

SENSITIVITY ANALYSIS OF WIND ENERGY RESOURCES WITH BAYESIAN NON-GAUSSIAN AND NONSTATIONARY FUNCTIONAL ANOVA

BY JIACHEN ZHANG^{1,a}, PAOLA CRIPPA^{2,c}, MARC G. GENTON^{3,d} AND STEFANO CASTRUCCIO^{1,b}

¹*Department of Applied and Computational Mathematics and Statistics, University of Notre Dame, ^ajzhang19@alumni.nd.edu, ^bscastruc@nd.edu*

²*Department of Civil and Environmental Engineering & Geosciences, University of Notre Dame, ^cpcrippa@nd.edu*

³*Statistics Program, King Abdullah University of Science and Technology, ^dmarc.genton@kaust.edu.sa*

The transition from nonrenewable to renewable energies represents a global societal challenge, and developing a sustainable energy portfolio is an especially daunting task for developing countries where little to no information is available regarding the abundance of renewable resources such as wind. Weather model simulations are key to obtain such information when observational data are scarce and sparse over a country as large and geographically diverse as Saudi Arabia. However, output from such models is uncertain, as it depends on inputs such as the parametrization of the physical processes and the spatial resolution of the simulated domain. In such situations a sensitivity analysis must be performed, and the input may have a spatially heterogeneous influence of wind. In this work we propose a latent Gaussian functional analysis of variance (ANOVA) model that relies on a nonstationary Gaussian Markov random field approximation of a continuous latent process. The proposed approach is able to capture the local sensitivity of Gaussian and non-Gaussian wind characteristics such as speed and threshold exceedances over a large simulation domain, and a continuous underlying process also allows us to assess the effect of different spatial resolutions. Our results indicate that: (1) the nonlocal planetary boundary layer scheme and high spatial resolution are both instrumental in capturing wind speed and energy (especially over complex mountainous terrain), and (2) the impact of planetary boundary layer scheme and resolution on Saudi Arabia's planned wind farms is small (at most 1.4%). Thus, our results lend support for the construction of these wind farms in the next decade.

1. Introduction. Global fossil fuel consumption has increased eight-fold since 1950 and remains the major contributor to global climate change and atmospheric air pollution (Dincer (2018), REN21 Secretariat (2021)). Given that oil is a nonrenewable energy resource, the rapid depletion of its reserves is a problem that humanity will face in the foreseeable future, prompting the development of alternative, renewable sources of energy. Wind energy, the focus of this work, grew significantly in recent decades and has already contributed to the reduction of greenhouse gas emissions and local air pollution. According to recent estimates, in 2019, 5% of the world's electricity was generated by wind, with the largest share coming from the United States and China (REN21 Secretariat (2021), Renewables (2018)). Further, although European countries have a smaller absolute installed wind capacity, their share in the energy portfolio is one of the highest, and in Denmark wind is by far the most widespread form of energy (56.3%).

Despite the growth of wind energy in many countries, in areas such as the Middle East and North Africa (MENA) the proportion of power generated by this renewable energy is

Received December 2021; revised April 2023.

Key words and phrases. Bayesian hierarchical model, stochastic partial differential equation, functional ANOVA, wind energy.

among the lowest worldwide (Mohseni-Cheraghloou (2021)). Saudi Arabia, in particular, is almost exclusively reliant on its abundant oil reserves for its internal energy demand; it has only recently outlined plans to diversify its energy portfolio using renewable energy sources. Given the country's latitude and climate, solar energy is expected to be the primary form of renewable energy (Aghahosseini, Bogdanov and Breyer (2020)), but the recently proposed "Vision 2030" plan aims to generate 16 GW of wind energy, which will place the country in a leading position for wind energy generation (NREP (2018), Nurunnabi (2017)). Unlike solar energy, which can only be harnessed during daytime, wind is always available and often peaks during nighttime, thus providing a promising opportunity to integrate complementary renewable energy resources and hence provide a continuous and reliable supply to the grid.

A comprehensive assessment of wind energy resources in a country as large as Saudi Arabia cannot be informed solely by observational records, which are very sparse in space, limited in time and challenging to retrieve. Instead, assessments must be integrated with climate model simulations, which provide spatially resolved, dynamic and physically consistent information on wind speed. Over the past few years, considerable progress has been made in this regard. The initial efforts were directed toward analyzing publicly available simulations from either global models (Jeong et al. (2018), Jeong et al. (2019)) or regional simulations (Chen et al. (2018)), such as the MENA COordinated Regional climate Downscaling Experiment (MENA CORDEX, Jones, Giorgi and Asrar (2011)). More recently, new and more detailed assessments have been performed by analyzing ad hoc high-resolution regional simulations (Crippa et al. (2021), Giani et al. (2020), Tagle et al. (2019), Zhang et al. (2021)) from a state-of-the-art model: the Weather Research and Forecasting (WRF, Skamarock et al. (2019)).

Numerical models can provide a comprehensive, spatially resolved assessment of wind over a large country, but each simulation depends on several inputs, including the parameterizations for physical processes (especially near the surface), boundary conditions as well as spatial and temporal resolution. Therefore, any assessment must be performed using a collection (*ensemble*) of simulations, with each ensemble member representing a different choice of the aforementioned input. To provide a reliable uncertainty quantification of wind energy, it is, therefore, crucial to understand if, to what extent, and where the final wind energy estimates depend on these model choices.

From a methodological point of view, sensitivity analysis of a variable dependent on factors with multiple levels is one of the oldest and most established problems in statistics, long before the formulation of numerical simulations. The significance of a factor is assessed by comparing its variability to the measurement error (analysis of variance, ANOVA). In its original formulation, ANOVA is aimed at sensitivity analysis of a single variable and a finite number of factors with different levels being observed independently. More flexible models have been proposed to perform ANOVA on more complex data structures. In this work we focus on functional ANOVA (FANOVA), a method developed to perform sensitivity analysis on spatial, temporal or even spatiotemporal process (Stone et al. (1997)). FANOVA has been used in many fields, such as public health (Huang et al. (2000), Ullah and Finch (2013), Zhang et al. (2019)), chemistry (Tarrío-Saavedra et al. (2011)) and geoscience (Kaufman and Sain (2010), Qu, Dai and Genton (2021), Sain, Nychka and Mearns (2011), Sun and Genton (2012)), and the majority of the applications focused on time series. Stone et al. (1997) provided a comprehensive review of FANOVA. More recently, Zhang, Hodges and Banerjee (2009) proposed a smoothed ANOVA model using a Bayesian framework that treated space as a factor in the model for multivariate observations with an areal spatial structure. Kaufman and Sain (2010) proposed a Bayesian framework for spatial FANOVA to perform sensitivity analysis of present and future regional climate simulations. Their work laid the foundation for a framework to perform the local sensitivity analysis of climate simulations but has several

limitations. First, the proposed model assumes Gaussian data, thereby limiting the analysis to continuous variables at a sufficiently high level of temporal aggregation. Second, the spatial dependence structure was restricted to stationary isotropic models. Although this is a useful simplifying assumption, it is unrealistic for large simulation domains (Cressie and Huang (1999), Jun and Stein (2008), Yue and Speckman (2010)). Finally, the assumption of a discrete spatial model without an underlying continuous process would not allow one to perform FANOVA to ensemble members with differing spatial resolution.

In this work we propose a nonstationary FANOVA model that is able to overcome the aforementioned methodological limitations. The key idea is that the latent process in the model can be regarded as the solution of a stochastic partial differential equation (SPDE). More specifically:

1. The proposed model relies on the flexible class of latent Gaussian models: the marginal distribution is assumed to be a member of the exponential family, with the expected values being spatially dependent through an appropriate link function and one or more latent Gaussian fields. Bayesian inference for this class of latent Gaussian models can be performed by deterministic approximation of the integrals in the posterior distribution using the integrated Laplace approximation (INLA) method (Rue, Martino and Chopin (2009)).

2. The definition of a spatial model through a solution of an SPDE allows one to naturally generalize FANOVA to a nonstationary setting by generalizing the differential operator and assuming its nonhomogeneity in space while still allowing a theoretically valid model. The SPDE chosen allows for an “explicit link” between the SPDE solution and the Gaussian Markov Random Field (Lindgren, Rue and Lindström (2011)). This allows likelihood evaluations with a sparse precision matrix and hence fast and affordable inference for very large spatial data.

3. The use of a FANOVA with a latent Gaussian field assumes an underlying continuous spatial model, which naturally accounts for datasets with different resolutions.

From a methodological standpoint, our work proposes a generalization of the FANOVA approach in Yue et al. (2019) by allowing a more flexible class of nonstationary models in the latent Gaussian effects and by quantifying the improvement against a stationary model with an extensive simulation study. Our work also explores the added value of space-time interaction.

While the proposed model is motivated by assessing the robustness of Saudi Arabia’s current plan for wind energy installation, its scope is more general, as it can be applied to any model ensemble resolved in space and time where one seeks a local sensitivity analysis with respect to a controlled number of factors.

In Section 2 we introduce the data set of numerical simulations used for this study. Sections 3 and 4 describe the proposed statistical methodology and inference, respectively. Section 5 validates the statistical model with a simulation study. In Section 6 we apply the proposed method to the WRF simulated data of wind speed to assess wind energy sensitivity across Saudi Arabia. Section 7 provides the conclusions of this study.

2. Data description. This study relies on an ensemble of high-resolution WRF simulations of the Arabian Peninsula during the years 2013–2016 (Giani et al. (2020)). For the ease of understanding, we show a topographical map of Saudi Arabia, along with the names of the regions that will be used throughout this work in Figure S1. The ensemble has been designed to explore model’s sensitivity toward different spatial resolutions and planetary boundary layer (PBL) schemes. The simulated domain comprises 339×299 and 549×499 grids at a resolution of nine km and six km, respectively. Wind speed is resolved on a vertical grid comprising 40 layers, which are more closely spaced near the surface and sparser

TABLE 1
Resolution and planetary boundary layer (PBL) setup of the four ensemble members

		Resolution	
		9 km ($j = 0$)	6 km ($j = 1$)
PBL	MYJ ($i = 0$)	Run 1	Run 4
	ACM2 ($i = 1$)	Run 2	Run 3

closer to the boundary layer height. In this work we consider wind between 10 and 110 m, which correspond to the heights of the majority of commercial wind turbines. The initial and boundary conditions used to drive WRF are obtained from high-resolution European Centre for Medium-Range Weather Forecast (HRES-ECMWF, [European Centre for Medium-Range Weather Forecasts \(2016\)](#)) model. The two PBL schemes adopted are: the Mellor–Yamada–Janjić (MYJ, [Janjić \(1994\)](#)) and the asymmetric convective model (ACM2, [Pleim \(2007\)](#)). MYJ is a local scheme where the vertical diffusion occurs only between neighboring cells, while ACM2 is a nonlocal scheme in which the diffusion also occurs between nonneighboring cells and counter-gradient fluxes. The spatial resolution of the simulation and PBL setup among the different runs are summarized in Table 1. Figure 1 shows the simulated domain and the wind speed averaged over the simulated time period for all four ensemble members. Discrepancies in the magnitude of the simulated wind speed are observed across different regions and runs, with higher resolution simulations being able to more accurately charac-




FIG. 1. Wind speed averaged across the years 2013–2016 simulated using the Weather and Research Forecasting (WRF) model in: (a) MYJ-9 km, (b) ACM2-9 km, (c) MYJ-6 km, and (d) ACM2-6 km.




FIG. 2. MYJ-9 km and ACM2-9 km monthly wind speed and fitted value according to model (7) at two locations P (Plain) and M (Mountain) indicated by the crosses in Figure 1(a).

terize the topography (and hence the wind speed) in mountainous regions and MYJ resulting in higher winds at finer resolution (see panels (c) and (d)). A spatial sensitivity analysis is, therefore, necessary to quantify the differences across resolutions and PBL schemes, and motivates the development of the FANOVA model in the next section.

In Figure 2 the monthly wind speed at two selected locations, represented by crosses in Figure 1(a), shows an interannual, seasonal behavior. Wind speeds are the highest during the summer and the lowest during winter at location P, a plain area within the Rub' al Khali region (see Figure S1 for an indication of Saudi Arabia's regions) and the opposite at location M in the Hijaz mountain region. The annual cycle observed is attributable to Saudi Arabia's location within the trade-wind belt of the Northern Hemisphere (Hasanean and Almazroui (2015)). The strong northerly flow contributes to the high wind speeds during the summer. In addition, the southeastern wind from the Indian Ocean travels to the southeastern part of the Arabian Peninsula during the monsoon season, further enhancing the summer wind. However, during winter, the north wind traveling from the Mediterranean to the Persian Gulf weakens.

Throughout this work N denotes the total number of locations, and the vector $\mathbf{Y}_{ij}(t) = (Y_{ij}(\mathbf{s}_1, t), \dots, Y_{ij}(\mathbf{s}_N, t))^T$ denotes the N -dimensional vector of the variable of interest (all aggregated at monthly level) at locations $\mathbf{s}_1, \dots, \mathbf{s}_N$, for the i th PBL settings ($i = 0, 1$ indicate MYJ and ACM2, respectively) the j th resolution ($j = 0, 1$ indicate nine km and six km, respectively), and the t th month over the four year time span ($t = 1, \dots, 4 \times 12$).

3. Model. We propose a Bayesian spatiotemporal non-Gaussian FANOVA model with nonstationary dependence. Because the ensemble introduced in Section 2 has PBL and resolution as inputs, we assume two predictors; however, the model can be generalized to any number,

$$\begin{aligned}
 \mathbf{Y}_{ij}(t) &\sim h(\boldsymbol{\mu}_{ij}(t), \boldsymbol{\theta}_{\text{MRG}}), \\
 \mathbf{g}(\boldsymbol{\mu}_{ij}(t)) &= \mathbf{f}^{(T)}(t) + \mathbf{f}_{ij}^{(F)} + i\boldsymbol{\beta}_{\text{PBL}} + j\boldsymbol{\beta}_{\text{RES}} + \boldsymbol{\epsilon}_{ij}(t), \\
 \mathbf{f}^{(T)}(t) &= \sum_{k=1}^K \left\{ \boldsymbol{\zeta}_k \sin\left(\frac{2\pi kt}{\delta}\right) + \boldsymbol{\zeta}'_k \cos\left(\frac{2\pi kt}{\delta}\right) \right\}, \\
 \mathbf{f}_{ij}^{(F)} &= \boldsymbol{\beta}_{\text{ALT}} \mathbf{A}_j,
 \end{aligned}
 \tag{1}$$

where bold indicates spatial vectors, for example, $\boldsymbol{\zeta}_k = (\zeta_{k;1}, \dots, \zeta_{k;N})^T$ and similarly for all other vectors. Each element of the vector $\mathbf{Y}_{ij}(t)$ has a distribution from the exponential

family $h(\cdot)$, with an expectation $\mu_{ij}(t)$ and possibly other marginal parameters θ_{MRG} . The vector $\mathbf{g}(\cdot)$ consists of the same link function $g(\cdot)$ for every location, which depends on the distribution (e.g., for the Gaussian distribution we have the identity, whereas for the Bernoulli distribution we have a logit). We assume that the vector of expectations $\mu_{ij}(t)$ is modeled using a location-specific N -dimensional vector of time effects $\mathbf{f}^{(T)}(t)$, a location-specific (possibly simulation-dependent) N -dimensional vector of time-invariant (fixed) effects $\mathbf{f}_{ij}^{(F)}$, two Gaussian N -dimensional spatial random vectors β_{PBL} , β_{RES} which represent the local contribution of PBL and resolution, respectively. In the Supplementary Material (Zhang et al. (2024)), we also consider a more general model with space-time interaction, which did yield very marginal improvements but was not implemented for the sake of simplicity and interpretability. Finally, $\epsilon_{ij}(t)$ represent a Gaussian white noise independent of factors, time and space with variance σ^2 .

The time effect $\mathbf{f}^{(T)}(t)$ is described using K annual harmonics controlled at each location $\mathbf{s}_n$ by different parameters $\theta_{\text{time},n} = \{\zeta_{k;n}, \zeta'_{k;n}, k = 1, \dots, K\}$ so that the total number of temporal parameters is $\theta_{\text{time}} = \{\theta_{\text{time},n}, n = 1, \dots, N\}$. In our case the time invariant effect $\mathbf{f}_{ij}^{(F)}$ represents the contribution of the altitude, which is expected to be linear, according to a parameter β_{ALT} , since wind is generally higher in mountainous regions with more complex terrain; see Figure S2 and associated diagnostics. In our application the altitude map $\mathbf{A}_j$ depends on the resolution of the simulation, hence the subscript j is used. In a standard two-way ANOVA, the two random effects β_{PBL} and β_{RES} will be independent in space, but in FANOVA they are assumed to be the realizations (independent in time) from a Gaussian random field,

$$(2) \quad \begin{aligned} \beta_{\text{PBL}} &\sim \mathcal{N}(\mathbf{0}, \Sigma(\theta_{\text{PBL}})), \\ \beta_{\text{RES}} &\sim \mathcal{N}(\mathbf{0}, \Sigma(\theta_{\text{RES}})), \end{aligned}$$

where θ_{PBL} and θ_{RES} are unknown parameters.

In this work we assume $\Sigma(\theta_\ell) = (\Sigma_{ij}(\theta_\ell))_{ij}$ where $\ell \in \{\text{PBL}, \text{RES}\}$. Instead of providing an explicit parametrization of the covariance matrix through a covariance function, we use a fundamental result in spatial statistics that links a class of covariance functions to the solution of an SPDE. Specifically, we focus on the Matérn function

$$(3) \quad \Sigma_{ij}(\theta_\ell) = \frac{1}{\tau_\ell 2^{\nu_\ell - 1} \Gamma(\nu_\ell)} (\kappa_\ell \|\mathbf{s}_i - \mathbf{s}_j\|)^{\nu_\ell} K_{\nu_\ell}(\kappa_\ell \|\mathbf{s}_i - \mathbf{s}_j\|),$$

where K_ν is the modified Bessel function of the second kind of order ν , $\|\mathbf{s}_i - \mathbf{s}_j\|$ is the Euclidean distance between two generic locations $\mathbf{s}_i, \mathbf{s}_j \in \mathbb{R}^2$ and the parameter vector is $\theta_\ell = (\tau_\ell, \kappa_\ell, \nu_\ell)^\top$. The parameter τ_ℓ controls the marginal precision, and κ_ℓ controls the range of the spatial dependence: when we consider a distance $\sqrt{8\nu_\ell}/\kappa_\ell$, the spatial correlation is near 0.1 for all ν_ℓ (Stein (1999)). Finally, ν_ℓ controls the degree of smoothness of the process and is usually fixed because of poor identifiability.

A well-known result (Whittle (1954), Whittle (1963)) stipulates that a Gaussian process with Matérn covariance is the (unique) stationary solution of a fractional diffusion-reaction SPDE,

$$(4) \quad (\kappa_\ell^2 - \Delta)^{\nu_\ell/2 + 1/2} (\tau_\ell \beta_\ell) = \mathcal{W}(\mathbf{s}),$$

where Δ is the Laplacian operator, τ_ℓ controls the variance and $\mathcal{W}(\mathbf{s})$ is a spatial Gaussian white noise with unit variance.

A stationary model resulting from solving (4) is overly simplistic over a large simulated domain such as the one used in our application. However, one of the main advantages of the SPDE approach is that it can be used as a baseline to formulate nonstationary models

that are automatically well-defined, a task considerably more challenging to perform when relying on covariance-based models. Several options are available to achieve nonstationarity, ranging from nested SPDEs (Bolin and Lindgren (2011)) to modification of the differential operator (Fuglstad et al. (2019)). Some recent works have proposed a nonstationary model based on a local deformation of the SPDE with a changing dependence structure across large geographical descriptors such as land and ocean in a Euclidean (Fuglstad et al. (2019)) and global (Fuglstad and Castruccio (2020), Hu, Fuglstad and Castruccio (2022)) domain. Since in this work the geography of the problem does not suggest a natural partition of the domain, we instead rely on a basis decomposition approach and assume that β_ℓ is a solution of a generalization of (4) with varying range and precision,

$$(5) \quad (\kappa_\ell^2(\mathbf{s}) - \Delta)^{\nu_\ell/2+1/2} \{\tau_\ell(\mathbf{s})\beta_\ell\} = \mathcal{W}(\mathbf{s}),$$

where ν_ℓ is fixed and $\tau_\ell(\mathbf{s})$ and $\kappa_\ell^2(\mathbf{s})$ change in space, according to some basis function,

$$(6) \quad \begin{aligned} \log(\tau_\ell(\mathbf{s})) &= b_{0;\ell}^{(\tau)}(\mathbf{s}) + \sum_{k=1}^p b_{k;\ell}^{(\tau)}(\mathbf{s})\theta_{k;\ell}^{(\tau)}, \\ \log(\kappa_\ell(\mathbf{s})) &= b_{0;\ell}^{(\kappa)}(\mathbf{s}) + \sum_{k=1}^p b_{k;\ell}^{(\kappa)}(\mathbf{s})\theta_{k;\ell}^{(\kappa)}, \end{aligned}$$

so that the total number of spatial parameters is $\theta_{\text{space}} = \{b_{k;\ell}^{(\kappa)}, b_{k;\ell}^{(\tau)}, k = 0, \dots, p, \ell \in \{\text{PBL}, \text{RES}\}\}$. All parameters are assumed to have independent vague Gaussian priors with mean zero and variance equal to 1000.

In the Supplementary Material, we have also tested the barrier model (Bakka et al. (2019)) to account for potentially abrupt changes between plains and mountains and obtained far superior results with the choice of nonstationary model in (5) and (6).

4. Inference. Inference is performed in two steps, to reduce the overall dimension of the parameter space at each step. First, $f^{(T)}(t)$ in (1) is estimated independently for each location to capture the annual periodicity of the (latent) Gaussian field. Second, inference on θ_{MRG} , β_{ALT} , θ_{space} is performed conditionally on the posterior mean of the parameters of $f^{(T)}(t)$. A numerical study of this two-step approach against joint space-time inference on a smaller subset is performed in the Supplementary Material, showing that the posteriors are similar for both approaches.

4.1. Step 1: Temporal structure. In the first step, we consider each location $\mathbf{s}_n$ independently, and a marginal time series version of (1) with no spatial and covariate effects,

$$(7) \quad \begin{aligned} Y_{ij}(\mathbf{s}_n, t) &\sim h(\mu_{ij}(\mathbf{s}_n, t), \theta_{\text{MRG}}), \\ g(\mu_{ij}(\mathbf{s}_n, t)) &= \sum_{k=1}^K \left\{ \zeta_{k,n} \sin\left(\frac{2\pi kt}{\delta}\right) + \zeta'_{k,n} \cos\left(\frac{2\pi kt}{\delta}\right) \right\}. \end{aligned}$$

The posterior distribution of $\theta_{\text{time},n} = \{\zeta_{k;n}, \zeta'_{k;n}, k = 1, \dots, K\}$ is then obtained, and in the following steps, these parameters are considered fixed at their posterior mean. This conditional approach has, in general, a small impact on the overall assessment of the uncertainty, as theoretical results (Edwards, Castruccio and Hammerling (2020)) have shown asymptotic correct results, and numerical results suggest an overall small impact on error propagation (Castruccio and Guinness (2017)). Inference can be performed independently on multiple cores on a laptop or a workstation. The fitted values, based on the model inference described above for two selected locations in Figure 1, are represented by black dashed line in Figure 2.

4.2. *Step 2: Covariates and spatial effects.* Once the temporal parameters are estimated, the parameters θ_{MRG} , β_{ALT} , θ_{space} associated to the marginal model, the covariate effect parameters and the spatial dependence in (1), respectively, must be estimated.

Inference in this step is especially challenging, given the large number of spatial locations, which imply a large covariance matrix in (2) and hence challenges in storing it and performing linear algebra operations. If the diffusion-reaction SPDE (1) as well as its generalization (5) have a smoothness parameter fixed to an integer number (in this work we assume $\nu_\ell = 1$), then the underlying process can be shown to have the Markov property (Lindgren, Rue and Lindström (2011)). This implies that the continuous solution can be conveniently discretized via finite volumes with a discrete Gaussian Markov random field with a sparse precision matrix, hence leading to fast and affordable inference.

The SPDE in (5) is solved by dividing the domain into a triangulation and then approximating the continuous solution by a piecewise linear function for each triangle, with the quality of the approximation being determined by the size of the triangulation (Lindgren, Rue and Lindström (2011)). Formally, the finite element approximation of the solution to the SPDE is

$$\beta_\ell(\mathbf{s}) \approx \sum_{k=1}^T \psi_k(\mathbf{s}) w_k,$$

where $\psi_k(\mathbf{s})$ are piecewise linear basis functions that are equal to one at vertex k , linearly decreasing to zero to nearby vertices and zero everywhere else. The weights w_k are Gaussian distributed, and $k = 1, \dots, T$ where T is the number of vertices in the triangulation. We define three $T \times T$ matrices as

$$\begin{aligned} \mathbf{C} &= \text{diag}(\mathbf{C}_{ii}), \quad \mathbf{C}_{ii} = \langle \psi_i, 1 \rangle, \\ \mathbf{G}_{ij} &= \langle \nabla \psi_i, \nabla \psi_j \rangle, \\ \mathbf{K} &= \kappa_\ell^2 \mathbf{C} + \mathbf{G}, \end{aligned}$$

where $\langle \cdot \rangle$ is the inner product and ∇ is the gradient. The sparse precision matrix of the joint distribution of the weights in the case of a stationary SPDE (4) can be written as (Lindgren, Rue and Lindström (2011))

$$(8) \quad \mathbf{Q} = \mathbf{K} \mathbf{C}^{-1} \mathbf{K} = \tau_\ell^2 (\kappa_\ell^4 \mathbf{C} + 2\kappa_\ell^2 \mathbf{G} + \mathbf{G} \mathbf{C}^{-1} \mathbf{G}).$$

If we further define $\mathbf{T}$ and $\mathbf{K}$ as diagonal matrices, where $\mathbf{T} = \text{diag}(\tau(\mathbf{s}))$ and $\mathbf{K} = \text{diag}(\kappa(\mathbf{s}))$, then the precision matrix for the nonstationary SPDE (5) can be written as

$$\mathbf{Q} = \mathbf{T} (\mathbf{K}^2 \mathbf{C} \mathbf{K}^2 + \mathbf{K} \mathbf{G} + \mathbf{G}^\top \mathbf{K} + \mathbf{G} \mathbf{C}^{-1} \mathbf{G}) \mathbf{T}.$$

4.3. *Bayesian inference for latent Gaussian models.* To further ease the computational burden, Bayesian inference for the proposed model (1) will be performed using the INLA approximation (Rue, Martino and Chopin (2009)), a deterministic method for fast approximation of high-dimensional integrals in Bayesian inference. A comprehensive review of INLA can be found in the works of Rue, Martino and Chopin (2009) and Bakka et al. (2018).

The INLA approach assumes a latent Gaussian model of which our FANOVA model (2) is a particular case, and throughout this section we omit the subscripts indicating the input and the temporal component. We, therefore, assume we have a vector $\mathbf{Y} = (Y(\mathbf{s}_1), \dots, Y(\mathbf{s}_n))^\top$ sampled at locations $\mathbf{s}_1, \dots, \mathbf{s}_n$ whose marginal distribution is from the exponential family, possibly described by the hyperparameter vector θ . We assume that, conditional on a latent spatial field $\mathbf{x}$, the observations are marginally independent,

$$\pi(\mathbf{y}|\mathbf{x}, \theta) = \prod_{i=1}^n \pi(y(\mathbf{s}_i)|x(\mathbf{s}_i)),$$

where $\mathbf{x} = (x(\mathbf{s}_1), \dots, x(\mathbf{s}_n))^T$ is a zero mean Gaussian field modeled with the SPDE approach in the previous section with precision matrix $\mathbf{Q}(\boldsymbol{\theta})$. Under this model the joint posterior distribution of the latent effects and hyperparameters can be written as

$$\begin{aligned}\pi(\mathbf{x}, \boldsymbol{\theta} | \mathbf{y}) &\propto \pi(\boldsymbol{\theta}) \pi(\mathbf{x} | \boldsymbol{\theta}) \prod_{i=1}^n \pi(y(\mathbf{s}_i) | x(\mathbf{s}_i), \boldsymbol{\theta}) \\ &\propto \pi(\boldsymbol{\theta}) |\mathbf{Q}(\boldsymbol{\theta})|^{1/2} \exp\left\{-\frac{1}{2} \mathbf{x}^T \mathbf{Q}(\boldsymbol{\theta}) \mathbf{x}\right\} \prod_{i=1}^n \pi(y(\mathbf{s}_i) | x(\mathbf{s}_i), \boldsymbol{\theta}),\end{aligned}$$

where $|\mathbf{Q}(\boldsymbol{\theta})|$ is the determinant of the precision matrix. To perform inference, it is of interest to obtain: (1) $\pi(\mathbf{x} | \mathbf{y})$ from which all the marginal distributions $\pi(x(\mathbf{s}_i) | \mathbf{y})$ can be obtained and (2) $\pi(\boldsymbol{\theta}_j | \mathbf{y})$, the marginal posterior distributions of the hyperparameters. Both distributions can be obtained from the following integrals:

$$\begin{aligned}\pi(x(\mathbf{s}_i) | \mathbf{y}) &= \int \pi(x(\mathbf{s}_i) | \boldsymbol{\theta}, \mathbf{y}) \pi(\boldsymbol{\theta} | \mathbf{y}) d\boldsymbol{\theta}, \\ \pi(\boldsymbol{\theta}_j | \mathbf{y}) &= \int \pi(\boldsymbol{\theta} | \mathbf{y}) d\boldsymbol{\theta}_{-j}.\end{aligned}$$

INLA predicates a Laplace approximation of the hyperparameter posterior $\tilde{\pi}(\boldsymbol{\theta} | \mathbf{y})$ using a Gaussian distribution. Such approximation would then allow us to obtain the posterior marginals of the latent parameter x_i as

$$\tilde{\pi}(x(\mathbf{s}_i) | \mathbf{y}) = \sum_k \tilde{\pi}(x(\mathbf{s}_i) | \boldsymbol{\theta}_k, \mathbf{y}) \times \tilde{\pi}(\boldsymbol{\theta}_k | \mathbf{y}) \times \Delta_k,$$

where Δ_k are the weights associated with a vector $\boldsymbol{\theta}_k$ of hyperparameters in a grid.

To summarize, fast Bayesian inference in this work is obtained by exploiting: (1) the latent Gaussian structure of the model, which allows an efficient deterministic Laplace approximation of the integrals determining the posterior distributions, and (2) a spatial model for the latent Gaussian field obtained through a solution of an SPDE, which results in a sparse precision matrix and allows us to take advantage of computationally efficient sparse linear algebra operations. In the Supplementary Material, we compare this inferential approach to a standard Markov chain Monte Carlo (MCMC) inference for this Bayesian model, and we experienced that INLA was more than twice as fast, while delivering approximately the same distributions; see Figure S3.

5. Simulation studies. In this section we compare the ability to capture the spatial structure of the proposed nonstationary approach in (5), which we denote as NSTAT, against a standard ANOVA with no spatial dependence (denoted as IND) as well as a spatial model with stationary SPDE in (4) (denoted as STAT). We show results for two cases: a Gaussian and a Bernoulli latent model. In Section 5.1 we outline the simulation design, in Section 5.2 we introduce the metrics used to compare the models, while in Sections 5.3 and 5.4, we show the results for the Gaussian and Bernoulli case, respectively.

5.1. Design. The simulation is designed as a one-factor model on a 15×15 regular grid, for a total of $N = 225$ locations. We assume a simplified setting of the WRF simulations with a single two-level factor, each with two replicates and with a deterministic structure over the spatial domain. Specifically, we denote our simulated response variable at location $\mathbf{s}_n$, $n = 1, \dots, N$, factor level $i \in \{0, 1\}$ and replicate $j \in \{1, 2\}$ as

$$\begin{aligned}(9) \quad Y_{ij}(\mathbf{s}_n) &\sim h(\mu_i(\mathbf{s}_n), \boldsymbol{\theta}_{\text{MRG}}), \\ g(\mu_i(\mathbf{s}_n)) &= \beta_0 + i\beta_1 I(\mathbf{s}_n \in \mathcal{D}) + \epsilon_{ij}(\mathbf{s}_n),\end{aligned}$$

where $Y_{ij}(\mathbf{s}_n)$ follows a distribution h from the exponential family, $g(\cdot)$ is the link function, the error $\epsilon_{ij}(\mathbf{s}_n) \sim \mathcal{N}(0, \sigma^2)$ is independent and identically distributed across level i , replicate j and location $\mathbf{s}_n$. We choose a constant nonzero value of β_1 for a subdomain $\mathcal{D}$ and zero outside of it. We choose four shapes for $\mathcal{D}$: square, zigzag, bar and U (see Figure S3). We report here only the results for the square, while the three other shapes are shown in the Supplementary Material. We evaluate: (1) a Gaussian model with identity link and $\beta_1 = (100, 10, 2, 1.5, 1.2, 1, 0.5, 0.3, 0.1)$, $\beta_0 = 0$ and $\sigma = 1$ and (2) a Bernoulli model with logit link with $(\beta_0, \beta_1) = ((-5, 100), (-5, 10), (-1, 2), (-1, 1), (-1, 0.5), (-1, 0.1))$ and $\sigma = 0.1$.

For each simulation, we fit the models assuming

$$(10) \quad g(\mu_i(\mathbf{s}_n)) = \beta_0 + i\beta_1(\mathbf{s}_n) + \epsilon_{ij}(\mathbf{s}_n),$$

where $\beta_1 = (\beta_1(\mathbf{s}_1), \dots, \beta_1(\mathbf{s}_N))^T$ is a Gaussian random field. IND is a standard ANOVA assuming independence across all the locations, STAT is the stationary SPDE model (4) and NSTAT is the nonstationary SPDE model obtained solving (5). A total of $n_{\text{sim}} = 1000$ simulations are performed for each shape and parameter combination.

5.2. Metrics. We assess the accuracy of the coefficient estimation using three metrics, two focused on the point estimates and one on the spatial smoothness of the estimated latent process. The first one is the proportion of 95% credibility intervals across the n_{sim} simulations containing zero for each location. Ideally, this proportion should be close to 0% inside $\mathcal{D}$ and 100% outside of it. In the Gaussian case, the second metric is the proportion of 95% credibility intervals that contains the true value of $\mu_i(\mathbf{s}_n)$ across the entire domain. The extent to which this proportion is closer to the nominal 95% level is a measure of a correct uncertainty quantification, and better results uniformly across models are expected as the signal to noise ratio of the true model (9) increases. In the case of the Bernoulli model, for the second metric we instead use the receiver operating characteristic (ROC) curve, obtained by plotting the true positive rate against the false positive rate at different thresholds. In particular, we focus on the area under the curve (AUC) of the ROC curve, which is between zero and one, with the latter representing perfect prediction.

Finally, we evaluate the model's ability to capture the smoothness of the posterior mean of the latent process by assessing the discretized gradient averaged across the four directions,

$$(11) \quad D_n = \frac{1}{|\partial\mathbf{s}_n|} \sum_{k \in \partial\mathbf{s}_n} \{\hat{\beta}(\mathbf{s}_n) - \hat{\beta}(\mathbf{s}_k)\},$$

where $\partial\mathbf{s}_n$ is the set of the nearest neighbors of $\mathbf{s}_n$, and $|\partial\mathbf{s}_n|$ is its cardinality, which is 4, 3 or 2 depending on whether the point is in the middle, at edge or on the corner of the grid, respectively. The spatially varying estimate $\hat{\beta}(\mathbf{s}_n)$ represents the posterior mean. Since the true underlying process, as specified in (9), is nonzero and constant across $\mathcal{D}$, a measure of performance of the models is the extent to which $\hat{\beta}(\mathbf{s}_n)$ have a small gradient D_n across all points in the domain but at the boundary of $\mathcal{D}$.

5.3. Results: Gaussian model. In Figure 3 we show the map of the proportion of 95% credibility intervals containing zero according to the three models for $(\beta_0, \beta_1) = (0, 2)$. The IND model in panel (a) is able to capture the change in value within the square, but the lack of spatial dependence results in noisy estimates, both inside and outside the square, as spatially closer estimates cannot borrow strength from each other. As a result, a substantial number of simulations will flag zero values inside the square, where the expected percentage is 0%. STAT and NSTAT in panels (b) and (c) both account for spatial dependence; hence, they




FIG. 3. Proportion of 95% credibility intervals containing zero for the square design. We compare the models assuming spatial (a) independence (b) stationarity and (c) nonstationarity.

avoid noisy spatial estimates and are able to capture a change in pattern at the border of the square as well as a nonzero signal inside of it.

In Figure 4(a) we show the second metric, that is, the empirical coverage of the 95% credibility intervals (averaged across the domain) as a function of the signal to noise ratio. As expected, all models show an improvement for a stronger signal, as the characterization of the latent process and its uncertainty quantification is more apparent from the data. The IND model has consistently the worst performance and is able to recover a coverage close to the nominal level only for a very strong signal. The presence of a spatial effect for STAT and NSTAT allows one instead to have a considerably closer coverage for weaker signals and consistently better results for NSTAT, given its ability to better capture the complex structure of the simulated process. Figures S4, S5 and S6 show similar results for the zigzag, bar and U shape, respectively, while Figure S7 shows similar results for all four shapes when we instead compare the predictive accuracy by means of the mean squared error.

Figure 5 shows a comparison of IND, STAT and NSTAT in terms of the discretized gradient D_n . In the case of IND in panel (a), the overall square shape is estimated, but the lack of spatial dependence results in nonsmooth estimators and occasional local discontinuities, with an associated nonzero gradient outside the boundaries of $\mathcal{D}$. The STAT model in panel (b) results instead in a smooth process and smaller gradients outside the square. The lack of flexibility in capturing local changes in the spatial structure, however, implies an overall lack




FIG. 4. (a) Proportion of 95% credibility intervals (in the Gaussian case) and (b) area under the curve (in the Bernoulli case) across the $n_{sim} = 1000$ simulations containing the true values against $(\beta_0 + \beta_1)/\sigma$ in the square design. The x axis is in the log scale.




FIG. 5. Map of the posterior mean discretized gradient D_n , as defined in (11), averaged across simulation from model (9) with $\beta = 2$ and a square design. Results are shown for the model (10) assuming: (a) IND (median value across locations 1.13, IRQ 0.029), (b) STAT (0.14, 0.06), (c) NSTAT (0.07, 0.05) (d) The map of estimated β values for NSTAT.

of ability in capturing the spatial patterns on the square and, in particular, the identification of a constant value inside it. The nonstationary model NSTAT in panel (c) is instead able to better capture the spatial structure both outside, inside and at the border of the square. Figures S7, S8 and S9 show similar results for the zigzag, bar and U shape, respectively.

5.4. Results: Bernoulli model. In Figure 4(b) we show the AUC as a function of the signal to noise ratio. For very weak signal, the three models have AUC close to 0.5, which corresponds to the predictive performance of a random guess. As the signal increases, however, all models increase their predictive ability. As in the Gaussian case, IND performs suboptimally, given its inability to borrow strength across locations in space, STAT has better predictions, and NSTAT further improves them. Figure S10 shows a comparison of D_n for IND, STAT and NSTAT with the Bernoulli simulations. As in the Gaussian case, the nonstationary model outperforms the other two by better capturing the spatial structure.

6. Application. In this section we use the proposed functional ANOVA approach to estimate how the PBL and resolution locally affect the wind characteristics in Saudi Arabia. In Section 6.1 we detail our approach to extrapolate the wind to turbine hub height. In Section 6.2 we evaluate the sensitivity of the ensemble runs with respect to PBL and resolution for both wind speed and wind energy. Section 6.3 further discusses a non-Gaussian sensitivity analysis of threshold exceedances.

6.1. *Wind speed extrapolation.* In this section we denoted by $W(s, h, t)$ the wind speed at location s , height h and time t . Since wind turbines are characterized by a hub height between 80 m to 110 m, extrapolation is necessary. Numerous literature studies have proposed models to characterize the vertical profile of wind; see [Gualtieri \(2019\)](#) for a recent review. A commonly used approach is to assume that the wind speed is directly proportional to the height through a power coefficient. The *power law* stipulates that

$$(12) \quad W(s_n, h_k, t) = W(s_n, h_r, t) \left(\frac{h_k}{h_r} \right)^{\alpha_{n,t}} e^{\eta(s_n, t)},$$

$$\eta(s_n, t) \sim \mathcal{N}(0, \sigma_{n,t}^2),$$

where h_k is the hub height to which we want to extrapolate and h_r is the reference height at which wind speed data are available. The *shear coefficient* $\alpha_{n,t}$ controls the magnitude of increase in the mean wind speed as the height increases ([Gualtieri \(2019\)](#)). In the absence of detailed meteorological observations, the shear coefficient is often assumed to be constant in space and time: $\alpha_{n,t} = \frac{1}{7}$. This choice is based on the assumption of a flat surface and neutral atmospheric conditions ([Peterson and Hennessey \(1978\)](#)). However, this simplification has been shown to be largely inappropriate in Saudi Arabia ([Crippa et al. \(2021\)](#)). To estimate the shear coefficient, our WRF simulations provide the wind speed vertical profile at six levels, at approximately equally spaced heights from 20 m to about 110 m, in addition to the wind speed at the reference height of 10 m. A simple log regression can be applied to estimate $\alpha_{n,t}$ using the six levels of wind speed for each of the four WRF runs. An example of the estimated shear coefficients, their standard deviation and the coefficients of determination R^2 for the run MYJ-6 km are displayed in Figure S11. In the mountainous Hijaz area (see Figure S1) in the west of Saudi Arabia, we observe very small estimated $\alpha_{n,t}$ values, even negative at some locations, indicating that the wind speed is on average lower at higher altitudes; this is an uncommon situation associated with large standard deviations and small R^2 values. These results suggest that the power law may not be suitable for some geographic locations, as discussed in previous studies ([Crippa et al. \(2021\)](#), [Gualtieri \(2019\)](#)). However, a previous study highlighted that the installation of a wind turbine over the complex terrain of the Hijaz region is not cost effective ([Giani et al. \(2020\)](#)). Therefore, the negative estimates in these areas are not a cause of concern in this study.

The wind speed at hub height is then converted to wind energy using turbine-specific *power curves* that transform wind speed to power. A power curve assumes a value of zero, until a minimum cut-in speed is reached, and the blades start rotating; then it keeps increasing with stronger wind speeds. Finally, the power curve reaches a cut-off speed and remains constant thereafter. We chose the turbine makes and models that were identified as the most cost effective for each location by [Giani et al. \(2020\)](#); see Figure S12. To compute the total wind energy that can be generated in each grid cell of our domain, we multiplied the power of a single turbine by the number of turbines that can fit in a specific grid cell. This number depends on the length of the turbine blades, as sufficient spacing must be provided to prevent excessive local turbulence.

6.2. *Sensitivity of wind speed and energy.* We applied the FANOVA model (1) with a Gaussian marginal distribution and an identity link to both ground wind speed and energy and assess the local sensitivity. A triangulation comprising $T = 53,387$ triangles across the domain was chosen for the SPDE approach; see Figure S13 for a plot of the triangulation in a subdomain. We chose $K = 3$ harmonics for the temporal component, as indicated by the model selection in Figures S14 and S15. The temporal trend can resolve the dynamics of the data, Figure S16 shows the boxplot of the skewness and excess kurtosis from the temporal residuals of the model for both wind speed and energy. Overall, the vast majority




FIG. 6. Posterior mean of the (a) PBL and (b) resolution coefficient of surface wind speed. Panel (c) shows the proportion of posterior variance from the PBL coefficient against that of resolution. Panel (d) shows the distribution of wind power for the 75 optimum wind farm locations in [Giani et al. \(2020\)](#) by sampling from the parameters' posterior distribution.

of locations show a symmetric distribution with no excess kurtosis, hence lending support for the Gaussian assumption. Figure S17 shows the posterior distribution of the nonspatial altitude effect β_{ALT} and the nugget effect σ^2 in the FANOVA model (1). Consistently with the diagnostics in Figure S2, β_{ALT} is positive with very high probability; hence, wind increases with altitude.

Figure 6 shows the estimated posterior mean of β_{PBL} and β_{RES} using the FANOVA model described in equation (1) for wind speed, and Figure S18 shows the corresponding results for wind energy. Figure 6(a) indicates that the ACM2 PBL scheme tends to simulate higher wind speed than the MYJ scheme over the northern portion of the domain and over the areas of complex terrain in the southwest area of Saudi Arabia. The difference between the two schemes is small over large areas of the domain, and MYJ dominates in central Saudi Arabia. Figure 6(b) shows similar spatial patterns for the simulation resolution, suggesting that both the higher resolution and the choice of the ACM2 PBL scheme result in higher winds in the same regions, some of which are characterized by complex topographic features that indeed require higher resolution and/or complex physics schemes to capture the flow patterns in these complex terrains accurately. This is expected as ACM2 combines the nonlocal and local turbulence scheme for unstable and stable conditions, respectively, unlike MYJ, which is strictly local. These results agree with prior studies showing higher accuracy of ACM2 than MYJ, also over complex terrain ([Hu, Nielsen-Gammon and Zhang \(2010\)](#), [Siuta, West and Stull \(2017\)](#)). The posterior means of the PBL and resolution for wind energy are shown in Figures S18(a) and (b), respectively. The spatial patterns are similar to those of wind speed with a generally higher posterior mean (and hence a stronger effect),

which is a feature attributable to the vertical extrapolation described in the previous section, as wind at hub height is generally higher. We also show in Figure 6(c) the ratio between the posterior variance of β_{PBL} against the total posterior variance from both effects (see Figure S18(c) for the corresponding results for wind energy), and it is readily apparent how PBL is more variable than the resolution for almost all points in the domain, except noticeably in the western coast, where a sharp transition between the Hijiaz mountain range and the sea can be better characterized by an increased resolution.

For the 75 locations indicated by [Giani et al. \(2020\)](#) as the most cost-effective for building wind farms, we can use the proposed model to assess the differences in total monthly wind power output across ensemble members. Indeed, we fix the parameters to their posterior mean, simulate 500 realizations of surface wind speed, extrapolate them and compute the power output according to the power curves. Figure 6(d) shows the boxplot of the wind power distribution across the simulations for all four ensemble members. The MYJ-9 km simulation would result in an average output of 2488.5 kW, whereas ACM2-9 km would increase this on average by 11.95 kW (0.4%), MYJ-6 km by 22.94 kW (0.9%), and ACM2-6 km by 34.89 kW (1.4%). The FANOVA, therefore, lends further support for the siting work in [Giani et al. \(2020\)](#), which was performed with MYJ-6 km (the simulation with closer validation metric with respect to some ground observations), by showing an overall robustness of the energy output with respect to both PBL and resolution.

The same analysis was performed using a standard ANOVA (i.e., assuming spatial independence) and a stationary FANOVA, as shown in Figure S19. In the case of ANOVA and in order to evaluate the impact of resolution, spatial interpolation of wind on the same is necessary, so we upscale the 6×6 km simulations to the 9×9 km grid of the other simulations in the ensemble. The stationary FANOVA model does not require any upscaling. The standard ANOVA suggests that ACM2 PBL scheme tends to simulate slower wind speed along the coast line (dark blue color in the map), and that resolution has smaller impact than the non-stationary FANOVA with near zero estimates at the majority of the locations. The stationary FANOVA results show that ACM2 PBL scheme tends to simulate high wind speeds in the northern region. With regard to resolution, we observe nonzero estimates at most of the locations. We can, therefore, conclude that the use of a nonstationary FANOVA emphasizes the role of the adopted resolution and PBL parameterization, especially in the context of complex terrains, which is consistent with findings from previous physics-based modeling studies.

We also assessed the smoothness of the latent spatial factors using the discrete Laplacian for the ANOVA model as well as the FANOVA stationary and nonstationary model. For wind speed a decrease of 33% (31% for wind energy) of the averaged discrete Laplacian was observed using FANOVA with stationary assumption, compared to the ANOVA model, and a further decrease of 25% (18%) was achieved using the nonstationary FANOVA model. Similar improvements in the smoothness were observed for wind energy data.

6.3. Sensitivity of threshold exceedances. To further assess the sensitivity of wind energy with respect to the PBL and resolution, we considered the threshold exceedance of wind power with respect to half of the maximum power output according to the power curve. The response variable $Y_{ij}(t)$, therefore, follows a Bernoulli distribution with a logit link function $g(\cdot)$ in (1). The temporal parameters $f^{(T)}(t)$ in the FANOVA model (1) were estimated independently at each location using Bayesian logistic regression. Figure S20 shows the binary values and fitted values of the Bernoulli model for the P and M locations, as indicated in Figure 1. The posterior means of the temporal effects were assumed to be fixed, and the spatial effect was estimated. Figures 7(a) and (b) show the posterior mean of β_{PBL} and β_{RES} , respectively. For the PBL coefficient, large negative values are observed in the Rub' al Khali area, indicated by the dark blue color on the map. This implies that, on average, the nonlocal ACM2 PBL scheme yields a decrease in the threshold exceedances. For the resolution




FIG. 7. Posterior mean of the (a) PBL and (b) resolution coefficient of the Bernoulli model.

coefficient, large positive values in the posterior mean are apparent in the same area as for the PBL. In other words, in the Rub' al Khali region, the choice of six km resolution tends to increase the odds of threshold exceedance. In the central Saudi Arabia and Hijaz region, both the ACM2 PBL scheme and six km resolution tend to increase the odds of threshold exceedance.

7. Conclusion. In this study we propose a model-based local sensitivity analysis of a climate ensemble, with spatially varying latent fields modeled using a nonstationary SPDE. The proposed approach allows us to capture the spatial dependence for a complex dataset on a large domain while simultaneously accounting for non-Gaussianity. Furthermore, the assumption of a continuous underlying process allows us to perform the sensitivity study for simulations with different spatial resolutions without any ad hoc upscaling. The simulation studies performed under a wide range of settings provide compelling evidence that the non-stationary model can capture more complex structures than a standard ANOVA or a stationary model. The proposed FANOVA approach is then used to provide insights about dependence of wind speed and energy for Saudi Arabia from PBL scheme and resolution.

The proposed model can be generalized to any sensitivity study consisting of spatiotemporally resolved ensembles with categorical but also quantitative input. From a methodological point of view, the assumption of a spatially varying precision and range through SPDE and basis decomposition allows us to capture complex patterns of spatial dependence while simultaneously allowing for a valid process. Alternative domain-specific approaches are also possible: if the geography suggests changes in the dependence structure dictated by descriptors, such as land/ocean, the SPDE operator can be modified to account for that with a more tailored spatial dependence structure (Fuglstad and Castruccio (2020), Hu, Fuglstad and Castruccio (2022)). Additionally, for applications where several variables are of simultaneous interest (e.g., temperature, wind and precipitation), a multivariate approach can be proposed by relying on the sparse approximations of multivariate SPDEs (Hu et al. (2013)), although the task of determining the dependence structure across both space and variables is currently limited by the dearth of sufficiently flexible models (Genton and Kleiber (2015)).

From the perspective of the application of interest in this work, our results indicate that both wind speed and energy are sensitive to the resolution and PBL scheme, with a nonlocal scheme and high resolution generally resulting in higher winds speeds over complex terrain, consistently with previous geoscience literature in other world areas. Additionally, this work indicates that the current plan for building turbines in Saudi Arabia is robust with respect to the input chosen in the ensemble, as the final output estimates change by at most 1.4% from the reference siting work (Giani et al. (2020)).

Finally, while the spatial sensitivity analysis offers scientific insights into the effect of the resolution and PBL schemes, the ultimate goal is to determine the extent to which these simulations offer an accurate representation of the true wind fields. The ensemble analyzed in this work has been validated with ground observations (Giani et al. (2020)); however, the spatial coverage of the observational network was sparse, and the diagnostic was limited to standard metrics in the geoscience literature. If a more comprehensive observational data set becomes available in the future, a more formal model-based approach can be proposed by assuming a true observational process in the FANOVA, thereby allowing both calibration and sensitivity analysis. The results from this work are useful to local policy makers in making decisions to install turbines in specific areas and hence demonstrate how our analyses can help developing countries to build renewable energy resources.

Funding. This study is based upon work supported by the King Abdullah University of Science and Technology (KAUST) Office of Sponsored Research (OSR) under Award ORFS-2022-CRG11-5069.2. Partial support for this work is provided by the National Science Foundation Grant No. DMS-2014166 to SC and by National Aeronautics and Space Administration under Grant No. 80NSSC21K1490 to SC.

SUPPLEMENTARY MATERIAL

Supplement to “Sensitivity analysis of wind energy resources with Bayesian non-Gaussian and nonstationary functional ANOVA” (DOI: [10.1214/23-AOAS1770SUPP](https://doi.org/10.1214/23-AOAS1770SUPP.pdf); .pdf). The supplement contains additional analyses and plots in support to the main findings in the paper. The code for this work is available at the following GitHub repository: github.com/Env-an-Stat-group/24.Zhang.AoAS.

REFERENCES

- AGHAHOSSEINI, A., BOGDANOV, D. and BREYER, C. (2020). Towards sustainable development in the MENA region: Analysing the feasibility of a 100% renewable electricity system in 2030. *Energy Strategy Reviews* **28** 100466.
- BAKKA, H., RUE, H., FUGLSTAD, G.-A., RIEBLER, A., BOLIN, D., ILLIAN, J., KRAINSKI, E., SIMPSON, D. and LINDGREN, F. (2018). Spatial modeling with R-INLA: A review. *Wiley Interdiscip. Rev.: Comput. Stat.* **10** e1443, 24. [MR3873676 https://doi.org/10.1002/wics.1443](https://doi.org/10.1002/wics.1443)
- BAKKA, H., VANHATALO, J., ILLIAN, J. B., SIMPSON, D. and RUE, H. (2019). Non-stationary Gaussian models with physical barriers. *Spat. Stat.* **29** 268–288. [MR3903698 https://doi.org/10.1016/j.spasta.2019.01.002](https://doi.org/10.1016/j.spasta.2019.01.002)
- BOLIN, D. and LINDGREN, F. (2011). Spatial models generated by nested stochastic partial differential equations, with an application to global ozone mapping. *Ann. Appl. Stat.* **5** 523–550. [MR2810408 https://doi.org/10.1214/10-AOAS383](https://doi.org/10.1214/10-AOAS383)
- CASTRUCCIO, S. and GUINNESS, J. (2017). An evolutionary spectrum approach to incorporate large-scale geographical descriptors on global processes. *J. R. Stat. Soc. Ser. C. Appl. Stat.* **66** 329–344. [MR3611690 https://doi.org/10.1111/rssc.12167](https://doi.org/10.1111/rssc.12167)
- CHEN, W., CASTRUCCIO, S., GENTON, M. G. and CRIPPA, P. (2018). Current and future estimates of wind energy potential over Saudi Arabia. *J. Geophys. Res., Atmos.* **123** 6443–6459.
- CRESSIE, N. and HUANG, H.-C. (1999). Classes of nonseparable, spatio-temporal stationary covariance functions. *J. Amer. Statist. Assoc.* **94** 1330–1340. [MR1731494 https://doi.org/10.2307/2669946](https://doi.org/10.2307/2669946)
- CRIPPA, P., ALIFA, M., BOLSTER, D., GENTON, M. G. and CASTRUCCIO, S. (2021). A temporal model for vertical extrapolation of wind speed and wind energy assessment. *Appl. Energy* **301** 117378.
- DINCER, I. (2018). 1.12 fossil fuels. In *Comprehensive Energy Systems* (I. Dincer, ed.) 521–567. Elsevier, Oxford.
- EDWARDS, M., CASTRUCCIO, S. and HAMMERLING, D. (2020). Marginally parameterized spatio-temporal models and stepwise maximum likelihood estimation. *Comput. Statist. Data Anal.* **151** 107018, 12. [MR4107728 https://doi.org/10.1016/j.csda.2020.107018](https://doi.org/10.1016/j.csda.2020.107018)
- EUROPEAN CENTRE FOR MEDIUM-RANGE WEATHER FORECASTS (2016). ECMWF IFS CY41r2 high-resolution operational forecasts. <https://doi.org/10.5065/D68050ZV>

- FUGLSTAD, G.-A. and CASTRUCCIO, S. (2020). Compression of climate simulations with a nonstationary global spatiotemporal SPDE model. *Ann. Appl. Stat.* **14** 542–559. [MR4117819](#) <https://doi.org/10.1214/20-AOAS1340>
- FUGLSTAD, G.-A., SIMPSON, D., LINDGREN, F. and RUE, H. (2019). Constructing priors that penalize the complexity of Gaussian random fields. *J. Amer. Statist. Assoc.* **114** 445–452. [MR3941267](#) <https://doi.org/10.1080/01621459.2017.1415907>
- GENTON, M. G. and KLEIBER, W. (2015). Cross-covariance functions for multivariate geostatistics. *Statist. Sci.* **30** 147–163. [MR3353096](#) <https://doi.org/10.1214/14-STS487>
- GIANI, P., TAGLE, F., GENTON, M. G., CASTRUCCIO, S. and CRIPPA, P. (2020). Closing the gap between wind energy targets and implementation for emerging countries. *Appl. Energy* **269** 115085.
- GUALTIERI, G. (2019). A comprehensive review on wind resource extrapolation models applied in wind energy. *Renew. Sustain. Energy Rev.* **102** 215–233.
- HASANEAN, H. and ALMAZROUI, M. (2015). Rainfall: Features and variations over Saudi Arabia, a review. *Climate* **3** 578–626.
- HU, W., FUGLSTAD, G.-A. and CASTRUCCIO, S. (2022). A stochastic locally diffusive model with neural network-based deformations for global sea surface temperature. *Stat* **11** Paper No. e431, 9. [MR4394991](#) <https://doi.org/10.1002/sta4.431>
- HU, X., SIMPSON, D., LINDGREN, F. and RUE, H. (2013). Multivariate Gaussian random fields using systems of stochastic partial differential equations. Available at [arXiv:1307.1379](#).
- HU, X.-M., NIELSEN-GAMMON, J. W. and ZHANG, F. (2010). Evaluation of three planetary boundary layer schemes in the WRF model. *J. Appl. Meteorol. Climatol.* **49** 1831–1844.
- HUANG, J. Z., KOOPERBERG, C., STONE, C. J. and TRUONG, Y. K. (2000). Functional ANOVA modeling for proportional hazards regression. *Ann. Statist.* **28** 961–999. [MR1810916](#) <https://doi.org/10.1214/aos/1015956703>
- JANJIĆ, Z. I. (1994). The step-mountain eta coordinate model: Further developments of the convection, viscous sublayer, and turbulence closure schemes. *Mon. Weather Rev.* **122** 927–945.
- JEONG, J., CASTRUCCIO, S., CRIPPA, P. and GENTON, M. G. (2018). Reducing storage of global wind ensembles with stochastic generators. *Ann. Appl. Stat.* **12** 490–509. [MR3773402](#) <https://doi.org/10.1214/17-AOAS1105>
- JEONG, J., YAN, Y., CASTRUCCIO, S. and GENTON, M. G. (2019). A stochastic generator of global monthly wind energy with Tukey g -and- h autoregressive processes. *Statist. Sinica* **29** 1105–1126. [MR3932511](#)
- JONES, C., GIORGI, F. and ASRAR, G. (2011). The coordinated regional downscaling experiment: Cordex—an international downscaling link to cmip5. *CLIVAR Exchanges* **16** 34–40.
- JUN, M. and STEIN, M. L. (2008). Nonstationary covariance models for global data. *Ann. Appl. Stat.* **2** 1271–1289. [MR2655659](#) <https://doi.org/10.1214/08-AOAS183>
- KAUFMAN, C. G. and SAIN, S. R. (2010). Bayesian functional ANOVA modeling using Gaussian process prior distributions. *Bayesian Anal.* **5** 123–149. [MR2596438](#) <https://doi.org/10.1214/10-BA505>
- LINDGREN, F., RUE, H. and LINDSTRÖM, J. (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: The stochastic partial differential equation approach. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **73** 423–498. [MR2853727](#) <https://doi.org/10.1111/j.1467-9868.2011.00777.x>
- MOHSENI-CHERAGHLOU, A. (2021). Fossil fuel subsidies and renewable energies in mena: An oxymoron? <https://www.mei.edu/publications/fossil-fuel-subsidies-and-renewable-energies-mena-oxymoron>.
- NREP (2018). Saudi arabia renewable energy targets and long term visibility. national renewable energy program.
- NURUNNABI, M. (2017). Transformation from an oil-based economy to a knowledge-based economy in Saudi Arabia: The direction of saudi vision 2030. *J. Knowl. Econ.* **8** 536–64.
- PETERSON, E. W. and HENNESSEY, J. P. (1978). On the use of power laws for estimates of wind power potential. *J. Appl. Meteorol.* **17** 390–394.
- PLEIM, J. E. (2007). A combined local and nonlocal closure model for the atmospheric boundary layer. Part I: Model description and testing. *J. Appl. Meteorol. Climatol.* **46** 1383–1395.
- QU, Z., DAI, W. and GENTON, M. G. (2021). Robust functional multivariate analysis of variance with environmental applications. *Environmetrics* **32** Paper No. e2641, 23. [MR4207556](#) <https://doi.org/10.1002/env.2641>
- REN21 SECRETARIAT (2021). Renewables 2021 – global status report. Paris, France.
- RENEWABLES, R. (2018). Global status report (Paris: Ren21 secretariat).
- RUE, H., MARTINO, S. and CHOPIN, N. (2009). Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **71** 319–392. [MR2649602](#) <https://doi.org/10.1111/j.1467-9868.2008.00700.x>
- SAIN, S. R., NYCHKA, D. and MEARNES, L. (2011). Functional ANOVA and regional climate experiments: A statistical analysis of dynamic downscaling. *Environmetrics* **22** 700–711. [MR2843137](#) <https://doi.org/10.1002/env.1068>

- SIUTA, D., WEST, G. and STULL, R. (2017). WRF hub-height wind forecast sensitivity to pbl scheme, grid length, and initial condition choice in complex terrain. *Weather Forecast.* **32** 493–509.
- SKAMAROCK, W. C., KLEMP, J. B., DUDHIA, J., GILL, D. O., LIU, Z., BERNER, J., WANG, W., POWERS, J. G., DUDA, M. G. et al. (2019). *A Description of the Advanced Research Wrf Model Version 4*. National Center for Atmospheric Research, Boulder, CO, USA.
- STEIN, M. L. (1999). *Interpolation of Spatial Data: Some Theory for Kriging*. Springer Series in Statistics. Springer, New York. MR1697409 <https://doi.org/10.1007/978-1-4612-1494-6>
- STONE, C. J., HANSEN, M. H., KOOPERBERG, C. and TRUONG, Y. K. (1997). Polynomial splines and their tensor products in extended linear modeling. *Ann. Statist.* **25** 1371–1470. MR1463561 <https://doi.org/10.1214/aos/1031594728>
- SUN, Y. and GENTON, M. G. (2012). Functional median Polish. *J. Agric. Biol. Environ. Stat.* **17** 354–376. MR2993271 <https://doi.org/10.1007/s13253-012-0096-8>
- TAGLE, F., CASTRUCCIO, S., CRIPPA, P. and GENTON, M. G. (2019). A non-Gaussian spatio-temporal model for daily wind speeds based on a multi-variate skew- t distribution. *J. Time Series Anal.* **40** 312–326. MR3946155 <https://doi.org/10.1111/jtsa.12437>
- TARRÍO-SAAVEDRA, J., NAYA, S., FRANCISCO-FERNÁNDEZ, M., ARTIAGA, R. and LOPEZ-BECEIRO, J. (2011). Application of functional anova to the study of thermal stability of micro–nano silica epoxy composites. *Chemom. Intell. Lab. Syst.* **105** 114–124.
- ULLAH, S. and FINCH, C. (2013). Applications of functional data analysis: A systematic review. *BMC Med. Res. Methodol.* **13**.
- WHITTLE, P. (1954). On stationary processes in the plane. *Biometrika* **41** 434–449. MR0067450 <https://doi.org/10.1093/biomet/41.3-4.434>
- WHITTLE, P. (1963). On the fitting of multivariate autoregressions, and the approximate canonical factorization of a spectral density matrix. *Biometrika* **50** 129–134. MR0161430 <https://doi.org/10.1093/biomet/50.1-2.129>
- YUE, Y., BOLIN, D., RUE, H. and WANG, X.-F. (2019). Bayesian generalized two-way ANOVA modeling for functional data using INLA. *Statist. Sinica* **29** 741–767. MR3931386
- YUE, Y. and SPECKMAN, P. L. (2010). Nonstationary spatial Gaussian Markov random fields. *J. Comput. Graph. Statist.* **19** 96–116. MR2654402 <https://doi.org/10.1198/jcgs.2009.08124>
- ZHANG, J., CRIPPA, P., GENTON, M. G. and CASTRUCCIO, S. (2021). Assessing the reliability of wind power operations under a changing climate with a non-Gaussian bias correction. *Ann. Appl. Stat.* **15** 1831–1849. MR4355078 <https://doi.org/10.1214/21-aos1460>
- ZHANG, J.-T., CHENG, M.-Y., WU, H.-T. and ZHOU, B. (2019). A new test for functional one-way ANOVA with applications to ischemic heart screening. *Comput. Statist. Data Anal.* **132** 3–17. MR3913131 <https://doi.org/10.1016/j.csda.2018.05.004>
- ZHANG, Y., HODGES, J. S. and BANERJEE, S. (2009). Smoothed ANOVA with spatial effects as a competitor to MCAR in multivariate spatial smoothing. *Ann. Appl. Stat.* **3** 1805–1830. MR2752159 <https://doi.org/10.1214/09-AOAS267>
- ZHANG, J., CRIPPA, P., GENTON, M. G. and CASTRUCCIO, S. (2024). Supplement to “Sensitivity analysis of wind energy resources with Bayesian non-Gaussian and nonstationary functional ANOVA.” <https://doi.org/10.1214/23-AOAS1770SUPP>