate time series with application to identifying regional wind regimes. *Journal of Agricultural, Biological, and Environmental Statistics* **20**: 192–217.

Lorente-Plazas R, Montávez J, Jimenez P, Jerez S, Gómez-Navarro J, Garca-Valero J, Jimenez-Guerrero P. 2015. Characterization of surface winds over the Iberian Peninsula. *International Journal of Climatology* **35**: 1007–1026.

Lütkepohl H. 2005. *New Introduction to Multiple Time Series Analysis*. Berlin: Springer Science & Business Media.

Mardia KV. 1970. Measures of multivariate skewness and kurtosis with applications. *Biometrika* **57**: 519–530.

Monbet V, Ailliot P. 2017. Sparse vector Markov switching autoregressive models. Application to multivariate time series of temperature. *Computational Statistics & Data Analysis* **108**: 40–51.

Morris SA, Reich BJ, Thibaud E, Cooley D. 2017. A space-time skew-*t* model for threshold exceedances. *Biometrics* **73**: 749–758.

Paciorek CJ, Schervish MJ. 2006. Spatial modelling using a new class of nonstationary covariance functions. *Environmetrics* **17**: 483–506.

Pinson P. 2013. Wind energy: forecasting challenges for its operational management. *Statistical Science* **28**: 564–585.
Pinson P, Madsen H. 2012. Adaptive modelling and forecasting of offshore wind power fluctuations with Markov-switching autoregressive models. *Journal of Forecasting* **31**: 281–313.
Rehman S, Ahmad A. 2004. Assessment of wind energy potential for coastal locations of the Kingdom of Saudi Arabia. *Energy* **29**: 1105–1115.
Rehman S, El-Amin I, Ahmad F, Shaahid S, Al-Shehri A, Bakhashwain J. 2007. Wind power resource assessment for Rafha, Saudi Arabia. *Renewable and Sustainable Energy Reviews* **11**: 937–950.
Sampson PD. 2010. Constructions for nonstationary spatial processes. In *Handbook of Spatial Statistics*, CRC/Chapman and Hall: Boca Raton, FL, 119–130.
Sampson PD, Guttorp P. 1992. Nonparametric estimation of nonstationary spatial covariance structure. *Journal of the American Statistical Association* **87**: 108–119.
Schweinberger M, Babkin S, Ensor KB. 2017. High-dimensional multivariate time series with additional structure. *Journal of Computational and Graphical Statistics* **26**: 610–622.
Shaahid S, Al-Hadhrami LM, Rahman M. 2014. Potential of establishment of wind farms in western province of Saudi Arabia. *Energy Procedia* **52**: 497–505.
Wilks DS. 2011. *Statistical Methods in the Atmospheric Sciences*. San Diego, CA: Academic Press.
Wilks DS, Wilby RL. 1999. The weather generation game: a review of stochastic weather models. *Progress in Physical Geography* **23**: 329–357.
WWEA. 2016. *WWEA Quarterly Bulleting Issue 1-206: Community Wind Special*. Bonn: World Wind Energy Association.
Yip CMA, Gunturu UB, Stenchikov GL. 2016. Wind resource characterization in the Arabian Peninsula. *Applied Energy* **164**: 826–836.
Zhu X, Genton MG. 2012. Short-term wind speed forecasting for power system operations. *International Statistical Review* **80**: 2–23.

APPENDIX: THE SKEW-*T* DISTRIBUTION

In this appendix we collect a series of relevant properties and results related to the skew-*t* distribution.

If $\mathbf{Y} = \boldsymbol{\xi} + \boldsymbol{\omega}\mathbf{Z}$, with $\mathbf{Z}$ being the standard ST_d variate, that is, $\mathbf{Z} \sim ST_d(\mathbf{0}, \bar{\boldsymbol{\Omega}}, \boldsymbol{\alpha}, \nu)$, then

$$\boldsymbol{\mu} = \mathrm{E}(\mathbf{Y}) = \boldsymbol{\xi} + \boldsymbol{\omega}\boldsymbol{\mu}_z, \quad \nu > 1,$$
$$\boldsymbol{\Sigma} = \mathrm{var}(\mathbf{Y}) = \frac{\nu}{\nu-2}\boldsymbol{\Omega} - \boldsymbol{\omega}\boldsymbol{\mu}_z\boldsymbol{\mu}_z^\top\boldsymbol{\omega}, \quad \nu > 2,$$

with $\boldsymbol{\mu}_z = b_\nu \boldsymbol{\delta}$, where

$$b_\nu = \frac{\sqrt{\nu}\,\Gamma\left\{\frac{1}{2}(\nu-1)\right\}}{\sqrt{\pi}\,\Gamma\left(\frac{1}{2}\nu\right)}, \quad \nu > 1, \quad \text{and} \quad \boldsymbol{\delta} = \left(1 + \boldsymbol{\alpha}^\top\bar{\boldsymbol{\Omega}}\boldsymbol{\alpha}\right)^{-1/2} \bar{\boldsymbol{\Omega}}\boldsymbol{\alpha}.$$

Mardia's measure of multi-variate kurtosis (Mardia, 1970), for a *d*-dimensional random variable $\mathbf{X}$, is given by

$$\gamma_2^M = \mathrm{E}\left[\{(\mathbf{X}-\boldsymbol{\mu})^\top\boldsymbol{\Sigma}^{-1}(\mathbf{X}-\boldsymbol{\mu})\}^2\right] - d(d+2)$$

and for a ST_d variate, assuming that $\nu > 4$,

$$\gamma_2^M = \frac{2d(d+2)}{\nu-4} + \frac{4(d+2)}{(\nu-3)(\nu-4)}\beta_0^2 + 2\left\{\frac{2\nu}{(\nu-3)b_\nu^2} - \frac{3(\nu-3)^2 - 6}{(\nu-3)(\nu-4)}\right\}\beta_0^4$$

where $\beta_0^2 = \boldsymbol{\mu}_0^\top\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}_0$, $\boldsymbol{\mu}_0 = \boldsymbol{\omega}\boldsymbol{\mu}_z$.