SPATIOTEMPORAL WILDFIRE MODELING THROUGH POINT PROCESSES WITH MODERATE AND EXTREME MARKS

By Jonathan Koh[1,2,a], François Pimont[3,b], Jean-Luc Dupuy[3,c] and Thomas Opitz[4,d]

[1]*Institute of Mathematics, EPFL*

[2]*Institute of Mathematical Statistics and Actuarial Science, Oeschger Centre for Climate Change Research, University of Bern,* [a]*jonathan.koh@unibe.ch*

[3]*URFM UR629, INRAE,* [b]*francois.pimont@inrae.fr,* [c]*jean-luc.dupuy@inrae.fr*

[4]*BioSP UR546, INRAE,* [d]*thomas.opitz@inrae.fr*

Accurate spatiotemporal modeling of conditions leading to moderate and large wildfires provides better understanding of mechanisms driving fire-prone ecosystems and improves risk management. Here, we develop a joint model for the occurrence intensity and the wildfire size distribution, by combining extreme-value theory and point processes within a novel Bayesian hierarchical model, and use it to study daily summer wildfire data for the French Mediterranean basin during 1995–2018. The occurrence component models wildfire ignitions as a spatiotemporal log-Gaussian Cox process. Burnt areas are numerical marks attached to points and are considered as extreme if they exceed a high threshold. The size component is a two-component mixture varying in space and time that jointly models moderate and extreme fires. We capture nonlinear influence of covariates (Fire Weather Index, forest cover) through component-specific smooth functions which may vary with season. We propose estimating shared random effects between model components to reveal and interpret common drivers of different aspects of wildfire activity. This increases parsimony and reduces estimation uncertainty, giving better predictions. Specific stratified subsampling of zero counts is implemented to cope with large observation vectors. We compare and validate models through predictive scores and visual diagnostics. Our methodology provides a holistic approach to explaining and predicting the drivers of wildfire activity and associated uncertainties.

1. Introduction. Wildfires are defined as uncontrolled fires of combustible natural vegetation, such as trees in a forest. Their activity usually shows seasonal cycles, as several conditions must coincide for their occurrence: the presence of combustible material as fuel, its easy flammability resulting from weather conditions such as droughts, and a trigger. Triggers include natural causes, such as lightning, but the majority of occurrences in Europe are caused by human activity, either intentional (arson), neglectful (cigarette stubs), or accidental (agriculture).

Wildfires represent major environmental and ecological risks worldwide. They provoke many human casualties and substantial economic costs and can trigger extreme air pollution episodes and important losses of biomass and biodiversity. While climate change is expected to exacerbate their frequency and extent (Jones et al. (2020)), wildfires themselves contribute an important fraction of global greenhouse gases that can accelerate climate change. To aid in wildfire prevention and risk mitigation, one must identify the factors contributing to wildfires and predict their spatiotemporal distribution. Prediction maps of various components of wildfire risk are relevant for the study of historical periods, for short-term forecasting, and for long-term projections.

The study of wildfire activity has led to a large body of statistical and machine learning literature on methods for identifying risk factors and producing risk maps (Pereira and Turkman (2019), Preisler et al. (2004), Xi et al. (2019)). Most studies focus on modeling either occurrence counts or sizes, the latter usually represented by the burnt areas of spatially and temporally contiguous wildfire events. In occurrence modeling, the spatial or spatiotemporal pattern of ignition points (or other representative points of separate wildfire events) can be analyzed with point process tools (Genton et al. (2006), Opitz, Bonneu and Gabriel (2020), Peng, Schoenberg and Woods (2005), Pereira and Turkman (2019), Serra et al. (2013), Tonini et al. (2017), Woolford et al. (2021), Xu and Schoenberg (2011)). Often, data are available, as presence/absence or counts, over dense spatial or spatiotemporal grids or have been transformed to such representations to facilitate modeling and to harmonize different spatial-temporal scales of wildfire and predictor data, such as weather conditions, land cover, and land use.

Burnt area, a key measure of wildfire impact, usually provides a good proxy for biomass loss and greenhouse gas emissions, and it allows impacts on ecosystem services, such as biodiversity or clean air, to be assessed. Many univariate probability distributions have been explored for modeling fire sizes (e.g., Cui and Perera (2008), Cumming (2001), Pereira and Turkman (2019), Schoenberg, Peng and Woods (2003)). Empirical distributions are usually heavy-tailed, as with the wildfire data we consider in Mediterranean France, and a few extreme wildfires account for a very large fraction of total burnt area. There is no consensus on which distribution provides the best fit (Pereira and Turkman (2019)). Early use of extreme-value methodology to study wildfires goes back to Moritz (1997), who analyzed extreme wildfires through descriptive approaches. Generative approaches, based on spatial models, suggested by extreme-value theory, using distributions, such as the generalized Pareto distribution (GPD), have also been studied (e.g., approaches by De Zea Bermudez et al. (2009), Mendes et al. (2010), Pereira and Turkman (2019), Turkman, Turkman and Pereira (2010)).

Joint statistical analyses of wildfire occurrence and sizes have been proposed and often use tools for marked point processes, where numerical marks represent burnt areas. Descriptive approaches (e.g., Tonini et al. (2017)) characterize different regimes of wildfire activity (i.e., numbers, sizes, spatialtemporal autocorrelation) by taking into account weather, land cover, fire management, and environmental factors. For explanatory and predictive modeling, Bayesian hierarchical models are useful; they can include latent Gaussian components to allow for observation and estimation uncertainty and capture nonlinear influences of covariates. One may consider only categorical information (e.g., small and large wildfires) without attempting to model the continuous distribution of sizes; for example, Serra et al. (2014) construct a Bayesian spatiotemporal "hurdle" model for large wildfires. As to continuous distributions, Ríos-Pena et al. (2018) implement MCMC inference for zero-inflated Beta-regression to model the occurrence of wildfires in spatial units, with absence corresponding to zero-inflation, while the positive area fraction covered by wildfires is captured through the Beta distribution. Joseph et al. (2019) estimate separate regression models with random effects for occurrence numbers in areal units and for sizes and study posterior predictive distributions for block maxima of wildfire sizes. Pimont et al. (2021) developed a marked spatiotemporal log-Gaussian Cox process model, called *Firelihood*, for daily data by applying the integrated nested Laplace approximation (INLA, Rue, Martino and Chopin (2009)) for Bayesian inference on most components of the model. Their distribution of wildfire sizes over positive values is based on estimating exceedance probabilities and excess distributions over a range of severity thresholds. Weather information is included through a nonlinear effect of the Fire Weather Index (FWI, van Wagner (1977)), constructed to yield high correlation with wildfire activity.

In this work we develop the following novelties to address key shortcomings of the works cited above. As large wildfires play a dominant and critical role for fire activity, due to the heavy tails of burnt areas, we focus on accurate modeling of their distribution and, in particular, its spatiotemporal variation. However, models constructed using only extreme wildfires would lead to high estimation uncertainty when inferring complex spatiotemporal structures. We, therefore, propose the joint estimation of extreme and nonextreme wildfires where the model borrows strength from the latter to help estimate the former; the large number of observations available for moderate fires improves the prediction of larger fires, so changes in extreme fire activity are better accounted for.

Complex models such as *Firelihood* require separate estimation of the occurrence and size model components, thus hampering inferences that exploit interactions between them. Temporal stochastic structures are often restricted to the spatiotemporal variability in covariates. In Pimont et al. (2021), simulated predictive distributions of wildfire activity for various divisions of the space-time domain failed to capture some very extreme events, specifically the year 2003. Here, we increase the flexibility of the spatiotemporal structure, especially for extremes.

Our new approach leverages a combination of marked point processes, defined over continuous space and time, and extreme-value theory to represent the mechanisms leading to wildfires exceeding a high severity threshold for burnt areas. The point pattern of extreme fires is viewed as a thinning of the full pattern, and we select a suitable threshold before using the GPD model for threshold excesses.

We also advocate sharing spatial random effects that affect several model components simultaneously: these effects are estimated for one response variable (e.g., wildfire counts), but we also include them with scaling coefficients in other response variables (e.g., wildfire size exceedances). This approach decreases uncertainty in the estimation of those regression equations whose vector of observed responses carries too little information to estimate complex predictive structures. We will highlight the improved inferences through sharing in our wildfire application. Besides increasing model parsimony, sharing also provides new scientific insight by highlighting joint drivers of different wildfire components.

The FWI quantifies the influence of weather drivers on wildfire activity and is often mapped as an index for fire danger, for instance, by the French weather service Météo France. Model diagnostics of Pimont et al. (2021) showed that the predictive power of FWI in France varies across seasons, and we estimate a more sophisticated seasonal nonlinear FWI effect.

Predictive model validation is intricate because of heavy tails and high prediction uncertainty for individual wildfires. Customary validation scores, such as means of squared or absolute errors, are not useful. In addition to visual diagnostics, we tackle this difficulty through joint assessment of several numerical criteria, either through scores for binary data (e.g., Area under the Curve, Fawcett (2006)) to assess the exceedance behavior over a relevant severity threshold or through comparison of probabilistic scores for continuous predictions, such as the scaled Continuous Ranked Probability Score (Bolin and Wallin (2020)).

We estimate our marked log-Gaussian Cox process in a Bayesian setting using INLA (Illian, Sørbye and Rue (2012)) by adopting penalized complexity (PC) priors for hyperparameters (Simpson et al. (2017b)). Gaussian process priors follow the Matérn covariance function, and we use the Stochastic Partial Differential Equation (SPDE) approach of Lindgren, Rue and Lindström (2011) for numerically efficient Gauss–Markov approximation. Fully Bayesian inference is out of reach, as we have several million observations of wildfire counts, so we devise a specific subsampling scheme for zero counts that keeps a relatively

larger proportion of observations with high FWI for which most wildfires occur. This allows joint Bayesian inference on all components, and we ensure that our subsample sizes allow the fitting of models on standard personal computers, in contrast to other recent approaches (e.g., Joseph et al. (2019), Opitz, Bonneu and Gabriel (2020), Pimont et al. (2021)) requiring high computer memory.

In the remainder of the paper, we first explore available data on wildfires and predictors in Section 2. We provide general background on extreme-value theory and point processes and on how to combine them in a Bayesian hierarchical model using the INLA-SPDE method, in Section 3. The specific hierarchical structure for the joint analysis of extreme and nonextreme wildfires is developed in Section 4. Estimation with subsampling of pixel-days without wildfire occurrences is detailed in Section 3.3. After a comparative analysis of models in Section 5.1, we highlight key findings and prediction of wildfire activity components in Sections 5.3 and 6, and we conclude in Section 7.

2. Wildfire data. Since 1973, wildfires occurring in the fire-prone French Mediterranean region have been recorded in the Prométhée database (www.promethee.com). Each wildfire occurrence is reported with its fire ignition cell in a 2×2 km^2 grid, day of detection, and burnt area in hectare (ha). Inconsistent reporting was found for small wildfires, especially smaller than 1 ha, and we keep only data with reported burnt area larger than 1 ha, that is, of escaped wildfires that could not quickly be extinguished. We use the observation period 1995–2018, for which gridded weather reanalysis data (*SAFRAN* model of *Météo France*) and information on forested area are available.

Figure 1 illustrates the heavy tails in the distribution of burnt areas and strong spatial variability in numbers and sizes of wildfires. It also shows the contours of administrative areas ("*départements*") in the study region. Small to moderately large wildfires dominate the pie charts for wildfire counts, while large wildfires dominate those for aggregated burnt area. Certain spatial patterns are similar in the distribution of numbers and sizes of wildfires (top and bottom display of Figure 1, respectively), but there are notable differences. For example, large wildfire numbers do not always entail large aggregated burnt areas, as we see for the Pyrénées-Orientales département in the southwest. The disparities show the need to model spatiotemporal structures in both wildfire numbers and sizes as well as in their interaction. Figure 2 (left panel) shows a histogram of burnt area values. The sum of burnt areas exceeding the empirical 99%-quantile is larger than the corresponding sum for the remaining wildfires.

The SAFRAN model provides gridded weather reanalyses at 8 km resolution. The joint influence of weather variables, such as temperature, precipitation, and wind speed on fire activity patterns is highly complex. Meteorological indices of fire danger have been constructed, such as the widely used unitless Fire Weather Index (FWI) that was originally defined for Canadian forests. Its values are often used for direct interpretation and fire danger mapping. Instead, we here study its relationship to components of fire risk, such as occurrence frequency and wildfire sizes. For our models we preprocess SAFRAN data to compute daily FWI, using the `cffdrs` library (Wang et al. (2017)), and utilize the SAFRAN grid by aggregating daily wildfire counts to its cells; Pimont et al. (2021) describe the full procedure of computing FWI as well as justifications for using this spatial-temporal resolution. Forest cover is another crucial explanatory variable. Around 60% of the study area has forested areas or vegetation types that ignite easily (shrubland or other natural herbaceous vegetation). Wildfires do not propagate easily through the other available land cover types. We consider relevant fuel material through the proportion covered by this vegetation in each SAFRAN grid cell (and day) based on CORINE Land Cover data (CLC). CLC dynamics are captured by linear temporal interpolation of several inventories. We refer to the resulting pixel-day predictor as forested area (FA) in %.
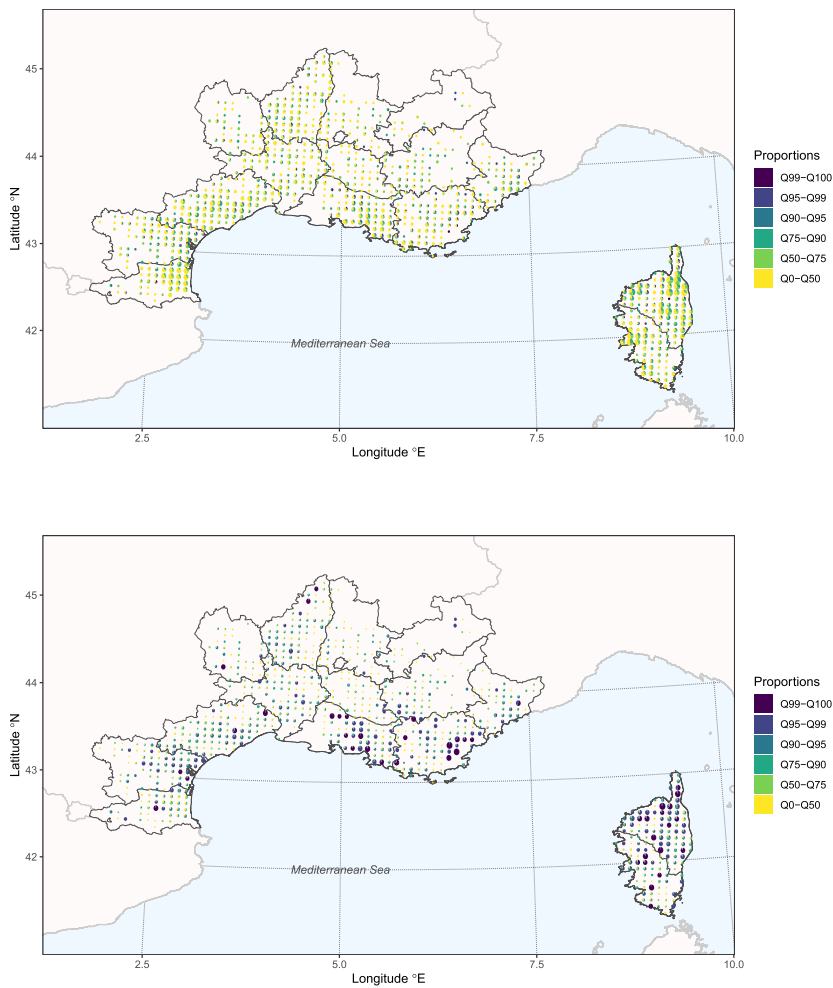
FIG. 1. *Maps of Prométhée data aggregated to the SAFRAN grid at 8 km resolution. The pie charts in the grid cells are based on 1 wildfire size classes with boundaries given by empirical quantile levels 0, 0.5, 0.75, 0.9, 0.95, 0.99, 1 of all burnt areas (June–October). Top display: Pie charts show relative count proportions over the six classes and have size increasing with increasing counts. Bottom display: Pie charts show relative burnt area proportions and have size increasing with increasing aggregated burnt area.*
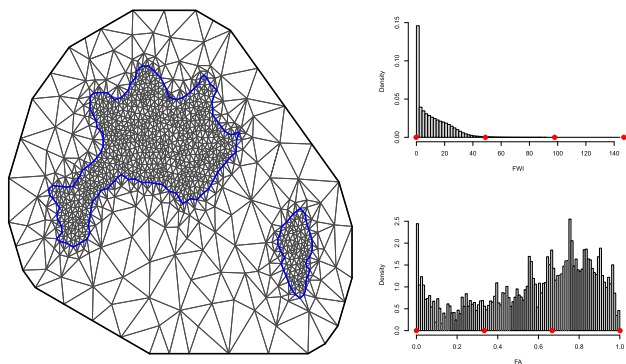

FIG. 2. *Burnt area distribution. Left: Histogram of burnt areas (ha) in base-10-logarithm. Middle: Parameter stability of the tail index. Right: P-values for the null hypothesis of a GPD distribution above the threshold; tick labels on top indicate the number of fires above the thresholds.*

3. Methods for point patterns with extreme marks.

3.1. *Extreme-value theory.* Given a random variable X whose distribution F satisfies mild regularity conditions, the generalized Pareto distribution (GPD) arises asymptotically for the positive excesses of X above a threshold increasing to $x^\star = \sup\{x : F(x) < 1\}$ (Coles (2001)). Therefore, given a large threshold $u < x^\star$, the tail behavior of a wide class of random variables X can be approximated as

$$(1) \qquad \Pr(X > x + u \mid X > u) \approx \mathrm{GPD}_{\sigma,\xi}(x) = \begin{cases} (1 + \xi x/\sigma)_+^{-1/\xi} & \xi \neq 0, \\ \exp(-x/\sigma) & \xi = 0, \end{cases} \quad x > 0$$

with shape parameter $\xi \in \mathbb{R}$ and scale parameter $\sigma = \sigma(u) > 0$, where $a_+ = \max(a, 0)$. The shape parameter determines the rate of tail decay with slow power-law decay for $\xi > 0$, exponential decay for $\xi = 0$, and polynomial decay toward a finite upper bound for $\xi < 0$. Writing $p_{\mathrm{exc}} = 1 - F(u)$ for the exceedance probability of X above u, we use (1) to approximate the cumulative distribution function F of X above the threshold u (Davison and Smith (1990)) as

$$(2) \qquad F(x) \approx 1 - p_{\mathrm{exc}} \, \mathrm{GPD}_{\sigma,\xi}(x - u), \quad x > u,$$

where ξ, σ and p_{exc} are parameters to be estimated. We account for dependence and nonstationarity among observations by including auxiliary variables and Gaussian random effects in σ and p_{exc}. Nonstationarity in ξ is often hard to identify, and we, therefore, keep ξ stationary.

Based on (2), we model the conditional GPD of fire size excesses and p_{exc}. To explore the tail behavior of all fire sizes pooled together and to choose an appropriate threshold u, we can use mean excess plots (see Section D of the Supplementary Material (Koh et al. (2023))) or the following threshold stability plot of parameters, here considered for the GPD shape ξ, estimated by maximum likelihood for thresholds $v_m > \cdots > v_1$. We use multiple statistical tests (Northrop and Coleman (2014)) to test the null hypotheses that the data come from a common truncated GPD on all intervals (v_k, v_{k+1}), $k = 1, \ldots, m$, where $v_{m+1} = \infty$. Using $m = 40$ equidistant intervals of length 5 ha for fire sizes, Figure 2 provides evidence that stability is reached above approximately the 95% quantile (79 ha), as we fail to reject the null hypothesis $\xi_k = \cdots = \xi_m$ for intervals with $v_k > 79$ ha and estimated shape $\hat{\xi}_k \approx 0.7$.

Joseph et al. (2019) modeled fire sizes in the contiguous United States and concluded that the GPD leads to overestimation of extreme fire sizes. However, they fitted the GPD to the full distribution; Figure 2 shows that we would have obtained $\hat{\xi} \approx 1.4$ for $u = 1$, corresponding to extremely slow tail decay.

3.2. *Mark-dependent thinning of point processes.* We consider the point pattern of fire ignitions and burnt areas as a realization of a spatiotemporal marked point process, that is, of a random count measure N that attributes values $N(B) \in \{0, 1, 2, \ldots\}$ to Borel sets $B \subset \mathcal{D} \subset \mathbb{R}^2 \times \mathbb{R}$. By abuse of notation, we also refer to the number of points in $\mathcal{D}$ as N. We model the intensity function $\lambda(x)$ of the point process in the observation window $\mathcal{D}$ and denote the expected number of points in the set B as

$$\Lambda(B) = \mathbb{E} N(B) = \mathbb{E} \sum_{i=1}^{N} 1(x_i \in B) = \int_B \lambda(x) \, \mathrm{d}x.$$

We focus on Poisson point processes characterized by the counts $N(B) \sim \mathrm{Pois}\{\Lambda(B)\}$. With two types of points, such as nonextreme and extreme points, the point pattern is a superposition of the two single-type patterns, and the intensity is a sum $\lambda = \lambda_1 + \lambda_2$. The points of a specific type, say type 2, are generated by thinning the full point pattern, that is, by removing the

points of other types (here type 1) using the thinning probability $p(x) = \lambda_2(x)/\lambda(x)$, $x \in \mathcal{D}$. Extreme events, characterized as points x_i whose magnitude mark y_i exceeds a fixed high value $u(x_i)$, are obtained by thinning the full point pattern. Given a point pattern $\{x_1, \ldots, x_N\}$, $N \geq 1$, we define variables $E_i = \mathbb{I}\{y_i > u(x_i)\} \sim \text{Bernoulli}\{p(x_i)\}$. An independently thinned Poisson process (i.e., E_i are independent) is again a Poisson process.

3.3. *Spatiotemporal log-Gaussian Cox processes.* Log-Gaussian Cox processes (LGCPs) are Poisson processes with log-Gaussian intensity function $\lambda(x)$. This random specification of the intensity function can explain spatiotemporal variability not captured by deterministic parameters, and provides a natural framework for the Bayesian modeling of point processes with Gaussian process priors. Two major challenges arise for likelihood-based inference in LGCPs: (i) intensity functions are conceptually defined over continuous space; (ii) the Gaussian random effects lead to an intractable likelihood with no general closed-form expression. Challenge (ii) requires estimation techniques to handle latent variables; see Section 3.5. As to (i), without considering the marks, LGCPs have no general closed-form expression for their probability densities

$$(3) \qquad (x_1, \ldots, x_N) \mapsto \mathbb{E}_\lambda \exp\left(-\int_\mathcal{D} \lambda(x)\, dx\right) \prod_{i=1}^{n} \lambda(x_i).$$

Different approximation strategies allow numerical computation of the integral $\int_\mathcal{D} \lambda(x)\, dx$ for a given intensity function. We discretize the observation window, using the SAFRAN grid, and assume that the intensity function does not vary within pixel-day grid cells. Conditional on λ, the number of points observed in a cell C_k, $k = 1, \ldots, K$, is Poisson distributed, so estimating the LGCP corresponds to performing a (mixed) Poisson regression with log-link,

$$(4) \qquad N_k \mid \lambda_k \stackrel{\text{ind}}{\sim} \text{Pois}(|C_k|\lambda_k), \qquad \log(\lambda_k) = \mu_k, \quad k = 1, \ldots, K,$$

where λ_k is the value of the constant intensity function in cell C_k, $|C_k|$ is the Lebesgue volume of the cell, $\bigcup_k^K C_k = \mathcal{D}$ and $C_{k_1} \cap C_{k_2} = \varnothing$ if $k_1 \neq k_2$. The linear predictor μ_k is a sum of fixed and random effects. Likelihood-based inference for latent Gaussian processes is often based on Laplace approximation (Tierney and Kadane (1986)). In particular, the INLA framework assumes conditional independence of the observations, given the latent Gaussian predictor, and is thus well suited for LGCPs, where the Poisson observations N_k are conditionally independent, given μ_k (Illian, Sørbye and Rue (2012), Opitz et al. (2020)). Other approaches for numerically approximating the integral in (3) exist and typically use appropriately weighted sums $\sum_k \omega_k \lambda(\tilde{x}_k)$ with discretization points $\tilde{x}_k$ and weights $\omega_k > 0$, leading to variants of Poisson and logistic regression (e.g., the Berman–Turner 1992 device); see Baddeley et al. (2010).

3.4. *Data aggregation and subsampling schemes.* Spatiotemporal hierarchical modeling is notoriously computer-intensive due to large datasets and numerical challenges with covariances. The R-INLA implementation (Rue et al. (2017)) can handle up to several hundred thousand observations. Stable inferences may require compromises with respect to the complexity of the latent model and the number of observations, which jointly determine the size and sparsity of the Gaussian precision matrices, which influence computation times, memory requirements, and well-conditioned numerical behavior. Even stronger restrictions arise with the use of methods, such as Markov chain Monte Carlo (MCMC), to achieve approximation quality comparable to INLA (Taylor and Diggle (2014), van Niekerk et al. (2019)). Krainski et al. ((2018), Section 8.4) develop strategies for LGCPs by aggregating the events to larger

mapping units and lowering spatial-temporal resolution of random effects to decrease computation times, though this impedes the modeling of structures with small spatiotemporal scales.

Another way to cope with this issue of having too many observations is subsampling (Baddeley et al. (2010), Baddeley et al. (2014), Baddeley and Turner (2000), Rathbun (2013), Rathbun, Shiffman and Gwaltney (2007)), where the model is estimated using an appropriately reweighted subsample of data points, which keeps the loss of information small. Since maximum likelihood estimation is equivalent to maximizing the empirical expectation of the log-density of observations, a subsampling scheme is appropriate, provided it ensures a faithful approximation of this expectation. Subsampling in likelihood-based estimation can be interpreted as importance sampling (Tokdar and Kass (2010)): the original sample with observation weight unity is replaced by a subsample with larger observation weights. Weighted subsampling theory goes back to Horvitz and Thompson (1952).

The Poisson intensities $\lambda_k = \exp(\mu_k)$ ($k = 1, \ldots, K$) in (4) are the parameters to be estimated, and we need a subsample N_{k_j} with weights ω_j ($j = 1, \ldots, J$) such that the subsample likelihood is close to the full density (3). The sample size K exceeds 5 million, as there are over 1000 daily-replicated spatial pixels. To enable R-INLA-based estimation, we devise a stratified subsampling scheme to reduce the number of observations one hundredfold. Observations $N_k > 0$ are not subsampled since they are rare and highly informative; we keep them with unit weights. For the zero wildfire occurrence counts, we link subsampling to Poisson additivity. The likelihood contribution $\exp(-\lambda_k)^{\omega_k} = \exp(-\omega_k \lambda_k)$ with weight $\omega_k \in \mathbb{N}$ is equal to the likelihood of the sum of ω_k observations with count 0; the size of the initial sample is divided by the factor ω_k. The predictors (covariates, random effects) and, therefore, the intensities λ_k differ between different pixel-days k in our models, so Poisson additivity cannot be applied without additional approximations. However, the values of such predictors may often be very similar for cells located close in space and time, so we control the loss of information by subsampling that preserves a representative coverage of space and time.

We partition our data by years and pixels and then apply subsampling within each partition. The subsample contains two observations for each year-pixel combination. We thus obtain approximately 50,000 observations in the subsample, in line with the rule of thumb of Baddeley et al. (2014), Baddeley, Rubak and Turner (2015) that the subsample should be at least a factor four larger than the number of event points. The resulting models can be run on standard desktop computers (16 Gb of memory). Within pixel-year combinations we use nonuniform random sampling to overweight specific parts of the predictor space. For inference on the FWI-month interaction, we set different sampling probabilities for FWI values above and below the empirical FWI-quantile at p_{FWI} for each pixel-year. Values above the threshold are expected to correspond to more fire-prone conditions, and we overrepresent them by fixing sampling probabilities $p_{\text{SS}} = 0.9$ for FWI values below the threshold. Large FWI values represent only a relatively small fraction of the observations; our procedure ensures that the model accurately discriminates between occurrences and nonoccurrences of fires for such large FWI values. To appropriately identify seasonal effects, we choose the month among June–October at random. For instance, high FWI values tend to be less frequent in October, but uniform subsampling of months gives them more weight. This stratification ensures that we have approximately equal numbers of observations for each month in the subsample. With this scheme we obtain a positive sampling probability $p_k > 0$ for each observation N_k in (4), and likelihood weights are $\omega_k = 1/p_k$ for the selected observations. Simulation experiments (see Section E of the Supplementary Material (Koh et al. (2023))) motivated taking $(p_{\text{FWI}}, p_{\text{SS}}) = (0.7, 0.9)$.
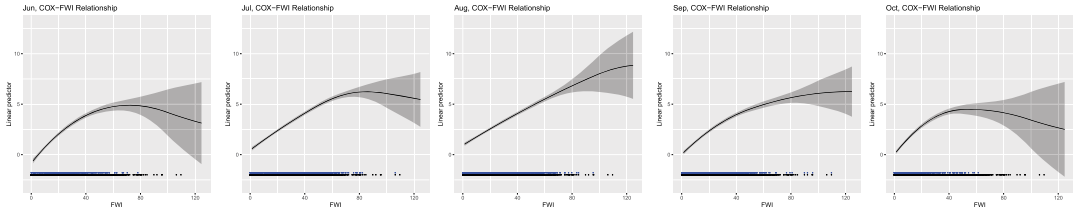
FIG. 3.   *Discretization of random effects with SPDE-based Gaussian prior processes. Left: Triangulation mesh of the study area (blue contours) for the SPDE approach. Neumann boundary conditions are set on the exterior (black) boundary to obtain a unique solution. The finite element solution defines a Gauss–Markov random vector with one variable in each node. Right: Histograms of FWI and FA values. The red points indicate where the spline knots are placed.*

3.5. *Fully Bayesian inference using INLA-SPDE.*   Integrated nested Laplace approximation (INLA Lindgren and Rue (2015), Opitz (2017), Rue, Martino and Chopin (2009)) is a Bayesian technique for fitting generalized additive models with Gaussian random effects. It uses astutely designed deterministic approximations for accurate posterior inference on model parameters, random effects and predictions conditional on data. INLA enables transfer of information across components, appropriate uncertainty assessment and estimation of shared effects. We implement penalized complexity priors (PC priors, Simpson et al. (2017b)) in our models to control the complexity of model components. Such priors penalize the distance (constructed using the Kullback–Leibler divergence) between the prior of a model component and a simpler baseline at a constant rate, that is, by using an exponential prior distribution for the distance.

Due to the large number of pixels in our problem, spatial Gaussian random effects and their conditional distributions must be tractable in this setting. We use the Matérn covariance function for random effects (denoted g), given as follows for two points s_1 and s_2:

$$\text{Cov}\{g(s_1), g(s_2)\} = \sigma^2 2^{1-\nu}(\kappa\|s_1 - s_2\|)^\nu K_\nu(\kappa\|s_1 - s_2\|)/\Gamma(\nu), \quad \sigma, \nu > 0$$

with Euclidean distance $\|\cdot\|$, gamma function Γ, modified Bessel function of the second kind K_ν, and standard deviation and smoothness parameters σ and ν. The empirical range at which the correlation drops to approximately 0.1 is $r = \sqrt{8\nu}/\kappa$. Numerically convenient representations by approximating Gauss–Markov random fields (GMRF, characterized by sparse precision, i.e., inverse covariance, matrices) are constructed by solving a stochastic partial differential equation (SPDE, Krainski et al. (2018), Lindgren, Rue and Lindström (2011)), where we fix the smoothness ν at unity. The discretization points of the triangulation in the SPDE approach are chosen as the nodes of a finite element representation (e.g., the triangulation of space for $d = 2$, or spline nodes for $d = 1$), which enables efficient inference for random effects representing spatial variation ($d = 2$) or nonlinear functions ($d = 1$ for the FWI and FA effects). Our spatial triangulation mesh in Figure 3 has 1114 nodes. It is less dense in the extended zone around the study area to ensure that SPDE boundary conditions have negligible influence on the study area. The four splines knots for FWI and FA are evenly spaced throughout the feature space.

4. Point processes with moderate and extreme marks.   Point processes govern the space-time point patterns of occurrences; size processes govern the moderate-level and extreme quantitative marks. We write N_{it} for the number of wildfire occurrences on day $t \in \{1, \ldots, n\}$ and over the 8 km $\times$ 8 km grid cell $i \in \{1, \ldots, 1143\}$ with centroid s_i, and $C_{i,t} \subset \mathcal{D}$ for the space-time cell with volume $|C_{i,t}| = 64$ (km^2 $\times$ day). If $N_{it} > 0$, we let $Y_{it} = (Y_{it,1}, \ldots, Y_{it,N_{it}}) \in (1, \infty)^{N_{it}}$ denote the corresponding quantitative marks. We write $z_k(s, t)$ ($k = 1, \ldots, K$) for known deterministic covariates.

We model data of escaped fires (> 1 ha), whose occurrence structure is captured by a regression component COX defining a LGCP. A logistic regression component BIN is used to classify fires into moderate (0) and large (1), according to their exceedance, or not above a fixed threshold u, that is, to provide the thinning of the point pattern and leave only extreme wildfires. Based on the statistical threshold selection procedure in Section 3.1 (see Figure 2), we consider a fire size $Y_{it,k}$ to be extreme if $Y_{it,k} > 79$ha ($k = 1, \ldots, N_{it}$); that is, we take $u = 79$. We write $\boldsymbol{R}_{it} = (R_{it,1}, \ldots, R_{it,N_{it}}) \in \{0, 1\}^{N_{it}}$ for the vector of binary exceedance indicators $R_{it,k} = \mathbb{I}(Y_{it,k} > u)$. Moderate wildfire sizes $Y_{it,k} \in (1, u]$ are modeled through a Beta regression component BETA applied to $(Y_{it,k} - 1)/(u - 1)$. The Beta distribution, usually parametrized by two shape parameters $a, b > 0$, is here parametrized through a precision parameter $\phi = a + b > 0$ and the mean $\mu_{it}^{\mathrm{BETA}} = a/(a + b) \in (0, 1)$ with logit-link function such that $a = \mu_{it}^{\mathrm{BETA}}\phi$ and $b = \phi(1 - \mu_{it}^{\mathrm{BETA}})$); it is a flexible location-shape family for interval-valued data and can be used with INLA. For large wildfires we use the extreme-value framework in Section 3.1 and model excesses $Y_{it} - u > 0$ above u through a GPD regression component GPD to characterize extreme wildfires. Following Opitz et al. (2018), we use a log-link function for the median μ_{it}^{GPD} of the GPD.

Some hyperparameters (e.g., precision parameters of priors for fixed effects) are fixed a priori, but those that may heavily influence the posterior model structure are estimated. The priors are fully detailed in Section F of the Supplementary Material (Koh et al. (2023)).

4.1. *Bayesian hierarchical multiresponse regression.* Our modeling assumptions in Section 3.3 give the linear COX predictor

$$\mu_{it}^{\mathrm{COX}} = \log \int_{\mathcal{C}_{it}} \lambda(s, t)\, \mathrm{d}(s, t) = \log \lambda(s_i, t) + \log |\mathcal{C}_{it}|.$$

We construct the system of regression equations in a Bayesian generalized additive mixed model (GAMM),

$$N_{it} \mid \mu_{it}^{\mathrm{COX}} \sim \mathrm{Poisson}\{\exp(\mu_{it}^{\mathrm{COX}})\},$$

$$R_{it,k} \mid \mu_{it}^{\mathrm{BIN}} \sim \mathrm{Bernoulli}\{\mathrm{logit}^{-1}(\mu_{it}^{\mathrm{BIN}})\}, \quad k = 1, \ldots, N_{it},$$

$$\{Y_{it,k} - u \mid R_{it,k} = 1, \mu_{it}^{\mathrm{GPD}}\} \sim \mathrm{GPD}\{\exp(\mu_{it}^{\mathrm{GPD}}), \xi\},$$

$$\{(Y_{it,k} - 1)/(u - 1) \mid R_{it,k} = 0, \mu_{it}^{\mathrm{BETA}}\} \sim \mathrm{Beta}\{\mathrm{logit}^{-1}(\mu_{it}^{\mathrm{BETA}}), \phi\};$$

$$\mu_{it}^{\mathrm{COMP}} = \sum_{k=1}^{K} g_k^{\mathrm{COMP}}\{z_k(s_i, t); \boldsymbol{\theta}^{\mathrm{COMP}}, \boldsymbol{\theta}^{\mathrm{SHR}}\},$$

$$\mathrm{COMP} \in \{\mathrm{COX, BIN, GPD, BETA}\};$$

$$\boldsymbol{\theta} = (\xi, \phi, \boldsymbol{\theta}^{\mathrm{COX}}, \boldsymbol{\theta}^{\mathrm{BIN}}, \boldsymbol{\theta}^{\mathrm{GPD}}, \boldsymbol{\theta}^{\mathrm{BETA}}, \boldsymbol{\theta}^{\mathrm{SHR}})$$

$$\sim \mathrm{Hyperpriors},$$

where terms g_k^{COMP} capture linear or nonlinear influence of the covariates in the corresponding model component. The specifics of $\boldsymbol{\theta}$ are discussed below.

The intensity function λ_{exc} of the point process of large fires satisfies $\lambda_{\mathrm{exc}}(s_i, t) \leq \lambda(s_i, t)$. The exceedance probability $\mathrm{logit}^{-1}\mu_{it}^{\mathrm{BIN}} = \lambda_{\mathrm{exc}}(s_i, t)/\lambda(s_i, t)$ defines the independent Bernoulli probability of the full point pattern in COX. Since $\lambda_{\mathrm{exc}}(s_i, t) = \exp(\mu_{it}^{\mathrm{BIN}}) \times \exp(\mu_{it}^{\mathrm{COX}})/\{1 + \exp(\mu_{it}^{\mathrm{BIN}})\}$ and typically $\exp(\mu_{it}^{\mathrm{BIN}}) \approx 0$, we obtain $\log \lambda_{\mathrm{exc}}(s_i, t) \approx \mu_{it}^{\mathrm{BIN}} + \mu_{it}^{\mathrm{COX}}$.

4.2. *Sharing latent effects.* For maximal flexibility we could incorporate mutually independent spatial effects into all model components. However, models would become overly complex, with too many spatial effects and hyperparameters to estimate and with high posterior uncertainties in the spatial effects of the BIN and GPD components due to the small number of large wildfires. We share spatial random effects between model components of the point and size processes with a preliminary model selection procedure (see Section 5.1) that avoids compromising the quality of model fit and predictions. We assign SPDE-based spatial GMRF priors $g^{\text{COX-BETA}}$, $g^{\text{COX-BIN}}$ and $g^{\text{BIN-GPD}}$ (recall Section 3.5) for the shared spatial effects. We use superscripts to indicate the two components into which we jointly incorporate an effect and write n to indicate the number of latent random variables for the corresponding effect (in superscript),

$$g^{\text{COX-BETA}}(s_i) \sim \mathcal{GP}_{\text{2D-SPDE}}(\boldsymbol{\omega}_1), \qquad n^{\text{COX-BETA}} = 1114,$$

$$g^{\text{COX-BIN}}(s_i) \sim \mathcal{GP}_{\text{2D-SPDE}}(\boldsymbol{\omega}_2), \qquad n^{\text{COX-BIN}} = 1114,$$

$$g^{\text{BIN-GPD}}(s_i) \sim \mathcal{GP}_{\text{2D-SPDE}}(\boldsymbol{\omega}_3), \qquad n^{\text{BIN-GPD}} = 1114,$$

where $\boldsymbol{\omega}_1$, $\boldsymbol{\omega}_2$, and $\boldsymbol{\omega}_3$ consist of r and σ with PC priors (Fuglstad et al. (2019)). Each shared effect is additively included in the linear predictor of the second component and then shared toward the first component with scaling factor $\beta \in \mathbb{R}$, with superscripts to denote the two components. We denote the vector of sharing-related hyperparameters by $\boldsymbol{\theta}^{\text{SHR}} = (\boldsymbol{\omega}_1, \boldsymbol{\omega}_2, \boldsymbol{\omega}_3, \beta^{\text{COX-BETA}}, \beta^{\text{COX-BIN}}, \beta^{\text{BIN-GPD}})$ and use flat, independent zero-centered Gaussian hyperpriors for the scaling factors.

Sharing allows modeling of residual spatial effect components that jointly affect multiple model responses, such as land use at the Wildland-to-Urban interface (Stewart et al. (2007)), where human activities intermingle with wildland vegetation. Accurate sharing improves parsimony of the model and borrows estimation strength for random effects across model components by simultaneously using different types of data. Expert knowledge should guide the choice of spatial effects to be shared between specific components; shared coefficients that differ from zero provide novel insight into the interplay of spatial structures.

4.3. *Prior structure of linear predictors.* We let $z_{\text{FWI}}(s_i, t)$ and $z_{\text{FA}}(s_i, t)$ denote the average FWI and FA on day t in grid cell i and by $a(t)$ and $m(t)$ the corresponding year and month of day t. Writing α for the intercept and g for the other GAMM components, the prior structure of the model component COX for escaped fire occurrences is

$$\mu_{it}^{\text{COX}} = \alpha^{\text{COX}} + g_1^{\text{COX}}(s_i) + \beta^{\text{COX-BETA}} g^{\text{COX-BETA}}(s_i)$$

$$+ \beta^{\text{COX-BIN}} g^{\text{COX-BIN}}(s_i)$$

$$+ g_2^{\text{COX}}\{z_{\text{FA}}(s_i, t)\} + g_3^{\text{COX}}\{z_{\text{FWI}}(s_i, t); m(t)\}$$

$$+ g_4^{\text{COX}}\{a(t)\} + g_5^{\text{COX}}\{m(t)\};$$

$$g_1^{\text{COX}}(s_i) \overset{\text{iid}}{\sim} \mathcal{N}\{0, 1/\tau_1\}, \qquad n_1^{\text{COX}} = 1143,$$

$$g_2^{\text{COX}}(\bullet) \sim \mathcal{GP}_{\text{1D-SPDE}}(\boldsymbol{\phi}_1), \qquad n_2^{\text{COX}} = 4,$$

$$g_3^{\text{COX}}(\bullet; m) \sim \mathcal{GP}_{\text{1D-SPDE}}(\boldsymbol{\phi}_2),$$

$$g_3^{\text{COX}}(z_{\text{FWI}}; \bullet) \sim \mathcal{GP}_{\text{RW1}}(1/\tau_2), \quad n_3^{\text{COX}} = 4 \times 5 = 20,$$

$$g_4^{\text{COX}}(\bullet) \sim \mathcal{GP}_{\text{RW1}}(1/\tau_3), \quad n_4^{\text{COX}} = 20,$$

$$g_5^{\text{COX}}(\bullet) \sim \mathcal{GP}_{\text{RW1}}(1/\tau_4), \quad n_5^{\text{COX}} = 5;$$

$$\boldsymbol{\theta}^{\text{COX}} = \{\alpha^{\text{COX}}, \boldsymbol{\phi}_1, \boldsymbol{\phi}_2, \tau_1, \tau_2, \tau_3, \tau_4\} \sim \text{Hyperpriors.}$$

Spatial occurrence hot spots (see Section C of the Supplementary Material (Koh et al. (2023))) may arise due to time-invariant landuse. Moreover, spatial variation may be shared from patterns in the BETA and BIN components through the components $g^{\text{COX-BETA}}(s_i)$ and $g^{\text{COX-BIN}}(s_i)$. The month and year effects, g_4^{COX} and g_5^{COX}, capture spatially homogeneous temporal variation in occurrence intensities. They are assigned first-order random-walk priors $\mathcal{GP}_{\text{RW1}}$ with a sum-to-zero constraint for identifiability so, for the yearly effect and for $a = 1995, \ldots, 2013$,

$$g_4^{\text{COX}}(a+1) - g_4^{\text{COX}}(a) \sim \mathcal{N}(0, 1/\tau_3), \qquad \sum_{i=1995}^{2014} g_4^{\text{COX}}(i) = 0.$$

The quadratic B-spline functions of FWI and FA are assigned priors $\mathcal{GP}_{\text{1D-SPDE}}$, constrained to zero at the left boundary 0 and to sum to zero, respectively. Most wildfires in the region are caused by human activity, possibly leading to a nonlinear relationship between FA and occurrence intensity, as dense forest areas are often exposed to low human activity. We model monthly variation of the nonlinear FWI effect through separate $\mathcal{GP}_{\text{1D-SPDE}}$-terms in g_3 for each month, linked across successive months with a $\mathcal{GP}_{\text{RW1}}$-structure in the prior model.

The regression equation used for the Bernoulli process is

$$\mu_{it}^{\text{BIN}} = \alpha^{\text{BIN}} + g^{\text{COX-BIN}}(s_i) + \beta^{\text{BIN-GPD}} g^{\text{BIN-GPD}}(s_i) + g_1^{\text{BIN}}\{z_{\text{FWI}}(s_i, t)\}$$
$$+ g_2^{\text{BIN}}\{z_{\text{FA}}(s_i, t)\} + g_3^{\text{BIN}}\{a(t)\};$$
$$g_k^{\text{BIN}}(\cdot) \sim \mathcal{GP}_{\text{1D-SPDE}}(\boldsymbol{\zeta}_k), \quad k = 1, 2, \qquad n_1^{\text{BIN}}, n_2^{\text{BIN}} = 5,$$
$$g_3^{\text{BIN}}(\cdot) \sim \mathcal{GP}_{\text{RW1}}(1/\tau_5), \qquad n_3^{\text{BIN}} = 5;$$
$$\boldsymbol{\theta}^{\text{BIN}} = \{\alpha^{\text{BIN}}, \boldsymbol{\zeta}_1, \boldsymbol{\zeta}_2, \tau_5\} \sim \text{Hyperpriors.}$$

The linear predictor of the Bernoulli probability has a simpler form than the occurrence component but still allows the capture of specific nonlinear effects of FWI and FA. In Figure 1 we discern hot-spot areas of large fire occurrences that differ substantially from the overall occurrence structure, and we aim to capture these residual effects through the shared spatial effects.

The prior structure for the two mixture components of quantitative marks is

$$\mu_{it}^{\text{BETA}} = \alpha^{\text{BETA}} + g^{\text{COX-BETA}}(s_i) + g_1^{\text{BETA}}\{z_{\text{FWI}}(s_i, t)\} + g_2^{\text{BETA}}\{z_{\text{FA}}(s_i, t)\},$$
$$\mu_{it}^{\text{GPD}} = \alpha^{\text{GPD}} + g^{\text{BIN-GPD}}(s_i) + g_1^{\text{GPD}}\{z_{\text{FWI}}(s_i, t)\} + g_2^{\text{GPD}}\{z_{\text{FA}}(s_i, t)\}$$
$$+ g_3^{\text{GPD}}\{a(t)\};$$
$$g_k^{\text{BETA}}(\cdot), g_k^{\text{GPD}}(\cdot) \sim \mathcal{GP}_{\text{1D-SPDE}}(\boldsymbol{\kappa}_k), \quad k = 1, 2, \qquad n_1^{\text{GPD}}, n_2^{\text{GPD}}, n_1^{\text{BETA}}, n_2^{\text{BETA}} = 5,$$
$$g_3^{\text{GPD}}(\cdot) \sim \mathcal{GP}_{\text{RW1}}(1/\tau_6), \qquad n_3^{\text{GPD}} = 5;$$
$$\boldsymbol{\theta}^{\text{MARK}} = \{\alpha^{\text{GPD}}, \alpha^{\text{BETA}}, \boldsymbol{\kappa}_1, \boldsymbol{\kappa}_2, \tau_6\} \sim \text{Hyperpriors.}$$

We assigned random-walk priors to year effects, included in some of the components (COX, BIN, GPD). In all components (BETA, BIN, COX, GPD), we model nonlinear relationships with respect to FWI or FA.

4.4. *Alternative model specifications.* We also consider size processes that do not model the moderate-level and extreme marks separately, that is, with no mixture representation of the size process. Similar models have been proposed in the literature (e.g., Joseph et al. (2019)), though without the sharing of random effects. We use either the Gamma distribution for the full range of marks, $Y_{it,k} \mid \mu_{it}^{\text{SIZE}} \sim \text{Gam}\{\exp(\mu_{it}^{\text{SIZE}}), \phi_{\text{Gam}}\}$, or the Normal distribution for the logarithmically transformed marks, $\log Y_{it,k} \mid \mu_{it}^{\text{SIZE}} \sim \mathcal{N}\{\exp(\mu_{it}^{\text{SIZE}}), \phi_{\mathcal{N}}\}$, where the distributions are parameterized by the link function μ_{it}^{SIZE} modeling the mean and precision parameters $\phi_{\text{Gam}} = \exp(\mu_{it}^{\text{SIZE}})^2 / \text{Var}(Y_{it,k})$ and $\phi_{\mathcal{N}} = 1/\text{Var}(\log Y_{it,k})$, respectively. In both cases

$$\mu_{it}^{\text{SIZE}} = \alpha^{\text{SIZE}} + g^{\text{SIZE-COX}}(s_i) + g_1^{\text{SIZE}}\{z_{\text{FWI}}(s_i, t)\} + g_2^{\text{SIZE}}\{z_{\text{FA}}(s_i, t)\}$$
$$+ g_3^{\text{SIZE}}\{a(t)\} + g^{\text{SIZE}}(s_i);$$
$$g_k^{\text{SIZE}}(\cdot) \sim \mathcal{GP}_{\text{1D-SPDE}}(\iota_k), \quad k = 1, 2, \qquad n_1^{\text{SIZE}}, n_2^{\text{SIZE}} = 5,$$
$$g_3^{\text{SIZE}}(\cdot) \sim \mathcal{GP}_{\text{RW1}}(1/\tau_7), \qquad n_3^{\text{SIZE}} = 5;$$
$$\boldsymbol{\theta}^{\text{SIZE}} = \{\alpha^{\text{SIZE}}, \iota_1, \iota_2, \tau_7\} \sim \text{Hyperpriors},$$

where the spatial effects $g^{\text{SIZE-COX}}(s_i)$ and $g^{\text{SIZE}}(s_i)$ are controlled by Matérn parameters $\boldsymbol{\omega}_4$ and $\boldsymbol{\omega}_5$, similar to those in Section 4.2.

5. Results.

5.1. *Model selection and comparison.* Estimation was carried out using the INLA-SPDE approach described in Section 3.5 by applying the subsampling scheme proposed in Section 3.4. In a preliminary analysis of the regression models described in Section 4, we used the Widely Applicable Information Criterion (WAIC, Watanabe (2010)) in a stepwise manner to compare nested models with different components in the regression equations (e.g., linear vs. nonlinear effects of explanatory variables) to choose their final forms.

We label the model with prior structure, detailed in Section 4.3 M1, and the model without spatial effects in the size and extreme occurrence components M2. We also considered other models from the recent wildfire modeling literature. We refer to model M2 but without monthly variation in the FWI effect as M3 which is similar to the model in Pimont et al. (2021). We let M4 and M5 denote the models with the point process model of M1 but with no mixture representation of the size process for which we use the log-Normal and Gamma distribution for the size distribution in the model structure described in Section 4.4. These models do not differentiate between extreme and nonextreme fires, but their response distributions were found to be good modeling candidates in Joseph et al. (2019), though their approach does not use shared random effects.

For the observed individual fires in the validation (2015–2018) periods, we generated posterior predictive distributions of each model based on 500 posterior simulations. First, we evaluated the models' ability to predict exceedances above the empirical 90% quantile of burnt areas using the AUC (Fawcett (2006)) and the Brier score (Brier (1950)). The severity threshold chosen here is sufficiently high for extreme risk assessment but low enough to retain enough observations to evaluate these scores with sufficient precision. As we considered the average predictive score across all observations in the validation set, we also computed the scaled Continuous Ranked Probability Score (sCRPS, Bolin and Wallin (2020)) which scales each observation's CRPS before calculating the average CRPS to account for nonidentical predictive distributions. For these analyses we kept the original locations of observed fires and simulated only from the size components. By combining posterior simulations of

Comparison of models using predictive scores (averaged over n observations) calculated with data from the validation period: sCRPS, Brier, and AUC scores for individual fires and sCRPS for the spatiotemporally aggregated burnt areas at month-département scale, based on 500 simulations of the posterior models, with p-values for a permutation test comparison with the best model M1. A lower score is better

		Model				
	Score	M1	M2	M3	M4	M5
Individual fires, $n = 823$	sCRPS	2.74	2.87	2.94	2.84	3.19
	p-value	–	<5%	<1%	<5%	<1%
	$Brier_{q90}$	0.0855	0.0868	0.0866	0.0944	0.0967
	p-value	–	<5%	6%	<1%	<1%
Dép-month, $n = 75$	$1 - AUC_{q90}$	0.3052	0.3502	0.3516	0.3184	0.3122
	p-value	–	<5%	<5%	40%	41%
	sCRPS	3.55	3.62	3.64	3.62	3.58
	p-value	–	7%	7%	9%	39%

the occurrence and size components, we also evaluated burnt area predictions aggregated at the month-département scale.

Table 1 shows good relative performance of M1 for all scores when evaluating wildfire predictions on the validation period. To better grasp the uncertainty in scores, we show p-values of a permutation test assessing the significance of negative values in the differences of scores between M1 and the other models, based on 2000 permutations. For the sCRPS of individual fires, the score differences are all significant at the 5% level. A general finding is that using sophisticated structures, such as the mixture representation of size processes, sharing and monthly variation of FWI effect, improves predictions; it further allows for the novel scientific insights presented in Section 6.

Comparison of M1 and M2 confirms the benefits of incorporating spatial random effects in the size model components in M1 using the sharing detailed in Section 4.2. M1 performs better than M2, and M2 and M3 give similar predictions for wildfire sizes and their aggregation. Model M1 performs better than M4 and M5, especially with respect to Brier and sCRPS scores, though in some cases improved scores have relatively high p-values. Models M4 and M5 perform better than M2 and M3 for some scores like the AUC and sCRPS at the month-département aggregation because of the additional sharing and spatial random effects in the size component, though they perform worse for the other scores due to having no components for extreme wildfires. Despite good scores of M4 and M5 on the training set (not shown), their worse results on the validation set suggest that the log-Normal and Gamma distributions for burnt areas do not predict the extremes in new data as well as M1 which does not show overfitting.

5.2. *Visual inspection of posterior predictive densities.* We also assess the predictive behavior of our chosen model M1 visually. First, we assess whether the size component correctly predicts extreme wildfires for specific départements. In Section A of the Supplementary Material (Koh et al. ((2023), Figure S1)), we use simulations from the posterior model at pixel-days, where fires have been observed to compare empirical and predicted excess probabilities over high thresholds, starting at 100 ha. Predictions are generally good, since most empirical exceedance probabilities fall within the interquantile range of simulations, except for the départements of Var and Haute-Corse in which there is slight underestimation at very large thresholds. These two départements have large continuous forest areas and saw unusually many large wildfires in the summer of 2017. Much of their land has acidic soils that favor

biomass production and is covered by tall and dense shrubland, so the fires that occurred in 2017 were harder to contain due to their higher heat release. Overall, the differences in tail behavior of fire-prone and less fire-prone regions are well captured.

Next, we compare the numbers of simulated and observed fires aggregated by year over the study region (Supplementary Material Koh et al. ((2023), Section A, Figure S2, left display)). Observed annual fire numbers for both test and training sets fall within the interquantile range of simulations for more than half of the study period. M1 captures the relatively high observed numbers of fires during 2001, 2003, and 1998 (training) and 2017 (test), while it also accurately predicts the sharp decrease in 2018.

Lastly, we jointly evaluate the size and occurrence components of our model M1. We aggregated simulated burnt areas by year, over the whole spatial region in Figure S2 (right display), and over départements in Figure S3. The global time trend in observed burnt areas is well captured in the Supplementary Material (Koh et al. (2023)), Figure S2, with interquantile coverage of 42%. M1 captures the exceptional peak in 2003, which is poorly predicted by M4 and M5 and the Firelihood model of Pimont et al. (2021). M1 also succeeds in accurately predicting the moderately high burnt areas in 2001 and 2017, and it generally discriminates well between fire conditions leading to small, moderate, large, and very large fire numbers. Figure S3 further shows that regional differences across départements are well captured by M1. Overall, our model captures spatiotemporal variation and provides satisfactory regionalized forecasts for operational purposes.

5.3. *Principal results of the main model M*1.

5.3.1. *Covariate effects.*   For the COX component, Figure 4 shows that the month-specific FWI effect varies significantly across months. For easier comparison, we have centred the curves so that the posterior mean is 0 for FWI = 0 in September. All the posterior means increase monotonically up to FWI of 75 and then flatten for higher values, especially at the beginning and end of the wildfire season, with a slight decrease of the curve toward the highest FWI.

The posterior partial effect of FA on the COX component in Figure 5 indicates a "bump"-shaped effect of FA, which is significant, based on pointwise credible intervals. High FA can be regarded as a good proxy for low human activity, while low FA means lack of fuel. Clearly, expected wildfire ignition numbers are not proportional to forest area.

As to temporal partial effects without spatial variation (Figure 5), the posterior year effect suggests a significant drop in wildfire activity after 2003, potentially related to policy changes after the exceptional 2003 events. The partial month effect (top-right display of Figure 5, corresponding to the intercept of its combined effect with FWI in Figure 4) is lowest at the start of the wildfire season and peaks in August.

As to the probability of occurrence of large fires (BIN), Figure 5 (bottom-middle display) highlights a strong positive posterior effect of FWI, increasing monotonically and significantly up to FWI values of around 75, before it dampens at very large FWI values, similar to the COX component: large wildfires are more frequent with moderate to high FWI values. The probability of large wildfires tends to increase with increasing FA in a grid cell (Figure 5, bottom-left display) which is reasonable because more FA fuel is available over large areas. The pointwise credible bounds of yearly effects across the study period suggest that the occurrence of large events was significantly higher around the peak in 2003.

The additive effects in the GPD and BETA components of the size distribution, seen in Figures 6, show similar posterior effects of FWI and forest area for extreme and moderate sizes. The posterior estimates imply that fires become larger when FWI increases until 60 but this effect flattens for higher FWI. Increasing FA leads to increasing wildfire size in

both components until 50%, then reaches a plateau after. For the year effect in the extreme component GPD, no clear trend arises, though 2003 has a significantly higher effect than 1998. The posterior mean estimate of the shape parameter ξ in the GPD component is 0.42; this suggests that the distribution of burnt areas remain heavy-tailed after incorporating our spatiotemporal effects into this component.

5.3.2. *Sharing effects induce correlated wildfire activity components.* Here, we focus only on the spatial effects shared between model components. The 95% credible intervals for the scaling parameters $\beta^{\text{COX-BETA}}$, $\beta^{\text{COX-BIN}}$, and $\beta^{\text{BIN-GPD}}$ do not cover 0; their posterior estimates for the triplet (2.5% quantile, mean, 97.5% quantile) are (6.4, 10.3, 14.0), $(-3.1, -1.8, -0.9)$, and (0.5, 1.0, 1.6), respectively. The posterior mean of $\beta^{\text{COX-BETA}}$ is positive and that of $\beta^{\text{COX-BIN}}$ is negative, confirming significant positive and negative sharing between COX and BETA, and COX and BIN, respectively; these findings provide new spatial insights for fire risk management described in Section 6. The posterior means for the effective range parameters of the shared spatial fields, $r^{\text{COX-BETA}}$, $r^{\text{COX-BIN}}$ and $r^{\text{BIN-GPD}}$, are 34.3 km, 26.2 km and 156.9 km, respectively. Posterior mean maps of their corresponding spatial random effects are shown in Section G of the Supplementary Material (Koh et al. (2023)).

Sharing decreases uncertainty by borrowing estimation strength between model components. The average lengths of 95% posterior credible intervals of variables constituting the random effect shrink by up to 30% (Figure 7) because of a higher observation-to-parameter ratio that enables us to better capture relevant spatial signals.

To identify the hot-spot regions of spatial random effects, we study credible sets for excursion regions (Bolin and Lindgren (2015)). We evaluate where the fields exceed or fall below the thresholds $u = 0.1$ and $-u$, respectively. These thresholds approximately correspond to
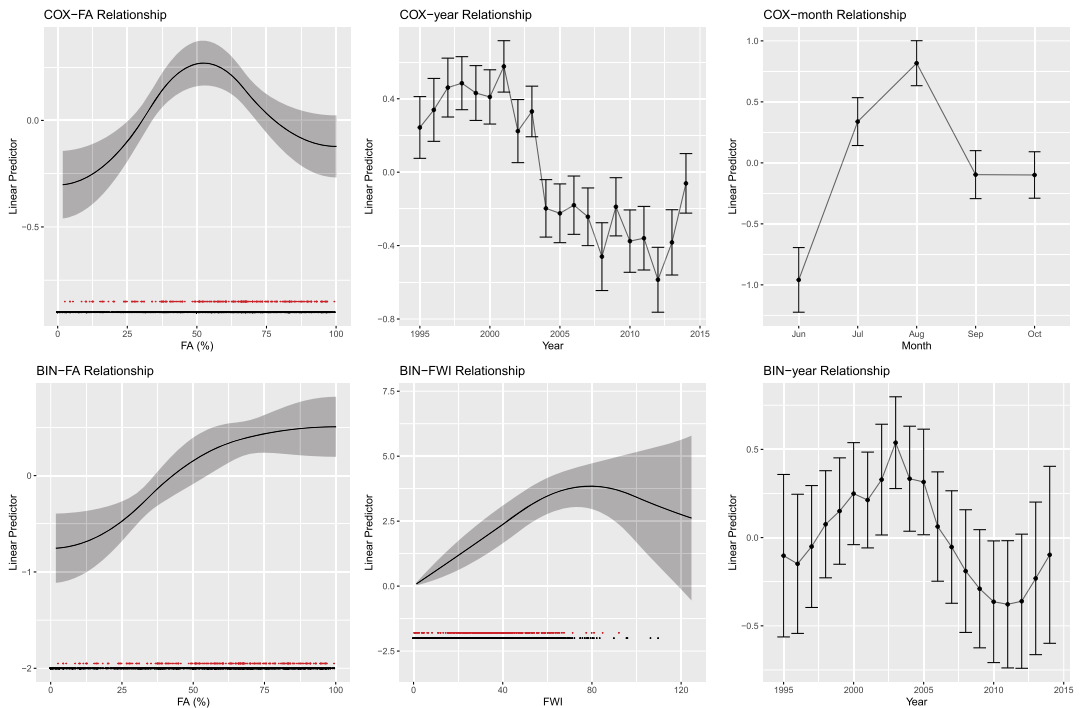

FIG. 4.   *Posterior estimates of $g_3^{\text{COX}}(\cdot; m) + g_5^{\text{COX}}(m)$, $m = 1, \ldots, 5$, the joint FWI-month effect, for June–October in the linear predictor of the point process (COX) component. The blanket of black and blue points at the bottom of each plot shows FWI values for pixel-days with fires in any month and the specific month, respectively.*
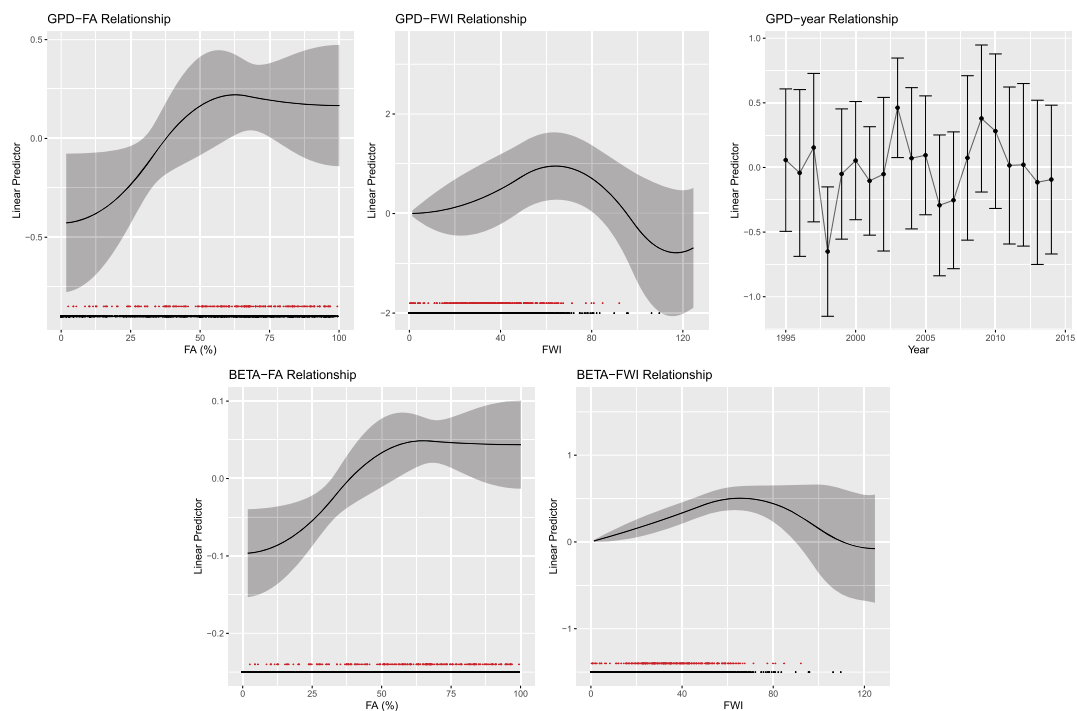
FIG. 5. *Posterior estimates of* $g_2^{\text{COX}}(\cdot)$ *(FA effect, top left panel),* $g_4^{\text{COX}}(\cdot)$ *(year effect, top middle panel),* $g_5^{\text{COX}}(\cdot)$ *(month effect, top right panel),* $g_2^{\text{BIN}}(\cdot)$ *(FA effect, bottom left panel),* $g_1^{\text{BIN}}(\cdot)$ *(FWI effect, bottom middle panel) and* $g_3^{\text{BIN}}(\cdot)$ *(year effect, bottom right panel) in the linear predictor of the point process (COX) component and large wildfire probability component (BIN). At the bottom of some displays, the blanket of black and red points shows FA/FWI values for pixel-days with moderate and large fires, respectively.*
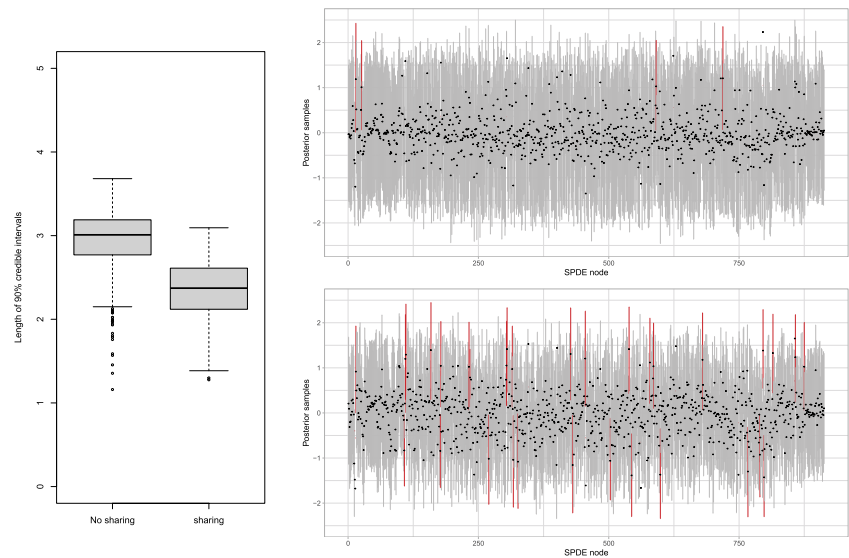
FIG. 6. *Panels as in Figure* 5. *Posterior estimates of* $g_2^{\text{GPD}}(\cdot)$ *(FA effect, top left),* $g_1^{\text{GPD}}(\cdot)$ *(FWI effect, top middle),* $g_3^{\text{GPD}}(\cdot)$ *(year effect, top right),* $g_2^{\text{BETA}}(\cdot)$ *(FA effect, bottom left) and* $g_1^{\text{BETA}}(\cdot)$ *(FWI effect, bottom right) in the linear predictor of the large wildfire size component (GPD) and moderate wildfire size component (BETA).*
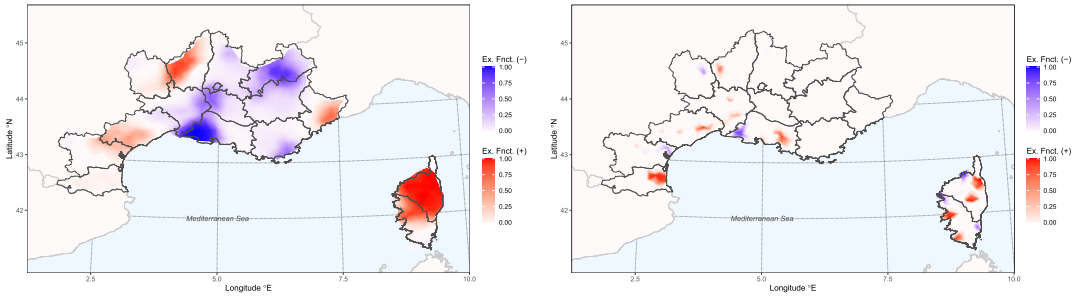
FIG. 7. *Lengths of the 90% credible intervals of spatial random effect variables at the SPDE triangulation nodes within the study area in the BIN component, based on 500 posterior simulations. Boxplots (left) and error bar plots for the models without (top right) and with sharing (bottom right). Red error bars indicate nodes where the intervals do not include zero.*

a 10% increase and decrease, respectively, on the scale of the response when taking into account the log or logistic link. The u-excursion set with probability α, $E_{u,\alpha}^{+}(X)$, is defined as the largest set for which the level u is exceeded at all locations in the set with probability $1 - \alpha$. The negative u excursion set with probability α, $E_{u,\alpha}^{-}(X)$, is defined as the largest set for which the process remains below the level $-u$ at all locations in the set with probability $1 - \alpha$. This approach determines the largest set contained in the exceedance set with a minimum probability threshold, and it assumes a parametric family for the exceedance sets. To visualize excursion sets simultaneously for all values of α, Bolin and Lindgren (2015) introduced the positive and negative excursion functions $F_u^+(s) = 1 - \inf\{\alpha \mid s \in E_{u,\alpha}^+\} \in [0, 1]$ and $F_u^-(s) = 1 - \inf\{\alpha \mid s \in E_{u,\alpha}^-\} \in [0, 1]$. Figure 8 highlights several hot-spot regions for the shared spatial effects, which we interpret with respect to wildfire management in Section 6.

6. New insights for wildfire science. Pimont et al. (2021) pointed out several critical divergences between simulations of their model and observed wildfire activity, and they have proposed hypotheses to explain them. The novel models developed here, especially M1, include components to estimate the sources of space-time variability conjectured by Pimont et al. (2021), leading to a better fit and more reliable inferences and predictions. Here, we outline the new insights.

6.1. *FWI and seasonal effects.* The estimated FWI effect on all wildfire components (COX, BIN, BETA, GPD) is nonlinear with a strong increase when moving from FWI $=$ 0 toward FWI ≈ 60–80, followed by a dampening and a slight decrease for extreme FWI values, though with wide credible bounds. Moreover, seasonal patterns emerge in the joint
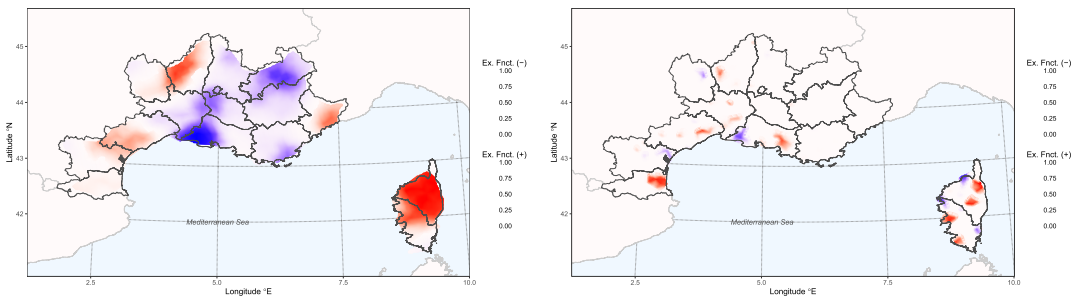

FIG. 8. *Excursion functions of posterior latent fields above 0.1 and below -0.1. Plots show* $\max\{F_{0.1}^+(\cdot), F_{0.1}^-(\cdot)\}$ *for the shared spatial random fields* $g^{\text{COX-BETA}}$ *(left panel) and* $g^{\text{COX-BIN}}$ *(right panel).*

FWI-month effect in the occurrence component COX. The practice of using FWI directly as a proxy for wildfire activity, without a nonlinear and seasonally varying transfer function as estimated here, would predict extreme wildfires badly and miss seasonally varying response of fire activity to this index.

This nonlinear, even decreasing, response to high FWI and seasonal biases can be attributed to the excessively sharp exponential response of FWI to wind speed in its upper range and to the limited ability of the Drought Code (a subcomponent of the FWI) to reproduce live fuel moisture dynamics in France (Ruffault et al. (2018)). In spring, vegetation budburst produces new foliage with a high water content that is maintained until the onset of the summer drought, typically in early July. The timing of periodic events in plant life cycles and stomatal control under drought might also explain why dynamics of soil and vegetation water content are unsynchronized at certain times. In our COX component we model not only a seasonal effect but also different responses of FWI across the months. The shapes of these monthly responses vary greatly, so seasonal variations cannot be handled solely through separate random effects. The response in August did not exhibit any saturation in the upper part of the FWI range, suggesting that higher values in mostly dry conditions correspond to increased fire activity; the contribution of wind to FWI could be adequate in these already-dry conditions. On the contrary, a flattening and notable decrease of the COX response to FWI was observed at FWI ≈ 45–50 for relatively moist conditions in June and October. This supports the hypotheses that the desynchronization of soil and fuel moistures caused by plant phenology in spring could be involved, and the response of the FWI to high wind would be inaccurate in such moist conditions. July and September, with their mixture of dry and moist days, show intermediate response levels to very high FWI. These findings confirm a need to develop better wildfire danger indices in the study region.

6.2. *Time trends during the study period.* The year 2003 was catastrophic in terms of fire sizes and burnt area. It has a pivotal role with a decrease of occurrence numbers and sizes afterward, as highlighted by the year component of our posterior model that captures temporal trends not explained through weather and land-cover related predictors. In 2003, a heat wave coincided with severe drought conditions, leading to an unusually high number of escaped fires (>1 ha), and of fires larger than 10 ha for several weeks, whose occurrence was not matched by very high values of FWI, due to its weaknesses outlined in Section 6.1. The drop in the estimated yearly effect after 2003 could be due to official policy measures that have slightly evolved after 2003, and to better prevention or suppression policies implemented by fire managers (Pimont et al. (2021)).

The yearly effect from the BIN component should interest wildfire managers, as it shows that the probability of observing a large fire increases toward the end of the study period after a decade of continuous decrease; the drivers of such subtle trends estimated by our model have not yet been identified. Our results also confirm those of Evin, Curt and Eckert (2018), who found no clear monotonic time trend for the probability of extreme fires.

6.3. *Shared spatial effects for improved regionalized predictions.* The shared spatial effects shown in Section 5.3 highlight regional differences in fire size distributions and provide quantitative interpretations of effects. They also reveal substantial regional variation in proportions of moderate and extreme fires. In particular, the sharing effect with significantly negative $\beta^{\text{COX-BIN}}$ allows for interpretation with respect to different wildland-to-urban interactions. The lowland area in the western Pyrénées-Orientales region, fairly densely populated with a large proportion of abandoned agricultural land intermixed with urban surfaces, appears to have high occurrence intensities, but its combustible area is strongly fragmented, so wildfires are mostly small. More fires than expected from weather/climate and forest area occur in densely populated or rural landscapes with significant human activities promoting fire

ignitions, while landscape fragmentation and landscape management reduce the likelihood of large fires. The COX-BETA sharing effect is highly positive in Corsica, where moderately large escaped fires become larger more often than elsewhere, perhaps due to longer arrival times of firefighters in remote Corsican forests and less frequent airborne firefighting. Moreover, extreme fires tend to be more frequent because of large contiguous forests. Further regional disparities in predictions are illustrated in the Supplementary Material (Koh et al. (2023)) where the right panel of Figure S4 highlights significant differences in threshold exceedance probabilities.

7. Conclusion. We have implemented a novel Bayesian spatiotemporal model for wildfire activity with specific components for extreme events, and with shared random effects to account for stochastic dependence among components not explained by covariates. The sophisticated structure of our fully Bayesian hierarchical model allows us to accurately disentangle the effects and interactions of various observed and unobserved drivers. The use of Gaussian random effects at high spatial resolution provides crucial benefits over purely frequentist generalized additive models, since fine-scale spatial variation and associated uncertainties can be identified properly.

Different sharing strategies respond to different considerations. If statistical stability is the focus, then sharing from well-identified model components toward those less informed by data is appropriate. If focus is on accurate inference of a specific component (e.g., extremes), then it is sensible to share effects from this component toward others. In both cases, component-specific effects without sharing remain important and should be included as far as data allow them to be estimated. In some applications, however, introducing common components by sharing is the only way to incorporate spatial effects in certain response variables. For example, had we chosen a threshold larger than 79 ha for large wildfires, we would have had even fewer observations available for the extreme fire size component. A separate spatial effect in this component would provide wider credible intervals than those in Figure 7 (top right) and be of less practical use.

Our findings improve decision support in wildfire management. Spatial and temporal random effects quantify the spatiotemporal variation in wildfire activity not explained by the available explanatory variables, that is, FWI and Forest Area. Our shared spatial effects explain how residual spatial variability is correlated across wildfire numbers and extreme sizes and allow us to provide maps of the significant disparities between regions. Weather forecasts, and the derived FWI forecasts, typically delivered at the département level, are currently the main components guiding fire detection/suppression resources and the temporary shutdown of forest areas to the public. However, FWI maps used for fire danger rating must be interpreted with care because of the strongly nonlinear and seasonal FWI effect on wildfire risk identified by our model. Moreover, the strong residual spatial effects estimated in our model could also hint at weather effects not captured by FWI. The precise regional forecasting of fire activity that our model can provide, especially of the expected number of fires and of the expected number and sizes of extreme fires, thus equips wildfire managers with additional objective criteria to aid decision-making.

Future work could better explore temporal trends (and how they vary over space) that are due to changes in land-use practices and wildfire management, by incorporating space-varying temporal random effects. Another extension of the model would be to include other informative predictors selected among a large variety of land-use, land-cover, and wildfire-management variables. This could better explain some of the spatiotemporal variation captured by the random effects of the current model. Another avenue for future models would be to include more information related to plant physiology and to specific subcomponents of the FWI (e.g., within the so-called Drought Code) to better explain the heterogeneous response

of wildfires to FWI. More generally, the construction of novel indices for fire-danger rating in the Mediterranean area remains an important task to be achieved through an appropriate mathematical combination of components of the FWI system, other fire-danger indices, and raw weather and land-related variables. For instance, the Fine Fuel Moisture Code (FFMC) component of the FWI system could provide improved danger rating for the ignition of wildfires, though it might be less relevant for their spread in Mediterranean summer conditions (and, therefore, for the exceedance of burnt area thresholds fixed at 1 ha or larger). An important constraint in this task is to maintain straightforward interpretations of how input variables contribute to new indices which excludes using black-box machine learning techniques.

SUPPLEMENTARY MATERIAL

Supplementary Material to "Spatiotemporal wildfire modeling through point processes with moderate and extreme marks" (DOI: 10.1214/22-AOAS1642SUPP; .zip). The supplementary material contains the following: a PDF document containing plots for the inspection of posterior predictive densities, plots showing regionalized predictions, kernel intensity plots for the spatial intensity function of the wildfire point process, mean excess plots for wildfire sizes, more details about the subsampling experiments, the hyperpriors used in the model and the posterior mean estimates of the spatial effects; moreover, R code and data files (including SAFRAN data with noise added) to reproduce the data analysis are provided.

REFERENCES

BADDELEY, A., BERMAN, M., FISHER, N. I., HARDEGEN, A., MILNE, R. K., SCHUHMACHER, D., SHAH, R. and TURNER, R. (2010). Spatial logistic regression and change-of-support in Poisson point processes. *Electron. J. Stat.* **4** 1151–1201. MR2735883 https://doi.org/10.1214/10-EJS581

BADDELEY, A., COEURJOLLY, J.-F., RUBAK, E. and WAAGEPETERSEN, R. (2014). Logistic regression for spatial Gibbs point processes. *Biometrika* **101** 377–392. MR3215354 https://doi.org/10.1093/biomet/ast060

BADDELEY, A., RUBAK, E. and TURNER, R. (2015). *Spatial Point Patterns*: *Methodology and Applications with R*. CRC Press/CRC Press, Boca Raton, FL.

BADDELEY, A. and TURNER, R. (2000). Practical maximum pseudolikelihood for spatial point patterns. *Aust. N. Z. J. Stat.* **42** 283–322. MR1794056 https://doi.org/10.1111/1467-842X.00128

BERMAN, M. and TURNER, T. R. (1992). Approximating point process likelihoods with GLIM. *J. R. Stat. Soc. Ser. C. Appl. Stat.* **41** 31–38.

BOLIN, D. and LINDGREN, F. (2015). Excursion and contour uncertainty regions for latent Gaussian models. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **77** 85–106. MR3299400 https://doi.org/10.1111/rssb.12055

BOLIN, D. and WALLIN, J. (2020). Scale dependence: Why the average CRPS often is inappropriate for ranking probabilistic forecasts. ArXiv Preprint. Available at arXiv:1912.05642.

BRIER, G. W. (1950). Verification of forecasts expressed in terms of probability. *Mon. Weather Rev.* **78** 1–3.

COLES, S. (2001). *An Introduction to Statistical Modeling of Extreme Values*. *Springer Series in Statistics*. Springer, London. MR1932132 https://doi.org/10.1007/978-1-4471-3675-0

CUI, W. and PERERA, A. H. (2008). What do we know about forest fire size distribution, and why is this knowledge useful for forest management? *Int. J. Wildland Fire* **17** 234–244.

CUMMING, S. (2001). A parametric model of the fire-size distribution. *Can. J. For. Res.* **31** 1297–1303.

DAVISON, A. C. and SMITH, R. L. (1990). Models for exceedances over high thresholds. *J. Roy. Statist. Soc. Ser. B* **52** 393–442. MR1086795

DE ZEA BERMUDEZ, P., MENDES, J., PEREIRA, J. M. C., TURKMAN, K. F. and VASCONCELOS, M. J. P. (2009). Spatial and temporal extremes of wildfire sizes in Portugal (1984–2004). *Int. J. Wildland Fire* **18** 983–991.

EVIN, G., CURT, T. and ECKERT, N. (2018). Has fire policy decreased the return period of the largest wildfire events in France? A Bayesian assessment based on extreme value theory. *Nat. Hazards Earth Syst. Sci.* **18** 2641–2651.

FAWCETT, T. (2006). An introduction to ROC analysis. *Pattern Recogn. Lett.* **27** 861–874.

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., BUTRY, D. T., GUMPERTZ, M. L. and PRESTEMON, J. P. (2006). Spatio-temporal analysis of wildfire ignitions in the St Johns River water management district, Florida. *Int. J. Wildland Fire* **15** 87–97.

HORVITZ, D. G. and THOMPSON, D. J. (1952). A generalization of sampling without replacement from a finite universe. *J. Amer. Statist. Assoc.* **47** 663–685. MR0053460

ILLIAN, J. B., SØRBYE, S. H. and RUE, H. (2012). A toolbox for fitting complex spatial point process models using integrated nested Laplace approximation (INLA). *Ann. Appl. Stat.* **6** 1499–1530. MR3058673 https://doi.org/10.1214/11-AOAS530

JONES, M. W., SMITH, A., BETTS, R., CANADELL, J. G., PRENTICE, I. C. and LE QUÉRÉ, C. (2020). ScienceBrief review: Climate change increases the risk of wildfires. In *Critical Issues in Climate Change Science* (C. Le Quéré, P. Liss and P. Forster, eds.).

JOSEPH, M. B., ROSSI, M. W., MIETKIEWICZ, N. P., MAHOOD, A. L., CATTAU, M. E., ST. DENIS, L. A., NAGY, R. C., IGLESIAS, V., ABATZOGLOU, J. T. et al. (2019). Spatiotemporal prediction of wildfire size extremes with Bayesian finite sample maxima. *Ecol. Appl.* **29** e01898.

KOH, J., PIMONT, F., DUPUY, J.-L and OPITZ, T. (2023). Supplement to "Spatiotemporal wildfire modeling through point processes with moderate and extreme marks." https://doi.org/10.1214/22-AOAS1642SUPP

KRAINSKI, E. T., GÓMEZ-RUBIO, V., BAKKA, H., LENZI, A., CASTRO-CAMILO, D., SIMPSON, D., LINDGREN, F. and RUE, H. (2018). *Advanced Spatial Modeling with Stochastic Partial Differential Equations Using R and INLA*. CRC Press/CRC, Boca Raton.

LINDGREN, F. and RUE, H. (2015). Bayesian spatial modelling with R-INLA. *J. Stat. Softw.* **63**.

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

MENDES, J. M., DE ZEA BERMUDEZ, P. C., PEREIRA, J., TURKMAN, K. F. and VASCONCELOS, M. J. P. (2010). Spatial extremes of wildfire sizes: Bayesian hierarchical models for extremes. *Environ. Ecol. Stat.* **17** 1–28. MR2594933 https://doi.org/10.1007/s10651-008-0099-3

MORITZ, M. A. (1997). Analyzing extreme disturbance events: Fire in los padres national forest. *Ecol. Appl.* **7** 1252–1262.

NORTHROP, P. J. and COLEMAN, C. L. (2014). Improved threshold diagnostic plots for extreme value analyses. *Extremes* **17** 289–303. MR3207718 https://doi.org/10.1007/s10687-014-0183-z

OPITZ, T. (2017). Latent Gaussian modeling and INLA: A review with focus on space-time applications. *J. SFdS* **158** 62–85. MR3720130

OPITZ, T., BAKKA, H., HUSER, R. and LOMBARDO, L. (2022). High-resolution Bayesian mapping of landslide hazard with unobserved trigger event. *Ann. Appl. Stat.* **16** 1653–1675. https://doi.org/10.1214/21-AOAS1561

OPITZ, T., BONNEU, F. and GABRIEL, E. (2020). Point-process based Bayesian modeling of space-time structures of forest fire occurrences in Mediterranean France. *Spat. Stat.* **40** 100429. MR4181143 https://doi.org/10.1016/j.spasta.2020.100429

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** 441–462. MR3855716 https://doi.org/10.1007/s10687-018-0324-x

PENG, R. D., SCHOENBERG, F. P. and WOODS, J. A. (2005). A space-time conditional intensity model for evaluating a wildfire hazard index. *J. Amer. Statist. Assoc.* **100** 26–35. MR2166067 https://doi.org/10.1198/016214504000001763

PEREIRA, J. M. C. and TURKMAN, K. F. (2019). Statistical models of vegetation fires: Spatial and temporal patterns. In *Handbook of Environmental and Ecological Statistics*. *Chapman & Hall/CRC Handb. Mod. Stat. Methods* 401–420. CRC Press, Boca Raton, FL. MR3889904

PIMONT, F., FARGEON, H., OPITZ, T., RUFFAULT, J., BARBERO, R., MARTIN-STPAUL, N., RIGOLOT, E. I., RIVIÈRE, M. and DUPUY, J.-L. (2021). Prediction of regional wildfire activity in the probabilistic Bayesian framework of firelihood. *Ecol. Appl.* e02316.

PREISLER, H. K., BRILLINGER, D. R., BURGAN, R. E. and BENOIT, J. (2004). Probability based models for estimation of wildfire risk. *Int. J. Wildland Fire* **13** 133–142.

RATHBUN, S. L. (2013). Optimal estimation of Poisson intensity with partially observed covariates. *Biometrika* **100** 277–281. MR3034343 https://doi.org/10.1093/biomet/ass069

RATHBUN, S. L., SHIFFMAN, S. and GWALTNEY, C. J. (2007). Modelling the effects of partially observed covariates on Poisson process intensity. *Biometrika* **94** 153–165. MR2367829 https://doi.org/10.1093/biomet/asm009

RÍOS-PENA, L., KNEIB, T., CADARSO-SUÁREZ, C., KLEIN, N. and MAREY-PÉREZ, M. (2018). Studying the occurrence and burnt area of wildfires using zero-one-inflated structured additive beta regression. *Environ. Model. Softw.* **110** 107–118.

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

RUE, H., RIEBLER, A., SØRBYE, S. H., ILLIAN, J. B., SIMPSON, D. P. and LINDGREN, F. K. (2017). Bayesian computing with INLA: A review. *Annu. Rev. Stat. Appl.* **4** 395–421.

RUFFAULT, J., MARTIN-STPAUL, N., PIMONT, F. and DUPUY, J.-L. (2018). How well do meteorological drought indices predict live fuel moisture content (LFMC)? An assessment for wildfire research and operations in Mediterranean ecosystems. *Agric. For. Meteorol.* **262** 391–401.

SCHOENBERG, F. P., PENG, R. and WOODS, J. (2003). On the distribution of wildfire sizes. *Environmetrics* **14** 583–592.

SERRA, L., JUAN, P., VARGA, D., MATEU, J. and SAEZ, M. (2013). Spatial pattern modelling of wildfires in Catalonia, Spain 2004–2008. *Environ. Model. Softw.* **40** 235–244.

SERRA, L., SAEZ, M., JUAN, P., VARGA, D. and MATEU, J. (2014). A spatio-temporal Poisson hurdle point process to model wildfires. *Stoch. Environ. Res. Risk Assess.* **28** 1671–1684.

SIMPSON, D., RUE, H., RIEBLER, A., MARTINS, T. G. and SØRBYE, S. H. (2017b). Penalising model component complexity: A principled, practical approach to constructing priors. *Statist. Sci.* **32** 1–28. MR3634300 https://doi.org/10.1214/16-STS576

STEWART, S. I., RADELOFF, V. C., HAMMER, R. B. and HAWBAKER, T. J. (2007). Defining the wildland–urban interface. *J. For.* **105** 201–207.

TAYLOR, B. M. and DIGGLE, P. J. (2014). INLA or MCMC? A tutorial and comparative evaluation for spatial prediction in log-Gaussian Cox processes. *J. Stat. Comput. Simul.* **84** 2266–2284. MR3223624 https://doi.org/10.1080/00949655.2013.788653

TIERNEY, L. and KADANE, J. B. (1986). Accurate approximations for posterior moments and marginal densities. *J. Amer. Statist. Assoc.* **81** 82–86. MR0830567

TOKDAR, S. T. and KASS, R. E. (2010). Importance sampling: A review. *Wiley Interdiscip. Rev.: Comput. Stat.* **2** 54–60.

TONINI, M., PEREIRA, M. G., PARENTE, J. and OROZCO, C. V. (2017). Evolution of forest fires in Portugal: From spatio-temporal point events to smoothed density maps. *Nat. Hazards* **85** 1489–1510.

TURKMAN, K. F., TURKMAN, M. A. A. and PEREIRA, J. M. (2010). Asymptotic models and inference for extremes of spatio-temporal data. *Extremes* **13** 375–397. MR2733939 https://doi.org/10.1007/s10687-009-0092-8

VAN NIEKERK, J., BAKKA, H., RUE, H. and SCHENK, L. (2019). New frontiers in Bayesian modeling using the INLA package in R. ArXiv Preprint. Available at arXiv:1907.10426.

VAN WAGNER, C. E. (1977). Conditions for the start and spread of crown fire. *Can. J. For. Res.* **7** 23–34.

WANG, X., WOTTON, B. M., CANTIN, A. S., PARISIEN, M.-A., ANDERSON, K., MOORE, B. and FLANNIGAN, M. D. (2017). cffdrs: An R package for the Canadian forest fire danger rating system. *Ecol. Process.* **6** 5.

WATANABE, S. (2010). Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory. *J. Mach. Learn. Res.* **11** 3571–3594. MR2756194

WOOLFORD, D. G., MARTELL, D. L., McFAYDEN, C. B., EVENS, J., STACEY, A., WOTTON, B. M. and BOYCHUK, D. (2021). The development and implementation of a human-caused wildland fire occurrence prediction system for the province of Ontario, Canada. *Can. J. For. Res.* **51** 303–325.

XI, D. D. Z., TAYLOR, S. W., WOOLFORD, D. G. and DEAN, C. B. (2019). Statistical models of key components of wildfire risk. *Annu. Rev. Stat. Appl.* **6** 197–222. MR3939518 https://doi.org/10.1146/annurev-statistics-031017-100450

XU, H. and SCHOENBERG, F. P. (2011). Point process modeling of wildfire hazard in Los Angeles County, California. *Ann. Appl. Stat.* **5** 684–704. MR2840171 https://doi.org/10.1214/10-AOAS401