



Castro-Camilo, D., Huser, R. and Rue, H. (2019) A spliced Gamma-Generalized Pareto model for short-term extreme wind speed probabilistic forecasting. *Journal of Agricultural, Biological and Environmental Statistics*, 24(3), pp. 517-534.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

<http://eprints.gla.ac.uk/190278/>

Deposited on: 11 September 2019

Enlighten – Research publications by members of the University of Glasgow  
<http://eprints.gla.ac.uk>

# A spliced Gamma-Generalized Pareto model for short-term extreme wind speed probabilistic forecasting

Daniela Castro-Camilo<sup>1</sup>, Raphaël Huser<sup>2</sup>, and Håvard Rue<sup>2</sup>

June 17, 2019

## Abstract

Renewable sources of energy such as wind power have become a sustainable alternative to fossil fuel-based energy. However, the uncertainty and fluctuation of the wind speed derived from its intermittent nature bring a great threat to the wind power production stability, and to the wind turbines themselves. Lately, much work has been done on developing models to forecast average wind speed values, yet surprisingly little has focused on proposing models to accurately forecast extreme wind speeds, which can damage the turbines. In this work, we develop a flexible spliced Gamma-Generalized Pareto model to forecast extreme and non-extreme wind speeds simultaneously. Our model belongs to the class of latent Gaussian models, for which inference is conveniently performed based on the integrated nested Laplace approximation method. Considering a flexible additive regression structure, we propose two models for the latent linear predictor to capture the spatio-temporal dynamics of wind speeds. Our models are fast to fit and can describe both the bulk and the tail of the wind speed distribution while producing short-term extreme and non-extreme wind speed probabilistic forecasts.

**Keywords:** Extreme value theory, Threshold-based inference, Latent Gaussian models, INLA, SPDE, Wind speed forecasting.

---

<sup>1</sup>School of Mathematics and Statistics, University of Glasgow, Glasgow G12 8QQ, UK. E-mail: daniela.castro.camilo@gmail.com.

<sup>2</sup>Computer, Electrical and Mathematical Sciences and Engineering (CEMSE) Division, King Abdullah University of Science and Technology (KAUST), Thuwal 23955-6900, Saudi Arabia. E-mails: raphael.huser@kaust.edu.sa, haavard.rue@kaust.edu.sa.

# 1 Introduction

The integration of renewable energies such as wind power into power systems has been developing on a large scale around the world in the last decade. Wind power generation is sustainable, emission-free, and its cost is nearly the same as that of coal or nuclear energy (Hering and Genton, 2010). These advantages are counter-balanced by several challenges, such as high variability, limited dispatchability, and non-storability (Pinson and Madsen, 2012; Hering and Genton, 2010). Accurate short-term forecasts of wind power are therefore crucial for power production planning and risk assessment.

Wind power forecasts can be made directly if power data are available (Lenzi *et al.*, 2017). However, wind speed forecasting can be more precise than wind power forecasting due to the spatial correlation of wind, and wind power forecasts can be easily obtained from the forecasted wind speeds (see, e.g., the deterministic cubic power curve in Figure 5 in Hering and Genton, 2010). In this paper, we focus on wind speed probabilistic forecasting based on 20 turbine towers measuring hourly average wind speed and wind direction. The towers are installed at the border between Oregon and Washington, along the Columbia River (see Figure 1). Each station encompasses between  $T = 21,306$  and  $T = 26,304$  hourly measurements of non-zero wind speed from January 2012 to December 2014. Basic exploratory analyses unveil different wind regimes among stations, high autocorrelation within stations, seasonal patterns, and persistence, which refers to the variable's tendency to maintain its current state. See Section 2 of the Supplementary Material for accompanying graphical results.

Over the last decades, statistical models have shown to be very effective in capturing the fluctuating characteristics of wind speed to produce accurate short-term wind speed forecasts (see, e.g., Zhu and Genton, 2012). Purely temporal models are built assuming that wind speed at each time point is partially predicted by wind speed in its near past. In this context, different time series models have been proposed for short-term forecasts, such as autoregressive

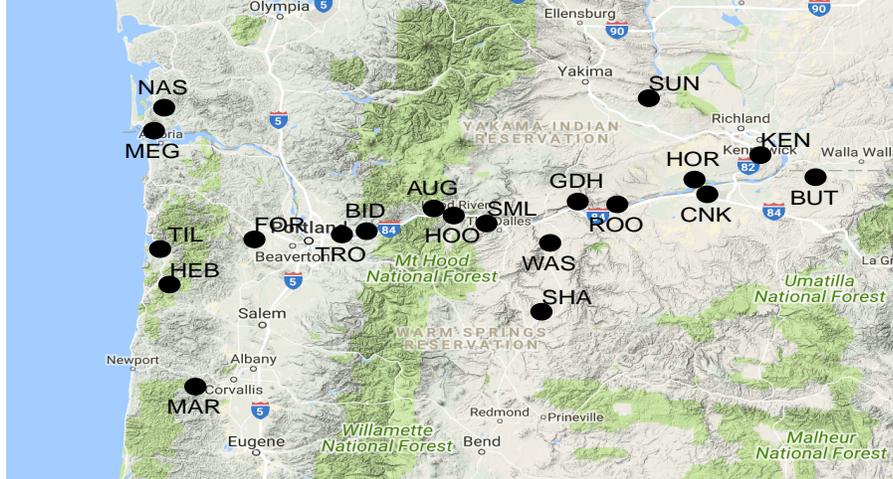


Figure 1: Map showing the 20 towers along the Columbia River.

models (Huang and Chalabi, 1995), autoregressive moving average models (Erdem and Shi, 2011), and autoregressive integrated moving average models (Palomares-Salas *et al.*, 2009). Wind speed is modeled using neural networks in Li and Shi (2010), while a model to assess wind potential using spectral analysis is proposed in Shih (2008). Models that incorporate both temporal and spatial correlations in the form of off-site information (i.e., information that is not collected at the site of interest, but at neighboring sites) have been found to increase accuracy over conventional time series models. In this framework, Alexiadis *et al.* (1999) use off-site predictors to improve wind speed and wind power forecasts, while Gneiting *et al.* (2006), Hering and Genton (2010), and Kazor and Hering (2015) assume that wind speeds follow a truncated normal distribution with regime-dependent mean and variance. See Zhu and Genton (2012) for a review on statistical wind speed forecasting models.

Our main goal is to develop a flexible space-time model designed to forecast future wind speeds at the turbine locations, with a particular focus on extreme wind speeds. The statistical modeling of extremes is challenging, as it usually focuses on the estimation of high (or low) quantiles, where the available data are limited. To tackle this, we rely on extreme-value theory, which offers a rigorous mathematical framework to develop techniques and models for

describing the tail of the distribution; see [Davison and Huser \(2015\)](#) for a review.

In the context of wind-powered energy, a prolonged period of large wind speed values may pose a considerable risk to the wind turbines, thus affecting wind energy production. In this context, a suitable definition of extremes is given by the *wind speed exceedances*, i.e., wind speeds that exceed a certain high threshold  $u > 0$ . The stochastic behavior of exceedances is described by the conditional probability  $\Pr(Y > u + y \mid Y > u) = \{1 - F(u + y)\} / \{1 - F(u)\}$ ,  $y > 0$ , where  $Y$  has distribution  $F$ . Under broad conditions ([Davison and Smith, 1990](#)), it can be shown that as  $u$  becomes large, this conditional probability can be approximated by  $1 - H_{GP}(y) = (1 + \xi y / \sigma_u)_+^{-1/\xi}$ , where  $H_{GP}$  denotes the generalized Pareto (GP) distribution function with scale parameter  $\sigma_u > 0$ , shape parameter  $\xi \in \mathbb{R}$ , and  $a_+ = \max(0, a)$ . As  $\xi \rightarrow 0$ , the GP distribution corresponds to the exponential distribution function  $1 - \exp(-y/\sigma_u)$ ,  $y > 0$ . The shape parameter determines the weight of the upper tail: a heavy tail is obtained with  $\xi > 0$ , a light tail with  $\xi \rightarrow 0$ , and a bounded tail with  $\xi < 0$ .

In this paper, we propose a hierarchical Bayesian model based on the spliced Gamma-Generalized Pareto (Gamma-GP) distribution to describe both non-extreme and extreme wind speeds in space and time. Our model is designed to provide upper tail forecasts beyond observed wind speed values. Hierarchical Bayesian models for threshold exceedances were introduced by [Casson and Coles \(1999\)](#) and [Cooley et al. \(2007\)](#) and fitted using expensive Markov chain Monte Carlo (MCMC) methods. Spliced extreme-value and alternative extended GP models were proposed by various authors ([Tancredi et al., 2006](#); [Naveau et al., 2016](#)) to model extreme and non-extreme data simultaneously; see, [Scarrott and MacDonald \(2012\)](#) for a review and [Opitz et al. \(2018\)](#) for a recent related contribution. We here explore two different latent processes that drive the space-time trends and dependence structure. The first one is a temporal model that incorporates spatial information in the form of off-site predictors, while the second one is a space-time model. As our Bayesian approach relies on latent Gaussian

processes, we can exploit the integrated nested Laplace approximation (INLA; [Rue et al., 2009](#)) of posterior distributions, conveniently implemented in the R-INLA package, which provides fast and accurate inference ([Rue et al., 2009](#)) for this class of models.

The remainder of the article is organized as follows. In [Section 2](#) we develop our modeling strategy, while in [Section 3](#) we explain our Bayesian estimation approach using INLA. [Section 4](#) presents the results of our modeling approaches in terms of wind speed forecasting. Conclusions and an outlook towards future research are given in [Section 5](#).

## 2 Modeling wind speeds using latent Gaussian models

### 2.1 General framework

Latent Gaussian models are a broad and flexible class of models, which are well-suited for modeling (possibly non-Gaussian and non-stationary) spatio-temporal data ([Rue et al., 2009](#)). They admit a hierarchical model formulation, whereby observations  $\mathbf{y}$  (with components  $y(i), i \in \mathcal{I}$ , for some index set  $\mathcal{I}$ ) are typically assumed to be conditionally independent given a latent Gaussian random field  $\mathbf{x}$  (with components  $x(i), i \in \mathcal{I}$ ) and hyperparameters  $\boldsymbol{\theta}_1$ , i.e.,

$$\begin{aligned} \mathbf{y} \mid \mathbf{x}, \boldsymbol{\theta}_1 &\sim \prod_{i \in \mathcal{I}} p(y(i) \mid x(i), \boldsymbol{\theta}_1), \\ \mathbf{x} \mid \boldsymbol{\theta}_2 &\sim \mathcal{N}(\boldsymbol{\mu}_{\boldsymbol{\theta}_2}, \mathbf{Q}_{\boldsymbol{\theta}_2}^{-1}), \\ \boldsymbol{\theta} = (\boldsymbol{\theta}_1^T, \boldsymbol{\theta}_2^T)^T &\sim p(\boldsymbol{\theta}), \end{aligned} \tag{2.1}$$

where  $p$  denotes a generic distribution. The latent Gaussian random field  $\mathbf{x}$  has mean vector  $\boldsymbol{\mu}_{\boldsymbol{\theta}_2}$  and precision matrix  $\mathbf{Q}_{\boldsymbol{\theta}_2}$ , which are controlled by the hyperparameters  $\boldsymbol{\theta}_2$ . It describes the trends and the underlying dependence structure of the data, and its specification in [\(2.1\)](#) is key to generate a flexible and versatile class of models. We here assume that  $y(i)$  only depends on a linear predictor  $\eta(i)$  which has an additive structure with respect to some fixed covariates

and random effects, i.e.,

$$\eta(i) = \mu + \sum_{j=1}^J \beta_j z_j(i) + \sum_{k=1}^K f_k(w_k(i)), \quad i \in \mathcal{I}. \quad (2.2)$$

Here,  $\mu$  is the overall intercept,  $z_j(i)$  are known covariates with coefficients  $\boldsymbol{\beta} = (\beta_1, \dots, \beta_J)^T$ , and  $\mathbf{f} = \{f_1(\cdot), \dots, f_K(\cdot)\}$  are specific (a priori independent) Gaussian processes defined in terms of a set of covariates  $\mathbf{w} = (w_1(i), \dots, w_K(i))^T$ . If we further assume that  $(\mu, \boldsymbol{\beta}^T)^T$  has a Gaussian prior, the joint distribution of  $\mathbf{x} = (\boldsymbol{\eta}^T, \mu, \boldsymbol{\beta}^T, \mathbf{f}^T)^T$  is Gaussian, and yields the latent field  $\mathbf{x}$  in the hierarchical formulation (2.1); see e.g., Rue *et al.* (2017) for more details. Note that in Section 2.3.1 the index set  $\mathcal{I}$  denotes time, while in Section 2.3.2, it denotes space and time.

The remainder of this section is organized following the hierarchical representation in (2.1). Therefore, Section 2.2 presents our model likelihood (with vector of parameters  $\boldsymbol{\theta}_1$ ), Section 2.3 introduces our two proposed latent structures, called the off-site model and the SPDE model (with vector of parameters  $\boldsymbol{\theta}_{2,\text{off-site}}$  and  $\boldsymbol{\theta}_{2,\text{SPDE}}$ , respectively), and Section 2.4 specifies prior distributions for all the parameters involved, namely  $\boldsymbol{\theta}_1$ ,  $\boldsymbol{\theta}_{2,\text{off-site}}$ , and  $\boldsymbol{\theta}_{2,\text{SPDE}}$ .

## 2.2 A spliced Gamma-GP model

Here we describe the top layer of the hierarchical representation in (2.1), namely, the model likelihood. Loosely speaking, our model assumes that wind speeds under a certain threshold follow a Gamma distribution, while those above the threshold follow a generalized Pareto (GP) distribution. The (non-stationarity) threshold is estimated first using Gamma quantile regression, and it is then used subsequently to estimate the proportion of observations above the threshold using a Bernoulli distribution, as well as to fit the GP distribution. For ease of exposition, we describe our model likelihood in three stages corresponding to the Gamma, Bernoulli, and GP models, respectively.

**Stage 1.** We assume that positive wind speeds  $y(i) > 0$  are characterized by a Gamma distribution parametrized in terms of an  $\alpha$ -quantile  $\psi_\alpha(i) > 0$  and a precision parameter  $\kappa > 0$ .

With this parametrization, the Gamma density is

$$h_{\text{Ga}}(y) = \frac{1}{\Gamma(\kappa)} \left\{ \frac{H_{\text{Ga}}^{-1}(\alpha; \kappa, 1)}{\psi_\alpha(i)} \right\}^\kappa y^{\kappa-1} \exp \left\{ -y \frac{H_{\text{Ga}}^{-1}(\alpha; \kappa, 1)}{\psi_\alpha(i)} \right\},$$

where  $H_{\text{Ga}}^{-1}(\alpha; \kappa, 1)$  is the quantile function of the Gamma distribution with unit rate and shape  $\kappa$ , evaluated at the probability  $\alpha \in (0, 1)$ .

**Stage 2.** The value  $\alpha$  corresponds to the probability that wind speeds exceed the threshold  $\psi_\alpha(i)$ , conditional on being positive. Here we estimate the unconditional probability of exceeding  $\psi_\alpha(i)$  to account for the few zero wind speed values that were previously excluded in Stage 1. Writing  $y(i)$  as the  $i$ -th wind speed observation, we model exceedance indicators  $\mathbf{1}\{y(i) > \psi_\alpha(i)\}$  using the Bernoulli distribution  $h_{\text{B}}(y) = p(i)^y \{1 - p(i)\}^{1-y}$ ,  $y \in \{0, 1\}$ , where  $p(i) = \Pr\{y(i) > \psi_\alpha(i)\}$ , and  $\psi_\alpha(i)$  has been previously estimated in Stage 1.

**Stage 3.** Since the tail of the Gamma distribution decays exponentially fast, probabilities associated with extreme events might be underestimated if the true underlying distribution is heavy-tailed. To more flexibly model the probabilities associated with extreme events, we correct the tail using a GP distribution; recall Section 1. Threshold exceedances defined as  $x(i) = \{y(i) - \psi_\alpha(i)\} \mid y(i) > \psi_\alpha(i)$ , are characterized by a GP distribution parametrized in terms of a  $\beta$ -quantile  $\phi_\beta(i) > 0$  and a shape parameter  $\xi \geq 0$ . With this parametrization, the GP distribution is

$$H_{\text{GP}}(y) = \begin{cases} 1 - [1 + \{(1 - \beta)^{-\xi} - 1\}y/\phi_\beta(i)]^{-1/\xi}, & \text{if } \xi > 0, \\ 1 - (1 - \beta)^{y/\phi_\beta(i)}, & \text{if } \xi = 0, \end{cases} \quad y > 0.$$

Note that the goal of the Gamma distribution in Stage 1 is twofold: to describe the distribution of non-extreme wind speeds and to obtain a suitable threshold  $\psi_\alpha(i)$  to define wind speed exceedances. The second and third stages model the frequency and intensity of extremes, respectively, and are connected to the first stage through the threshold  $\psi_\alpha(i)$ .

## 2.3 Latent structure

Here we detail two approaches to describe the middle layer of the hierarchical representation in (2.1), namely, the latent Gaussian process specification. The first approach, called the off-site latent model, corresponds to a temporal model fitted at each station separately, with off-site information included in the form of covariates. In this case,  $\mathcal{I} \equiv \mathcal{T} = \{1, \dots, T\}$  (with  $T$  the number of time points) and we denote by  $y_{\mathbf{s}}(t)$  the wind speed at some fixed location  $\mathbf{s} \in \mathcal{S} \subset \mathbb{R}^2$  and time  $t \in \mathcal{T}$ . The second approach, called the SPDE latent model, corresponds to a proper space-time model stemming from a particular stochastic partial differential equation (SPDE). In this case,  $\mathcal{I} \equiv \mathcal{S} \times \mathcal{T}$  and we denote by  $y(\mathbf{s}, t)$  the wind speed at location  $\mathbf{s} \in \mathcal{S}$  and time  $t \in \mathcal{T}$ .

### 2.3.1 Off-site latent model

For each fixed location  $\mathbf{s} \in \mathcal{S}$ , we propose the following linear predictor:

$$\eta_{\mathbf{s}}^{(1)}(t) = \mu_{\mathbf{s}} + \sum_{j=1}^{|\mathcal{N}_{\mathbf{s}}|} \beta_j y_{\mathbf{s}_j}(t-1) + f_1(t; \rho_{\mathbf{s},1}, \tau_{\mathbf{s},1}) + f_2(w_2(t); \tau_{\mathbf{s},2}), \quad t = 1, \dots, T, \quad (2.3)$$

where for each  $j \in \{1, \dots, |\mathcal{N}_{\mathbf{s}}|\}$ ,  $y_{\mathbf{s}_j}(t-1)$  is the lagged time series of wind speeds at the  $j$ -th neighbor of  $\mathbf{s}$ ,  $\mathcal{N}_{\mathbf{s}}$  is the set of neighbors of  $\mathbf{s}$  of cardinality  $|\mathcal{N}_{\mathbf{s}}|$ . The coefficients  $\{\beta_j\}$  (fixed effects) quantify the effect that wind speeds at the  $j$ -th neighbor observed at time lag one have on the response. The random effects  $f_1(t) \equiv f_1(t; \rho_{\mathbf{s},1}, \tau_{\mathbf{s},1})$  and  $f_2(w(t)) \equiv f_2(w_2(t); \tau_{\mathbf{s},2})$  account for unobserved heterogeneity within each station. Specifically,  $f_1(t)$  captures the autocorrelation structure of order one within each wind speed time series, and is assumed to be described by a zero-mean autoregressive Gaussian process of first order, i.e.,

$$\begin{aligned} f_1(1) &\sim \mathcal{N}(0, \{\tau_{\mathbf{s},1}(1 - \rho_{\mathbf{s},1}^2)\}^{-1}), \\ f_1(t) &= \rho_{\mathbf{s},1} f_1(t-1) + \epsilon(t), \quad |\rho_{\mathbf{s},1}| < 1, \quad \epsilon(t) \sim \mathcal{N}(0, \tau_{\mathbf{s},1}^{-1}), \quad t = 2, \dots, T, \end{aligned}$$

where  $\tau_{s,1} > 0$  is a precision parameter. Our experiments show that including greater lags do not significantly improve the fit or predictive power of our model. The random effect  $f_1(t)$  along with the lagged covariates account for temporal dependence and persistence.

The random effect  $f_2(w(t))$  represents the sub-daily variation of wind speeds, and is similar in spirit to the semiparametric splines of [Hering \*et al.\* \(2015\)](#). It is assumed to be described by a cyclic Gaussian random walk of second order with precision  $\tau_{s,2} > 0$ , defined over each of the 24 hours within a day ([Rue and Held, 2005](#), Ch. 3). Let  $w_2(t) \in \{1, \dots, 24\}$  denote the hour associated to time  $t$ , then:

$$f_2(w_2(t)) - 2f_2(w_2(t+1)) + f_2(w_2(t+2)) \sim \mathcal{N}(0, \tau_{s,2}^{-1}), \quad t = 1, \dots, T-2.$$

Both random effects are constrained to sum to zero, and their precision hyperparameters  $\tau_{s,l}$ ,  $l = 1, 2$ , have the main goal of controlling the strength of dependence among neighboring covariate classes, i.e., the smoothness of the corresponding random effect.

The parameters of our spliced Gamma-GP model are linked to different linear predictors of the form (2.3). For consistency with the off-site model notation, we now write  $\psi_{s,\alpha}(t)$ ,  $p_s(t)$ , and  $\phi_{s,\beta}(t)$  to indicate time-varying but location-specific Gamma  $\alpha$ -quantile, Bernoulli probability, and GP  $\beta$ -quantile, respectively, and  $\kappa_s$  and  $\xi_s$  to indicate time-constant, but location-specific Gamma and GP shape parameters, respectively. The time-varying parameters are finally linked to the off-site latent model as follows:

$$\begin{aligned} \text{Stage 1:} \quad & \psi_{s,\alpha}(t) = \exp\{\eta_{s,\text{Gamma}}^{(1)}(t)\}, \quad \text{where} \quad \Pr\{y_s(t) \leq \psi_{s,\alpha}(t)\} = \alpha, \\ \text{Stage 2:} \quad & p_s(t) = \exp\{\eta_{s,\text{Ber}}^{(1)}(t)\} / [1 + \exp\{\eta_{s,\text{Ber}}^{(1)}(t)\}], \\ \text{Stage 3:} \quad & \phi_{s,\beta}(t) = \exp\{\eta_{s,\text{GP}}^{(1)}(t)\}, \quad \text{where} \quad \Pr\{y_s(t) \leq \phi_{s,\beta}(t)\} = \beta. \end{aligned} \quad (2.4)$$

### 2.3.2 SPDE latent model

The linear predictor in (2.3) corresponds to a temporal model that introduces spatial information using lagged off-site predictors. Here we describe a proper space-time model, which

explicitly accounts for spatial dependence amongst wind speeds at different stations. Specifically, we assume that the space-time dependence between wind speeds at different wind towers can be described by a spatio-temporal term  $u(\mathbf{s}, t)$  that varies in time according to a first-order autoregressive structure. Specifically, we assume that  $u(\mathbf{s}, t) = \rho_2 u(\mathbf{s}, t - 1) + z(\mathbf{s}, t)$ , where  $|\rho_2| < 1$ , and  $z(\mathbf{s}, t)$  is a zero-mean, temporally independent Gaussian field, that is completely determined by a stationary Matérn covariance function with marginal variance  $\sigma^2 > 0$ , range  $r > 0$ , and fixed smoothness parameter. This gives rise to our second linear predictor:

$$\eta^{(2)}(\mathbf{s}, t) = \mu + u(\mathbf{s}, t) + f_2(w_2(t); \tau_2), \quad \mathbf{s} \in \mathcal{S}, t = 1, \dots, T, \quad (2.5)$$

where  $\mu$  is an intercept,  $f_2(w_2(t)) \equiv f_2(w_2(t); \tau_2)$  is the cyclic random effect described in Section 2.3.1, capturing the sub-daily variations that are common to all the stations, and  $\tau_2 > 0$  is its precision parameter. By contrast with (2.3), the notation  $\eta^{(2)}(\mathbf{s}, t)$  in (2.5) emphasizes that this second linear predictor is a function of space and time.

The latent process associated with (2.5) has dense covariance and (possibly) precision matrices, which implies that any attempt to make Bayesian inference can be computationally demanding. To avoid this issue, we use the Stochastic Partial Differential Equations (SPDE) approach introduced by Lindgren *et al.* (2011). The SPDE approach consists in constructing a continuous approximation to the Gaussian field  $z(\mathbf{s}, t)$  by using a continuous SPDE latent model defined over the entire study area. It can be shown that this SPDE has a Gaussian field with a Matérn covariance function as stationary solution. Under certain conditions (see, e.g., Lindgren *et al.*, 2011), the continuous SPDE solution has the Markovian property. This property produces sparse precision matrices that can be easily factorized, and are the focus of the INLA methodology. An approximate discrete solution of the SPDE in a bounded domain, defined in our case by the location of the wind towers, can be obtained using a finite element method that allows for flexible boundaries and different levels of accuracy for the discretization; see Section 3 of the Supplementary Material for details on the mesh used to discretize the study

region. For more details on spatial modeling using the SPDE approach, see the recent review by Bakka *et al.* (2018).

Because the Cascade Mountains illustrated in Figure 1 naturally divide the study region into two sub-regions, we consider two independent sub-models for  $u(\mathbf{s}, t)$ , one for stations to the East and another for stations to the West of the mountains. We fit each set of stations separately but using the same latent structure specified in (2.5).

The link between this new linear predictor and the model likelihood in Section 2.2 is the same as in (2.4), with the suitable change of notation. Note that by pooling stations together, the space-time parameters  $\psi_\alpha(\mathbf{s}, t)$  and  $\phi_\beta(\mathbf{s}, t)$  vary in space and time, whereas the Gamma and GP shape parameters  $\kappa$  and  $\xi$ , respectively, are now fixed across locations to limit the number of hyperparameters to be estimated using INLA.

## 2.4 Prior specification

Here we describe the bottom layer of the hierarchical representation in (2.1), namely, the specification of prior distributions. We need to specify priors for the likelihood hyperparameters  $\boldsymbol{\theta}_1$  ( $\boldsymbol{\theta}_1 = \kappa$  for the Gamma model and  $\boldsymbol{\theta}_1 = \xi$  for the GP model; recall that  $\kappa = \kappa_{\mathbf{s}}$  and  $\xi = \xi_{\mathbf{s}}$  for the off-site model), as well as for the hyperparameters of the two latent structures described in Section 2.3, namely  $\boldsymbol{\theta}_{2,\text{off-site}} = (\rho_{\mathbf{s},1}, \tau_{\mathbf{s},1}, \tau_{\mathbf{s},2})^T$  and  $\boldsymbol{\theta}_{2,\text{SPDE}} = (\sigma^2, r, \rho_2, \tau_2)^T$ .

When little expert knowledge is available, a common practice is to assume non-informative priors. Alternatively, informative priors can be proposed using Penalized Complexity (PC) priors (Simpson *et al.*, 2017). In this framework, model components are considered to be flexible extensions of simpler base models. Priors are then developed in such a way that the components shrink towards their base models, thus preventing overfitting. Simpson *et al.* (2017) propose to use the Kullback-Leibler divergence to measure the squared “distance” from the base model to its flexible extension, and to penalize this “distance” at constant rate.

For our likelihood hyperparameters, we assume a slightly informative prior over the Gamma shape  $\kappa$ , by considering a Gamma distribution with shape 10 and rate 1, which gives a high probability to values between 5 and 15. A moderately strong PC prior is assumed for the shape parameter of the GP distribution  $\xi$ ; since large values of the shape parameter are usually unrealistic for wind speeds, we here assume that  $\Pr(\xi > 0.4) \approx 0.01$ .

Regarding the hyperparameters of the off-site latent model, we assume fairly informative PC priors for the correlation hyperparameter of the AR(1) process, and the precision hyperparameter of the random walk of order 2. Specifically,  $\Pr(\rho_{\mathbf{s},1} > 0.9) = 0.95$  and  $\Pr(1/\sqrt{\tau_{\mathbf{s},2}} > \text{sd}_{\text{wind}}) = 0.01$ , where  $\text{sd}_{\text{wind}}$  denotes the empirical standard deviation of the temporally aggregated wind speeds. A diffuse prior is assumed for  $\tau_{\mathbf{s},1}$ , the precision hyperparameter of the AR(1) component.

PC priors on the parameters of the Gaussian field in the SPDE latent model, namely the marginal variance  $\sigma^2 > 0$  and the range of dependence  $r > 0$ , are chosen such that the variance is shrunk towards zero, whereas the range is shrunk towards infinity (Fuglstad *et al.*, 2018). Specifically, we set  $\Pr(\sigma > 2 \times \text{sd}_{\text{wind}}) = 0.01$  and  $\Pr(r < r_{\text{median}}) = 0.5$ , where  $r_{\text{median}}$  is the median of the distances between stations. For stations to the East of the Cascade Mountains,  $r_{\text{median}} = 94.6$  km, and for stations to the West,  $r_{\text{median}} = 113.3$  km. A PC prior is also chosen for the correlation coefficient of the autoregressive term in  $u(\mathbf{s}, t)$ , specifically  $\Pr(\rho_2 > 0.9) = 0.95$ . The PC prior for  $\tau_2$  is the same as for  $\tau_{\mathbf{s},2}$  in the off-site latent model.

In general, our prior specification tries to reflect some characteristics observed in the data. As part of the model selection, we conducted a small sensitivity analysis using a subset of the data. The results show that the prior specification does not affect the results considerably.

### 3 Inference based on INLA

Here we describe the form of the joint posterior distribution for each stage of our spliced Gamma-GP model detailed in Section 2. In the following, we reuse the generic notation in Section 2.1 for ease of exposition. Let  $\mathbf{y}$  denote the vector of observations for any of the three stages detailed in Section 2.2, with associated hyperparameters  $\boldsymbol{\theta}_1 = \kappa$  (Gamma likelihood) or  $\boldsymbol{\theta}_1 = \xi$  (GP likelihood). As in Section 2.1, let  $\mathbf{x} = (\boldsymbol{\eta}^T, \mu, \boldsymbol{\beta}^T, \mathbf{f}^T)^T$  be the latent Gaussian random field,  $\boldsymbol{\theta}_2$  be the vector of hyperparameters of any of the latent models detailed in Section 2.3, and  $\boldsymbol{\theta} = (\boldsymbol{\theta}_1^T, \boldsymbol{\theta}_2^T)^T$ . Then, from (2.1), the joint posterior distribution of parameters and hyperparameters for any of the three stages, can be written as

$$\begin{aligned} p(\mathbf{x}, \boldsymbol{\theta} \mid \mathbf{y}) &\propto p(\boldsymbol{\theta})p(\mathbf{x} \mid \boldsymbol{\theta}_2) \prod_{i \in \mathcal{I}} p(y(i) \mid x(i), \boldsymbol{\theta}_1) \\ &\propto p(\boldsymbol{\theta}) |\mathbf{Q}_{\boldsymbol{\theta}_2}|^{1/2} \exp \left( -\frac{1}{2} (\mathbf{x} - \boldsymbol{\mu}_{\boldsymbol{\theta}_2})^T \mathbf{Q}_{\boldsymbol{\theta}_2} (\mathbf{x} - \boldsymbol{\mu}_{\boldsymbol{\theta}_2}) + \sum_{i \in \mathcal{I}} \log p(y(i) \mid x(i), \boldsymbol{\theta}_1) \right), \end{aligned} \quad (3.1)$$

where  $\boldsymbol{\mu}_{\boldsymbol{\theta}_2}$  and  $\mathbf{Q}_{\boldsymbol{\theta}_2}$  are the mean and precision matrix of  $\mathbf{x}$ , respectively. Note that for the off-site latent model,  $|\mathcal{I}| = T$  (number of time points) for Stages 1 and 2, and  $|\mathcal{I}|$  is the number of exceedances in Stage 3, whereas for the SPDE latent model,  $|\mathcal{I}| = ST$  (number of space-time points, where  $S$  corresponds to the number of stations) for Stages 1 and 2, whereas  $|\mathcal{I}|$  is the number of space-time exceedances in Stage 3. We emphasize that each stage is fitted separately, and therefore (3.1) applies for each stage independently. The main objectives of the statistical inference are to extract from (3.1) the marginal posterior distributions for each of the elements of the linear predictor vector, and for each element of the hyperparameter vector, i.e.,

$$p(x(i) \mid \mathbf{y}) = \int p(\boldsymbol{\theta} \mid \mathbf{y}) p(x(i) \mid \boldsymbol{\theta}, \mathbf{y}) d\boldsymbol{\theta}, \quad p(\theta_k \mid \mathbf{y}) = \int p(\boldsymbol{\theta} \mid \mathbf{y}) d\theta_{-k}, \quad (3.2)$$

from which predictive distributions may be derived.

In a Bayesian framework, model estimation is typically performed using simulation-based techniques, such as MCMC methods. Alternatively, approximate methods can be used to cope with the computation of high-dimensional integrals needed to obtain posterior distributions. One of such approaches that has become increasingly popular in the last decade is the integrated nested Laplace approximation (INLA; Rue *et al.*, 2009), where posterior distributions of interest are numerically approximated using the Laplace approximation, therefore avoiding the usually complex updating schemes, long running times, and diagnostic convergence checks of simulation-based MCMC. INLA is designed for latent Gaussian models and therefore, it can be successfully used in a wide variety of applications (see for instance Riebler *et al.*, 2012; Lombardo *et al.*, 2018; Krainski *et al.*, 2019). The R-INLA package (Bivand *et al.*, 2015) is a convenient interface to the INLA methodology. A wide variety of models are already implemented, and the package is continuously maintained and updated. In particular, our work motivated the INLA implementation of the Gamma and Weibull quantile regressions, both now available to the users. Details regarding the numerical approximations to (3.2) given by INLA are provided in Section 1 of the Supplementary Material.

## 4 Wind speed probabilistic forecasting results

### 4.1 Automatic off-site predictor selection based on wind direction

An important step to fit our off-site latent model described in Section 2.3.1 is to select a suitable set of off-site predictors  $N_{\mathbf{s}}$ ,  $\mathbf{s} \in \mathcal{S}$ . Here, we develop a data-driven approach for identifying dominant wind directions, which we then use subsequently to automatically choose the off-site predictors. This procedure is similar in spirit to that in Kazor and Hering (2015) who identify wind regimes by fitting a Gaussian mixture model to the wind vector.

For each station  $\mathbf{s} \in \mathcal{S}$ , we fit a mixture of von Mises circular distributions to the time series of wind directions,  $\theta \in [0, 2\pi)$ . The von Mises density with location parameter  $\mu \in \mathbb{R}$  and

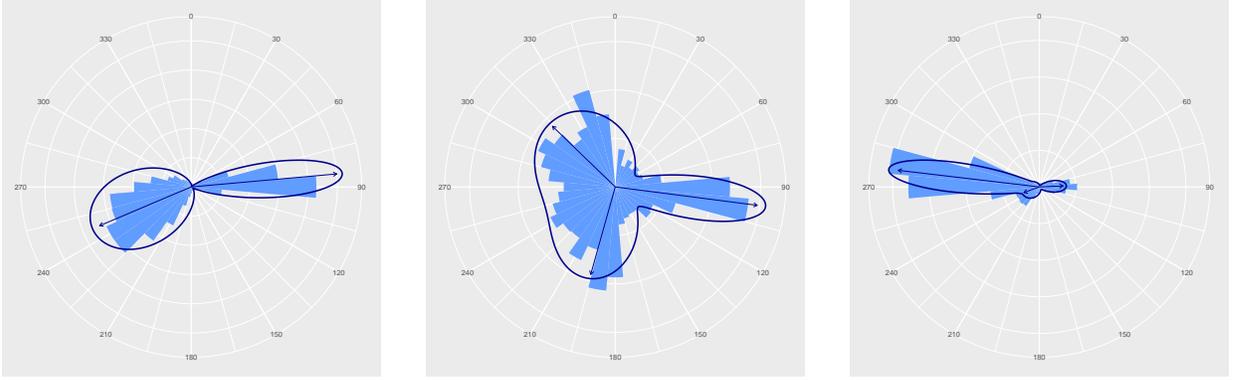


Figure 2: Circular histograms and fitted von Mises distribution mixtures with two and three dominating wind directions for Biddle Butte (BID), Megler (MEG), and Wasco (WAS).

concentration parameter  $v > 0$  is given by  $f(\theta | \mu, v) = [\exp\{v \cos(\theta - \mu)\}] / \{2\pi \mathbf{I}_0(v)\}$ , where  $\mathbf{I}_0(v)$  is the modified Bessel function of order 0. For a mixture of  $M$  von Mises distributions, we identify the dominant wind directions with the locations parameters  $\mu_{s,m}$ ,  $m = 1, \dots, M$ , and construct the sets of angles  $\mathcal{R}_{s,m}(\alpha) = \{\theta \in [0, 2\pi) : \mu_{s,m} - \alpha \leq \theta \leq \mu_{s,m} + \alpha\}$ , for some  $\alpha \in [0, \pi/4]$ , defining station-specific directions of influence. Then, the set of off-site predictors for the station  $\mathbf{s} \in \mathcal{S}$  is chosen as all the stations whose angle with  $\mathbf{s}$  lies in  $\cup_{m=1}^M \mathcal{R}_{s,m}(\alpha)$ , such that their distance to  $\mathbf{s}$  is less than some maximum distance  $d_{\max}$ . We selected the number of components for the mixture of von Mises distributions via the Bayesian Information Criterion (BIC), also guided by the wind roses displayed in Figure 1 of the Supplementary Material. The angle  $\alpha$  depends on the geographical conditions and the distance between stations. We choose  $\alpha = \pi/8$  for all stations. We take  $d_{\max} = 176\text{km}$ , which corresponds to the mean distance among all stations. To illustrate our approach, Figure 2 displays the fitted dominant wind directions for Biddle Butte (BID), Megler (MEG), and Wasco (WAS) stations. The number of off-site predictors for all the stations varies between 2 and 10, while the distance between off-site predictor locations and  $\mathbf{s}$  ranges between 13.3km and 175.3km, for all  $\mathbf{s} \in \mathcal{S}$ . Because the procedure described above is fast, it would be possible to extend it in order to dynamically select potentially different sets of suitable off-site predictors over time.

## 4.2 Posterior predictive distributions

Here we briefly explain how to obtain posterior predictive distributions for the 1-hour, 2-hour, and 3-hour ahead probabilistic forecasts of hourly wind speeds, produced by our spliced Gamma-GP hierarchical Bayesian model, using the linear predictors introduced in Section 2.3. For the SPDE latent model, we use a rolling training period of length 5 days, whereas for the off-site latent model, we multiply this period by the number of stations in each side of the Cascade Mountains, as a way to balance the effective sample sizes of the SPDE and the off-site latent models. We generate 10,000 samples from the posterior predictive distribution, for each station, each forecasting time horizon, and each latent model, as follows: we extract the posterior means of the linear predictor and hyperparameters for each stage, and use (2.4) to obtain 10,000 samples for the Gamma, Bernoulli, and GP predictive distributions. We replace Gamma samples by threshold exceedances (GP samples) whenever the threshold is exceeded, i.e., whenever the associated Bernoulli sample is equal to 1. In other words, the tail of the Gamma distribution is *corrected* by the GP distribution in the presence of exceedances. Note that a more realistic sampling scheme would be to sample non-exceedances from a truncated Gamma distribution, which implies an additional step between Stages 1 and 2, where a truncated Gamma distribution is fitted up to the threshold  $\psi_{s,\alpha}(t)$ . In this case, all values above the threshold would come strictly from the generalized Pareto distribution.

Estimates for the GP shape parameter  $\xi$  are obtained by fitting the GP distribution to each training set. The sequence of estimates computed using all the training data indicate that the wind speed distribution exhibits a moderately heavy upper tail with  $\hat{\xi} \approx 0.07\text{--}0.1$ . Since we restrict ourselves to  $\xi \geq 0$ , our model cannot produce predictive distributions with bounded tails, but this does not seem to be a major concern here; see Figure 7 in Section 5 of the Supplementary Material.

### 4.3 Forecast evaluation for extreme and non-extreme wind speeds

We consider performance measures that describe the ability of our spliced Gamma-GP model to forecast extreme and non-extreme wind speeds. To measure of the overall forecasting ability, we consider the continuous ranked probability score (CRPS) introduced by [Gneiting \*et al.\* \(2007\)](#). The CRPS is a proper scoring rule that quantifies the calibration and sharpness of the forecast ([Gneiting \*et al.\*, 2007](#)). In the following, we describe the CRPS using the notation for the off-site model, but the same definition applies to the SPDE model with the corresponding change of notation. The CRPS is a function of the observed wind speed  $y_s(t)$  and the corresponding predictive distribution  $\hat{F}_s(t)$ ,  $t \in \mathcal{T}$ , and it is defined as

$$\text{CRPS}\{\hat{F}_s(t), y_s(t)\} = \int_{-\infty}^{\infty} [\hat{F}_s(t)(x) - \mathbb{1}\{y_s(t) \leq x\}]^2 dx \quad (4.1)$$

The average CRPS for the  $h$ -hour ahead forecast is

$$\overline{\text{CRPS}}(h) = \frac{1}{T} \sum_{t=1}^T \text{CRPS}\{\hat{F}_s(t+h), y_s(t+h)\}. \quad (4.2)$$

To assess the predictive performance in the upper tail of the distribution, we use the quantile loss (QL) function and the threshold-weighted continuous ranked probability score (twCRPS) ([Gneiting and Ranjan, 2011](#)). The quantile loss function measures the performance of a model to estimate a specific quantile  $\tau \in (0, 1)$ , and it is defined as

$$\text{QL}_\tau(y, q) = \begin{cases} \tau(y - q), & \text{if } y \geq q, \\ -(1 - \tau)(y - q), & \text{if } y < q. \end{cases}$$

If  $Y \sim F$  then  $\arg \min_q E\{\text{QL}_\tau(Y, q)\} = F^{-1}(\tau)$ , so this loss function has been used extensively in non-parametric quantile regression; see [Koenker \(2005\)](#).

The twCRPS is defined as

$$\text{twCRPS}\{\hat{F}_s(t), y_s(t)\} = \int_{-\infty}^{\infty} [\hat{F}_s(t)(x) - \mathbb{1}\{y_s(t) \leq x\}]^2 \omega(x) dx,$$

where  $\omega(x)$  is a non-negative weight function on the real line. For  $\omega(x) \equiv 1$ , the twCRPS reduces to the CRPS in (4.1). We select two different weight functions that highlight our interest in the right tail. We set  $\omega_1(x) = \mathbf{1}\{x \geq r\}$ , with  $r$  equal to the 95% quantile of the wind speed distribution, and  $\omega_2(x) = \Phi(x | r, 1)$ , where  $\Phi(\cdot)$  denotes the standard normal distribution. While  $\omega_2(x)$  yields a proper scoring rule,  $\omega_1(x)$  does not because the distribution is left-truncated; see [Lerch \*et al.\* \(2017\)](#). The average quantile loss and twCRPS for the  $h$ -hour ahead forecast is defined as in (4.2). Lower values of these criteria are better.

Since our main goal is to accurately forecast extreme wind speeds, we compare our proposed spliced Gamma-GP models against a baseline Gamma model that forecasts wind speeds using only the first stage described in Section 2.2 (hence, without correcting the upper tail). This baseline model assumes the two latent structures detailed in Section 2.3. Given that our proposed model is fairly complex, the purpose of this comparison is to check if the GP correction of the tail improves the forecast of strong wind speeds. Table 1 shows performance measures for one-hour ahead forecasts for the baseline and the Gamma-GP models at ten selected stations, as well as average prediction skills at all stations. Throughout all the fits, we set the probabilities defined in Section 2.2 as  $\alpha = 0.8$  and  $\beta = 0.5$ . The value of  $\alpha$  was chosen as a compromise between having a good approximation to the tail of the wind speed distribution and having enough exceedances to fit the GP approximation in each time window. The value of  $\beta$  was chosen pragmatically in order to reduce the correlation between the estimated GP parameters (see Section 6 of the Supplementary Material for further details regarding the GP parametrization). The performance measures are the average CRPS, average twCRPS using  $\omega_1(x)$  and  $\omega_2(x)$  as defined before, and average quantile loss (QL). From Table 1, we can see that the SPDE latent model performs better than the off-site latent model at predicting strong values of wind speeds. The difference might be due to the difficulty of the off-site latent model at estimating the GP shape parameter at each station, while a single shape parameter is

assumed in the SPDE latent model, reducing dramatically the estimated posterior predictive uncertainty by borrowing strength across all stations. Both the off-site and the SPDE latent models appear to be better than their baseline counterparts when focusing on the upper tail of the distribution, showing that the GP correction may be useful to improve the forecasting of strong wind speeds, although further diagnostics would be needed to draw firm conclusions.

Table 1: Performance measures for one-hour ahead forecast using (left to right) the off-site model, the off-site baseline model, the SPDE model, and the SPDE baseline model at ten selected stations. We report the average continuous ranked probability score (CRPS); average threshold-weighted continuous ranked probability score (twCRPS) using the indicator weight function  $\omega_1(x)$  and the normal weight function  $\omega_2(x)$ ; and average quantile loss (QL) with  $\tau = 0.99$ . The bottom line is the average for each performance measure over the 20 stations. The best model is highlighted in bold.

Station	twCRPS			
	CRPS	$\omega_1(x) = \mathbb{1}\{x \geq r\}$	$\omega_2(x) = \Phi(x   r, 1)$	QL ( $\tau = 0.99$ )
AUG	0.92/0.93/0.92/0.92	0.06/0.06/0.06/0.05	0.06/0.06/0.06/0.09	0.80/0.81/0.76/0.80
BID	0.90/1.09/0.91/0.92	0.09/0.09/0.08/0.10	0.09/0.09/0.06/0.11	0.77/1.08/0.71/1.06
FOR	0.39/0.39/0.38/0.39	0.04/0.04/0.04/0.06	0.03/0.02/0.02/0.07	0.38/0.40/0.39/0.39
HOO	0.61/0.62/0.60/0.61	0.04/0.04/0.04/0.07	0.05/0.05/0.04/0.09	0.56/0.58/0.53/0.61
KEN	1.14/1.15/1.01/1.13	0.09/0.09/0.06/0.10	0.05/0.05/0.05/0.08	0.95/0.98/0.90/0.10
MEG	0.67/0.68/0.67/0.67	0.07/0.07/0.04/0.08	0.07/0.07/0.06/0.10	0.62/0.65/0.61/0.82
NAS	0.69/0.69/0.50/0.75	0.06/0.06/0.05/0.09	0.06/0.06/0.07/0.09	0.65/0.67/0.64/0.85
SML	0.94/0.95/0.50/0.60	0.06/0.06/0.04/0.10	0.07/0.07/0.07/0.07	0.91/0.93/0.81/0.13
SUN	1.16/1.17/0.71/1.92	0.09/0.09/0.09/0.13	0.08/0.08/0.07/0.12	1.11/1.14/1.01/1.19
WAS	0.82/0.82/0.81/0.81	0.06/0.05/0.05/0.07	0.06/0.06/0.05/0.06	0.77/0.79/0.70/0.81
Avg.	0.82/0.85/ <b>0.75</b> /0.84	0.07/0.07/ <b>0.06</b> /0.08	0.07/0.07/ <b>0.06</b> /0.08	0.76/0.79/ <b>0.73</b> /0.78

We assess the calibration of our probabilistic forecasts using reliability diagrams (see, e.g., [Lenzi et al., 2017](#)). Reliability refers to the ability of the model to match the observation frequencies. The diagram is constructed as follows: for every station, we compute the nominal coverage rate, which is the proportion of times that the cumulative distribution of our spliced Gamma-GP model is below a certain threshold. Our model is well calibrated if this proportion is close to the observed frequencies. Figure 3 shows reliability diagrams for the off-site latent model (coral line) and the SPDE latent model (cyan line), as well as for their baseline counterparts (in green and purple, respectively). We can see that both models tend to overestimate the wind speed quantiles smaller than the median. The off-site latent model

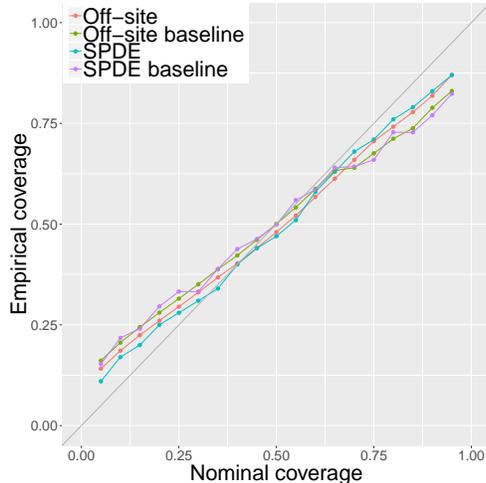


Figure 3: Reliability diagram for the off-site model (coral line), off-site baseline model (green line), SPDE model (cyan line), and SPDE baseline model (purple line).

underestimates the wind speed values larger than the median, whereas the SPDE latent model is better calibrated at higher quantiles.

Finally, to compare the ability of our proposed models to forecast extreme wind speeds, we compute the pseudo-uniform scores  $\hat{u}_{\mathbf{s},k}(t+h) = \hat{F}_{\mathbf{s},k}(t+h)(y_{\mathbf{s}}(t+h))$  for all forecasted values  $y_{\mathbf{s}}(t+h)$ ,  $h = 1, 2, 3$ , where  $\hat{F}_{\mathbf{s},k}(t+h)$  is the predictive distribution for location  $\mathbf{s} \in \mathcal{S}$  and time  $(t+h)$  based on the  $k$ -th training set. Note that we here use the notation for the off-site model, but the same definition applies to the SPDE model with the corresponding change of notation. We then plot the histogram of  $\{\hat{u}_{\mathbf{s},k}(t+h)\}_{k \geq 1}$  conditional on being greater than 0.6. We refer to such diagnostics as conditional Probability Integral Transform (PIT) plots. The results for 1h ahead forecasts are shown in Figure 4. Because we condition on  $\hat{u}_{\mathbf{s},k}(t+h) > 0.6$ , the conditional PIT plots are informative about the model ability to forecast moderately strong wind speeds using the Gamma model described in Stage 1, and strong wind speeds using the GP model in Stage 3; recall Section 2.2. We can see that both the off-site model and the SPDE model tend to outperform their baseline counterparts, while the SPDE model seems to perform slightly better than the off-site model at forecasting strong wind speeds.

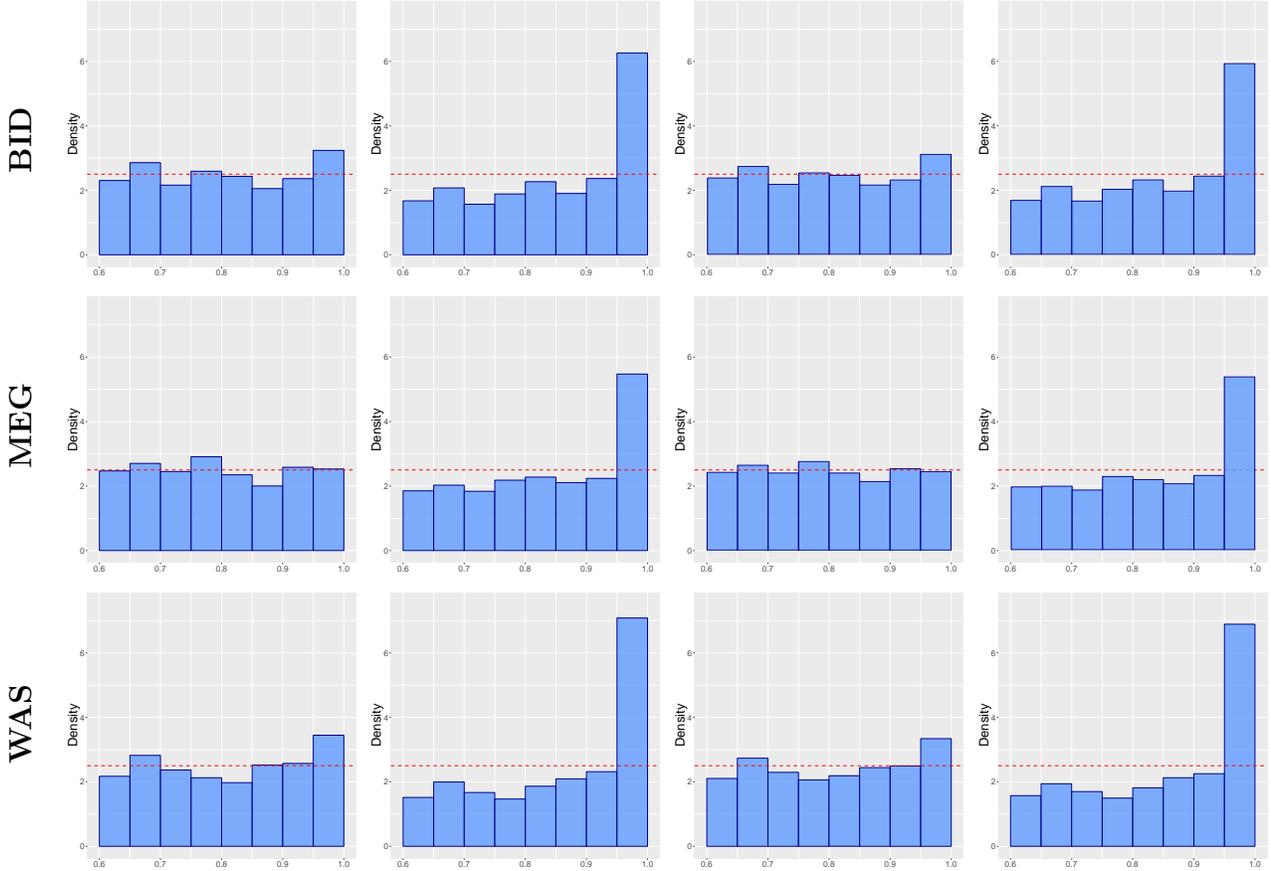


Figure 4: Conditional PIT plots for 1h ahead forecasts for the off-site model (first column), the off-site baseline model (second column), the SPDE model (third column), and the SPDE baseline model (fourth column) at three selected stations (different rows). Uniform histograms imply good forecasting ability.

## 5 Conclusion

In this work, we have explored a hierarchical Bayesian spliced Gamma-GP model designed to forecast extreme wind speeds. Our model corrects the tail of the Gamma distribution by a generalized Pareto distribution in the presence of exceedances. Each stage of our model belongs to the class of latent Gaussian models, for which the integrated nested Laplace approximation (INLA; [Rue et al., 2009](#)) method is well-suited. Considering an additive latent structure, we proposed two types of linear predictors describing the spatio-temporal dynamics of wind speeds. The first linear predictor includes off-site information in terms of lagged wind speeds from neighboring stations. To select the neighbors, we propose an automatic method to identify

dominant wind direction patterns based on mixtures of von Mises circular distributions. The second linear predictor considers a spatio-temporal term with Matérn covariance structure (driven by a stochastic partial differential equation), that varies in time according to a first-order autoregressive dynamics. In terms of forecasting extreme wind speeds, both models seem to perform decently well, although the SPDE latent model is better calibrated at high quantiles, because it better exploits spatial information. It would be interesting to explore the potential of the SPDE latent model to predict extreme wind speeds at unobserved locations, which could be helpful for optimal design of wind farms.

Thanks to the very powerful and fast INLA estimation approach, we can implement in a reasonable amount of time complex hierarchical spatial models that are well suited to wind speed data. Specifically, each set of 1h–3h ahead forecasts using a single core took less than 2 minutes for the off-site model, and less than 20 minutes for the SPDE model. The parallelization of these computations was done using resources for distributed computing.

By selecting a suitable distribution in the first stage, our spliced Gamma-GP model for exceedances can be easily adapted to model and forecast other types of environmental data. We fitted the bulk and the tail of the wind speed distribution separately, as extreme events usually behave differently from low and moderately large events, and therefore only extreme observations may give information about the tail of the distribution (Rootzén *et al.*, 2018). If the model for the bulk is misspecified, then the  $\alpha$ -quantile might not be well estimated. But unless the fitted proportion of exceedances deviates considerable from the truth (i.e., the parameter  $\alpha$ ), this should not considerably affect the fit for the tail.

Improving the quality of wind power forecasts is a constant challenge, but there are many possible directions on how to incorporate a rigorous extreme value analysis to the estimation of wind speed. For instance, a better representation of the physical phenomena involved in the generation of strong wind speeds may be achieved by introducing outputs from numerical

climate models. Alternatively, non-stationary spatial models with local anisotropies informed by wind direction could be developed. Moreover, for a broader description of the tail of the wind speed distribution, it would be interesting to fit the GP distribution with  $\xi \in \mathbb{R}$  (i.e., not only  $\xi \geq 0$ ), which could be useful for lighter-tailed data. Nevertheless, designing suitable shrinkage priors for  $\xi < 0$  is still an open question. For computational convenience (and because of the constraints with INLA), we have assumed conditional independence of the data given the latent process. In cases where strong (tail) dependence prevails, more specialized extreme-value models should be considered, such as generalized Pareto processes (Thibaud and Opitz, 2015), max-stable models (Huser and Davison, 2014), and flexible copula models (Castro-Camilo and Huser, 2019). Although these models are attractive from a theoretical viewpoint, they are very cumbersome to fit, especially in high dimensions.

## References

- Alexiadis, M., Dokopoulos, P., Sahsamanoglou, H. *et al.* (1999) Wind speed and power forecasting based on spatial correlation models. *IEEE Transactions on Energy Conversion* **14**(3), 836–842.
- Bakka, H., Rue, H., Fuglstad, G. A., Riebler, A., Bolin, D., Illian, J., Krainski, E., Simpson, D. and Lindgren, F. (2018) Spatial modelling with R-INLA: A review. *WIREs Computational Statistics*, 10: null. doi: 10.1002/wics.1443.
- Bivand, R., Gómez-Rubio, V. and Rue, H. (2015) Spatial data analysis with R-INLA with some extensions. *Journal of Statistical Software* **63**(1), 1–31.
- Casson, E. and Coles, S. (1999) Spatial regression models for extremes. *Extremes* **1**(4), 449–468.
- Castro-Camilo, D. and Huser, R. (2019) Local likelihood estimation of complex tail dependence structures, applied to us precipitation extremes. *arXiv preprint arXiv:1710.00875* Submitted.
- Cooley, D., Nychka, D. and Naveau, P. (2007) Bayesian spatial modeling of extreme precipitation return levels. *Journal of the American Statistical Association* **102**(479), 824–840.
- Davison, A. C. and Huser, R. (2015) Statistics of extremes. *Annual Review of Statistics and its Application* **2**, 203–235.
- Davison, A. C. and Smith, R. L. (1990) Models for exceedances over high thresholds. *Journal of the Royal Statistical Society. Series B (Methodological)* **52**(3), 393–442.
- Erdem, E. and Shi, J. (2011) ARMA based approaches for forecasting the tuple of wind speed and direction. *Applied Energy* **88**(4), 1405–1414.

- Fuglstad, G.-A., Simpson, D., Lindgren, F. and Rue, H. (2018) Constructing priors that penalize the complexity of Gaussian random fields. *Journal of the American Statistical Association* pp. 1–8.
- Gneiting, T., Balabdaoui, F. and Raftery, A. E. (2007) Probabilistic forecasts, calibration and sharpness. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **69**(2), 243–268.
- Gneiting, T., Larson, K., Westrick, K., Genton, M. G. and Aldrich, E. (2006) Calibrated probabilistic forecasting at the stateline wind energy center: The regime-switching space-time method. *Journal of the American Statistical Association* **101**(475), 968–979.
- Gneiting, T. and Ranjan, R. (2011) Comparing density forecasts using threshold-and quantile-weighted scoring rules. *Journal of Business & Economic Statistics* **29**(3), 411–422.
- Hering, A. S. and Genton, M. G. (2010) Powering up with space-time wind forecasting. *Journal of the American Statistical Association* **105**(489), 92–104.
- Hering, A. S., Kazor, K. and Kleiber, W. (2015) A markov-switching vector autoregressive stochastic wind generator for multiple spatial and temporal scales. *Resources* **4**(1), 70–92.
- Huang, Z. and Chalabi, Z. (1995) Use of time-series analysis to model and forecast wind speed. *Journal of Wind Engineering and Industrial Aerodynamics* **56**(2–3), 311–322.
- Huser, R. and Davison, A. (2014) Space-time modelling of extreme events. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **76**(2), 439–461.
- Kazor, K. and Hering, A. S. (2015) The role of regimes in short-term wind speed forecasting at multiple wind farms. *Stat* **4**(1), 271–290.
- Koenker, R. (2005) *Quantile Regression*. Cambridge University Press, Cambridge UK.
- Krainski, E. T., Gómez-Rubio, V., Bakka, H., Lenzi, A., Castro-Camilo, D., Simpson, D., Lindgren, F. and Rue, H. (2019) *Advanced Spatial Modeling with Stochastic Partial Differential Equations using R and INLA*. CRC press. Github version [www.r-inla.org/spde-book](http://www.r-inla.org/spde-book).
- Lenzi, A., Pinson, P., Clemmensen, L. H. and Guillot, G. (2017) Spatial models for probabilistic prediction of wind power with application to annual-average and high temporal resolution data. *Stochastic Environmental Research and Risk Assessment* **31**(7), 1615–1631.
- Lerch, S., Thorarinsdottir, T. L., Ravazzolo, F., Gneiting, T. *et al.* (2017) Forecaster’s dilemma: Extreme events and forecast evaluation. *Statistical Science* **32**(1), 106–127.
- Li, G. and Shi, J. (2010) On comparing three artificial neural networks for wind speed forecasting. *Applied Energy* **87**(7), 2313–2320.
- 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. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **73**(4), 423–498.
- Lombardo, L., Opitz, T. and Huser, R. (2018) Point process-based modeling of multiple debris flow landslides using INLA: an application to the 2009 Messina disaster. *Stochastic Environmental Research and Risk Assessment* **32**(7), 2179–2198.

- Naveau, P., Huser, R., Ribereau, P. and Hannart, A. (2016) Modeling jointly low, moderate, and heavy rainfall intensities without a threshold selection. *Water Resources Research* **52**(4), 2753–2769.
- Opitz, T., Huser, R., Bakka, H. and Rue, H. (2018) INLA goes extreme: Bayesian tail regression for the estimation of high spatio-temporal quantiles. *Extremes* **21**(3), 441–462.
- Palomares-Salas, J., De La Rosa, J., Ramiro, J., Melgar, J., Aguera, A. and Moreno, A. (2009) Arima vs. neural networks for wind speed forecasting. In *Computational Intelligence for Measurement Systems and Applications, 2009. CIMSA'09. IEEE International Conference on*, pp. 129–133.
- Pinson, P. and Madsen, H. (2012) Adaptive modelling and forecasting of offshore wind power fluctuations with Markov-switching autoregressive models. *Journal of forecasting* **31**(4), 281–313.
- Riebler, A., Held, L., Rue, H. *et al.* (2012) Estimation and extrapolation of time trends in registry data—borrowing strength from related populations. *The Annals of Applied Statistics* **6**(1), 304–333.
- Rootzén, H., Segers, J. and Wadsworth, J. L. (2018) Multivariate peaks over thresholds models. *Extremes* **21**(1), 115–145.
- Rue, H. and Held, L. (2005) *Gaussian Markov random fields: theory and applications*. CRC press.
- Rue, H., Martino, S. and Chopin, N. (2009) Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *Journal of the Royal statistical society: Series B (Statistical Methodology)* **71**(2), 319–392.
- Rue, H., Riebler, A., Sørbye, S. H., Illian, J. B., Simpson, D. P. and Lindgren, F. K. (2017) Bayesian computing with INLA: a review. *Annual Review of Statistics and Its Application* **4**, 395–421.
- Scarrott, C. and MacDonald, A. (2012) A review of extreme value threshold estimation and uncertainty quantification. *REVSTAT—Statistical Journal* **10**(1), 33–60.
- Shih, D. C.-F. (2008) Wind characterization and potential assessment using spectral analysis. *Stochastic Environmental Research and Risk Assessment* **22**(2), 247–256.
- Simpson, D., Rue, H., Riebler, A., Martins, T. G., Sørbye, S. H. *et al.* (2017) Penalising model component complexity: A principled, practical approach to constructing priors. *Statistical Science* **32**(1), 1–28.
- Tancredi, A., Anderson, C. and O’Hagan, A. (2006) Accounting for threshold uncertainty in extreme value estimation. *Extremes* **9**(2), 87.
- Thibaud, E. and Opitz, T. (2015) Efficient inference and simulation for elliptical pareto processes. *Biometrika* **102**(4), 855–870.
- Zhu, X. and Genton, M. G. (2012) Short-term wind speed forecasting for power system operations. *International Statistical Review* **80**(1), 2–23.