



Gradient boosting with extreme-value theory for wildfire prediction

Jonathan Koh^{1,2} 

Received: 3 January 2022 / Revised: 28 October 2022 / Accepted: 31 October 2022 /
Published online: 21 January 2023
© The Author(s) 2023

Abstract

This paper details the approach of the team *Kohrrelation* in the 2021 Extreme Value Analysis data challenge, dealing with the prediction of wildfire counts and sizes over the contiguous US. Our approach uses ideas from extreme-value theory in a machine learning context with theoretically justified loss functions for gradient boosting. We devise a spatial cross-validation scheme and show that in our setting it provides a better proxy for test set performance than naive cross-validation. The predictions are benchmarked against boosting approaches with different loss functions, and perform competitively in terms of the score criterion, finally placing second in the competition ranking.

Keywords Cross-validation · Generalized Pareto distribution · Gradient boosting · Loss likelihood · Machine learning · Wildfire prediction

AMS 2000 Subject Classifications 62G32 · 62J99 · 62P12

1 Introduction

Wildfires occur in every season of the year and are a natural phenomenon of the forest ecosystem, important for clearing out decayed vegetation and helping plants to reproduce. However, they have the potential to become conflagrations—intense, destructive fires—that may have huge environmental and ecological impacts. Apart

✉ Jonathan Koh
jonathan.koh@unibe.ch

¹ Institute of Mathematics, EPFL, Lausanne, Switzerland

² Institute of Mathematical Statistics and Actuarial Science, Oeschger Centre for Climate Change Research, University of Bern, Bern, Switzerland

from human casualties, these fires can lead to substantial economic losses; global insured claims due to wildfire events have increased dramatically in recent years, from below \$10 billion in 2000–2009 to \$45 billion in the subsequent decade¹.

Wildfires are complex dynamic processes: their occurrences and behaviour are the product of interconnected factors that include the ignition source, fuel composition, topography and the weather. For example, the wind plays a big role in the spread and ease of fire containment, but its effect is magnified in the presence of accumulated biomass on a hilly boreal forest after a prolonged dry spell. The modelling of wildfires is made even more complicated by the need to model the Wildland-to-Urban interface (Stewart et al. 2007), as 90% of fires are caused by human activity.

An important measure of wildfire impact and size is the burned area of wildfire events, commonly used by government agencies and aggregated at different spatial and temporal scales for reporting purposes, e.g., National Interagency Fire Center (2021). Although there is a positive relationship between wildfire counts and sizes, with perfect agreement of the lowest value of zero in both variables, more wildfire occurrences in a region do not necessarily translate to higher burned areas. In 2020, nearly 26,000 wildfires burned approximately 9.5 million acres (ac) in the west, compared with the over 33,000 fires that burned just under 0.7 million ac in the east. Similarly, although the numbers of wildfires have fallen since the 1990s, the average annual acreage burned since 2000 has more than doubled.

Many statistical approaches have been developed to aid in wildfire prevention and risk mitigation, with most studies modelling wildfire occurrences and sizes separately (Taylor et al. 2013; Xi et al. 2019; Pereira and Turkman 2019; Jain et al. 2020), though models that identify latent factors affecting both have been proposed (e.g., Koh et al. 2023). Point processes are natural models for the spatiotemporal pattern of occurrences (Peng et al. 2005; Genton et al. 2006; Tonini et al. 2017; Opitz et al. 2020). Cumming (2001), Cui and Perera (2008) and Pereira and Turkman (2019) suggested modelling fire sizes with various probability distributions. As data usually show heavy-tailed behaviour, only a small fraction of wildfires account for the vast majority of the area burned. Obvious candidates to capture this stem from extreme-value theory, such as the generalized Pareto distribution (GPD) for modelling threshold exceedances (De Zea Bermudez et al. 2009; Turkman et al. 2010; Pereira and Turkman 2019).

Both Bayesian and frequentist methods have been used for explanatory modeling, the former predominantly for hierarchical mixed effect models (Joseph et al. 2019; Pimont et al. 2020; Koh et al. 2023) and the latter within the generalized additive modelling (GAM, Wood 2017) framework (Preisler et al. 2004; Brillinger et al. 2006; Vilar et al. 2010; Woolford et al. 2011). Covariates include weather variables such as humidity, temperature, precipitation and meteorologically-based fire danger indices such as the Canadian Fire Weather Index (van Wagner 1977). When available, land-use or locally observed anthropogenic variables like population density and the distance to the nearest train line are used to help to explain

¹ <https://www.swissre.com/risk-knowledge/mitigating-climate-risk/yet-more-wildfires.html>

human-induced occurrences; spatiotemporal random effects have been incorporated as surrogates for these variables.

If accurate prediction is of primary interest, then machine learning (ML) techniques offer an attractive alternative to the statistical modelling approaches described above. Since the 1990s, the surge in the availability of data and covariates has spurred the use of these techniques to predict wildfire behaviour. Jain et al. (2020) found 127 journal papers or conference proceedings published up to the end of 2019 on ML applied to fire occurrence, susceptibility and risk; of these adversarial neural networks (ANN) were the most prominent (eg., Dutta et al. 2013; Shidik and Mustofa 2014; Liang et al. 2019). For wildfire occurrences, most studies focus on classification tasks instead of count modelling. Sakr et al. (2010) used meteorological variables with support vector machines to predict a four-class fire risk index based on the daily number of fires in Lebanon. Dutta et al. (2013) compared ten ANN based cognitive imaging systems to determine the relationship between monthly fire incidence and climate for Australia. Mitsopoulos and Mallinis (2017) and Xie and Peng (2019) showed that ensemble learning methods like random forests and boosting trees performed well in estimating area burned or classifying wildfire sizes in Greece and Portugal, respectively.

Gradient boosting techniques (Friedman 2001) have exploded in popularity over the last decade, in part due to the development and dissemination of open-source packages such as `gbm` (Greenwell et al. 2020) and `xgboost` (Chen and Guestrin 2016). A key ingredient of gradient boosting is the loss function used to train these models, and choices for these functions have largely been restricted to those that emphasize good prediction of the distributional bulk instead of the tails. For example, squared loss, the default when modelling wildfire sizes, implicitly presupposes that the target functional of interest is the conditional mean given the covariates, which may be inappropriate if the focus is predominantly on extreme values. The Poisson loss is popular for modelling wildfire counts within a grid cell, but the zero-inflated nature and potentially heavy tails of count distributions suggest that this loss may be unsuitable. Evaluation metrics should also reflect the non-linear impact of wildfire events; e.g., in many cases, predicting a false negative occurrence is much costlier than predicting a false positive.

As ML methods are prone to overfitting, it is imperative to evaluate models with held-out datasets using robust validation schemes. A realistic approach in the forecasting context (when there are no trends) is to leave out the most recent portion of the dataset (eg., Woolford et al. 2011; Dutta et al. 2013; Joseph et al. 2019; Koh et al. 2023). Dutta et al. (2013) explored different combinations of training-testing splits to identify the best possible paradigm to maximize the generalization capability of their ANN architecture. K -fold cross-validation is also popular (Shidik and Mustofa 2014; De Angelis et al. 2015; Mitsopoulos and Mallinis 2017; Xie and Peng 2019), but it may give overly optimistic evaluations for spatially dependent data (Roberts et al. 2017). An alternative is spatial cross-validation (Pohjankukka et al. 2017), but it is still unclear how best to construct spatial folds in this context, and doing so anyway ignores relevant inter-variable or time dependencies.

Our work aims to tackle the limitations of the studies mentioned above, and does so in the context of the Extreme Value Analysis 2021 data challenge (Opitz 2021). We develop novel gradient boosting models trained with loss functions appropriate

for predicting extreme values. Our chosen model for fire counts is a discrete generalized Pareto distribution (dGPD, Shimura 2012) relying on a covariate-dependent parameter that models a chosen high quantile of the distribution, and a shape hyperparameter selected by cross-validation. The dGPD provides added flexibility in modelling the upper tail of the count distribution, especially when compared to other more standard count distributions like the Poisson. The model for fire sizes has three components and covariate-dependent probabilities. The first component models the probability of observing no fires, and the others model the probabilities of observing medium-sized and extreme fires. We approximate the conditional distribution above a high threshold with a GPD, and the conditional distribution below the threshold with a truncated log-gamma distribution.

To improve our models, we also engineered new covariates that incorporate more spatial information into the climatic and land-use covariates provided by averaging them across neighbouring grid cells each month. With a smart imputation method for replacing missing data, we also use the wildfire counts as a covariate when predicting wildfire sizes, and vice versa.

We develop a spatiotemporal cross-validation scheme that provides a better proxy for our models' test set performance than the naive scheme. This involves fitting a space-time latent Gaussian model to pseudo-binary observations that indicate whether a grid cell was masked in a particular month, and then simulating from the fitted model to generate folds of train-test regimes.

In the remainder of the paper, we first explore the data on wildfires and their covariates, and then introduce the problem set out by the data challenge in Sect. 2. We provide general background on extreme-value theory and gradient boosting and on how to combine them in Sect. 3. Our spatiotemporal cross-validation scheme is developed in Sect. 3.4 and the specific model structure is detailed in Sect. 4. We highlight the prediction of wildfire activity components in Sect. 4.2, and compare them to related and competing approaches. We conclude with a discussion and outlook in Sect. 5.

2 Data and exploratory analyses

The Extreme Value Analysis 2021 data challenge dealt with the prediction of monthly wildfire counts and burned areas at 3503 grid cells across the contiguous US over the period 1993–2015. As fuel moisture is an integral of past precipitation and evaporation mediated by soil field capacity, temporal scales longer than hourly or daily (e.g., monthly in our case) are appropriate for predicting fire risk from climatic covariates.

The data comprise the monthly numbers of wildfires (CNT) and the aggregated burned area (BA) in each grid cell based on a $0.5^\circ \times 0.5^\circ$ grid of longitude and latitude coordinates (roughly 55 km \times 55 km) covering the study area, from March to September each year. Figure 1 shows that the grid cells with the highest averaged CNT tend to be clustered towards the west (California) and southeast (North and South Carolina) corners of the study region, while clusters in the west (near the border of Idaho and Nevada), southwest (Arizona, New Mexico and Texas), and southeast (Florida) are observed for BA.

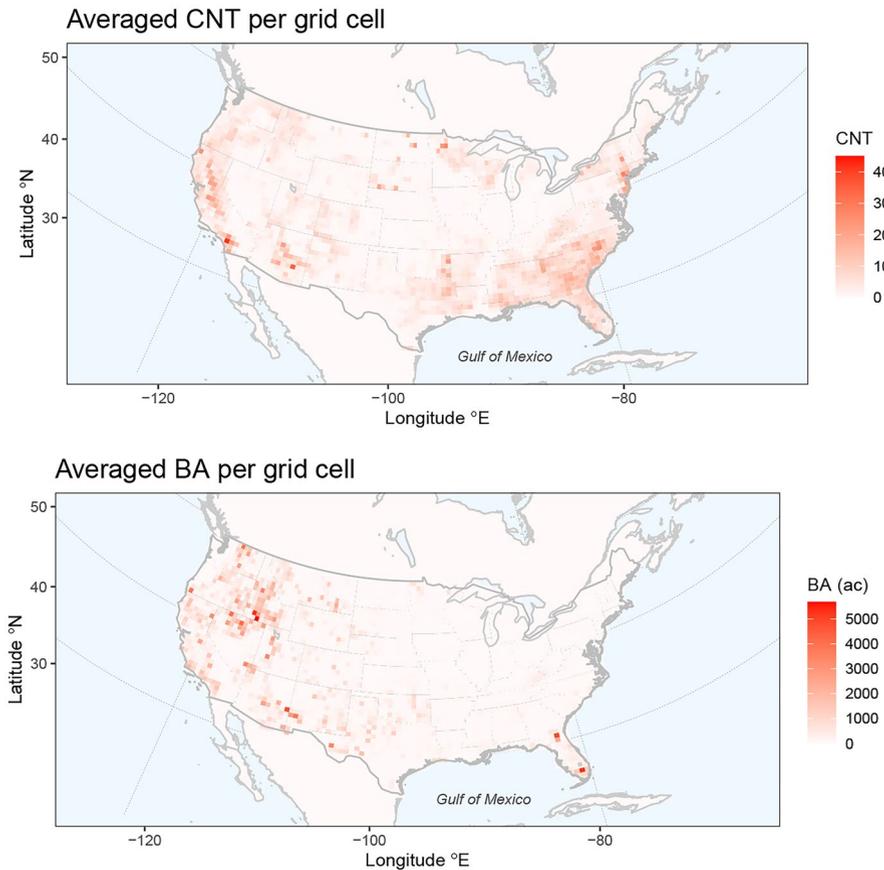


Fig. 1 Maps of CNT (top) and BA (bottom) averaged across all months and years for each grid cell

Thirty-five auxiliary variables related to land cover, weather and altitude are provided at the same spatial and temporal resolution, and can be used for modelling. Figure 2 hints at the importance of some of these variables, such as meteorological covariate 5 (clim5; potential evaporation, measured in meter water equivalent – mwe) and land cover covariate 7 (lc7; tree needleleaf evergreen closed to open, measured in % of the grid cell), for predicting high BA. However it also indicates that their effect differs over space, and that the interactions between these covariates may be important for predicting large burned areas. The Rocky Mountain Area and Great Basin, two regions defined formally as ‘Geographic Area Coordination Centers’ by the United States Department of Agriculture, have empirical exceedance probabilities that respond differently to both covariates. For instance, the probability tends to decrease, then increase (after the 75% quantile) with potential evaporation in the Great Basin, while the opposite holds for the Rocky Mountain Area, though the associated large uncertainties suggest that there is substantial heterogeneity within each region.

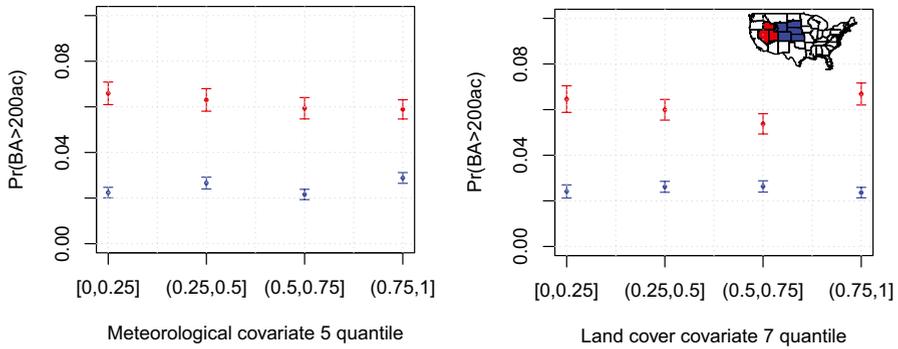


Fig. 2 Empirical frequency $\Pr(\text{BA} > 200\text{ac})$ with 95% error bars, as a function of the potential evaporation (in m, left) and tree needleleaf evergreen closed to open (in %, right) covariates, grouped by observations within four empirical quantile ranges, for grid cells within the Rocky Mountain Area (blue) and Great Basin (red) coordination centers. The two regions are highlighted in blue and red in the map inserted in the right panel

The original dataset has no gaps, but missing data were artificially created to compare predictive approaches; the full dataset was split into training and testing subsets to evaluate participants’ test scores. No data were masked in the odd years, but 80,000 observations of each variable were masked in the even years. Figure 3 shows that the spatial and temporal positions of test data are clustered in space, and the test grid cells for BA and CNT are correlated. This masking is reminiscent of a real-world situation in which two related and spatially dependent processes could render both CNT and BA unobservable at small spatial clusters every month (e.g., from a potentially spatially dependent measurement error), and one could only use the available covariates and responses from the surrounding non-masked regions for prediction. For example, datasets generated using satellite-based remote sensing of wildfires have known misclassification issues, very often due to cloud occlusion.

The evaluation metrics used for the competition (see Opitz 2021) require estimates of the probability $\Pr(\text{BA} < u_{\text{BA}})$ and $\Pr(\text{CNT} < u_{\text{CNT}})$ for 28 thresholds u_{CNT} and u_{BA} . The metrics are variants of weighted ranked probability scores and put relatively strong weight on good prediction in the extremes of the distributions of counts and burned areas. As

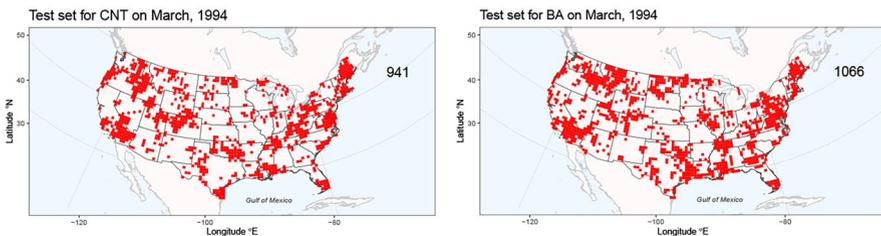


Fig. 3 The test set grid cells (in red) for CNT (left) and BA (right) in March 1994. The number at the top right indicates the number of masked grid cells

such, we expect that models that emphasize accurate modelling of the largest counts and burned areas will perform better, and we achieve this by appealing to extreme-value theory.

3 Methodology

3.1 Extreme-value theory

The generalized Pareto distribution (GPD) arises asymptotically for excesses above a large threshold of a random variable $X \sim F$, when the distribution F satisfies mild regularity conditions. Let $x^* = \sup\{x : F(x) < 1\}$. The excess distribution above $u < x^*$ can be approximated (Davison and Smith 1990) as

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

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 towards a finite x^* , for $\xi < 0$. When the approximation (1) is exact asymptotically (i.e., when $u \rightarrow x^*$), we say that the random variable X lies in the maximum domain of attraction of a generalized Pareto distribution with shape parameter ξ , written $X \in \text{MDA}_\xi$.

3.1.1 Positive discrete random variables

We say that a discrete non-negative random variable Y lies in the discrete maximum domain of attraction, $Y \in \text{dMDA}_\xi$, if there exists a random variable $X \in \text{MDA}_\xi$ with $\xi \geq 0$ such that $\Pr(Y \geq k) = \Pr(X \geq k)$, for $k = 0, 1, \dots$. The random variable X is called the *version* of Y , and many popular discrete distributions such as the geometric, Poisson and negative binomial distributions lie in dMDA_ξ .

For $Y \in \text{dMDA}_\xi$ and large integers u ,

$$\begin{aligned} \Pr(Y - u = k \mid Y \geq u) &\approx \text{GPD}_{\sigma, \xi}(k + 1) - \text{GPD}_{\sigma, \xi}(k) \\ &= (1 + \xi k / \sigma)^{-1/\xi} - (1 + \xi(k + 1) / \sigma)^{-1/\xi}, \end{aligned} \tag{2}$$

where the last term is the probability mass function of the discrete generalized Pareto distribution (Shimura 2012). Several studies have used this distribution to model count data; Prieto et al. (2014) modelled numbers of road accidents and Hitz et al. (2017) modelled numbers of extreme tornadoes per outbreak and multiple births.

3.2 Gradient tree boosting

The generic gradient boosting estimator (Bühlmann and Hothorn 2007) is a sum of base procedure estimates. Regression trees (CART, Breiman et al. 1984) are popular base procedures, as they include non-linear covariate interactions by construction, and are invariant under monotone transformations of covariates, so the user need not search for good data transformations.

Let $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}$, $\mathbf{x}_i \in \mathbb{R}^p$, $y_i \in \mathbb{R}$, $i = 1, \dots, n$, be a dataset with n observations and p covariates. A *binary split* of the covariate space uses a splitting variable indexed by $j \in \{1, \dots, p\}$ and a split point $v \in \mathbb{R}$ to partition the space into the pair of half-spaces $\{\mathbf{x} \in \mathbb{R}^p : x_j \leq v\}$ and $\{\mathbf{x} \in \mathbb{R}^p : x_j > v\}$, where x_j is the j -th component of \mathbf{x} .

By successive binary splits, a regression tree partitions the covariate space into a set of L disjoint regions A_1, \dots, A_L , and fits a simple model such as a constant in each region. The regions created by the splits are called nodes; a terminal node is called a *leaf* and an interior node is called a *branch*. We index leaves by $l \in \{1, \dots, L\}$, with leaf l representing region A_l . The simplest tree is one with two leaves, known as a stump. A learning algorithm needs to decide the tree structure, i.e., the splitting variables and split points.

Suppose that L leaves with regions A_1, \dots, A_L have been chosen and we model the response as a *score* $c_l \in \mathbb{R}$ in each region. A regression tree is a function

$$f(\mathbf{x}_i) = \sum_{l=1}^L c_l \mathbb{I}(\mathbf{x}_i \in A_l),$$

where \mathbb{I} is the indicator function. A gradient tree boosting model uses T such trees to model the *boosting estimate*

$$\hat{\theta}_i = \sum_{t=1}^T f_t(\mathbf{x}_i), \quad f_t \in \mathcal{T}, \quad (3)$$

where \mathcal{T} is the space of regression trees. In this paper, $\hat{\theta}_i$ will always represent a parameter estimate in a given model for the conditional distribution of y_i given \mathbf{x}_i . The boosting estimate using stumps will be additive in the original covariates, because every base estimate is a function of a single covariate. A boosting model that has trees with at most L leaves has interactions of order at most $L - 2$. Thus, constraining the maximum number of nodes in the base procedure controls the complexity of the model.

Gradient tree boosting learns the set of trees used in (3) by minimizing a regularized objective function in a greedy iterative fashion; at each iteration we add the tree that most improves our model according to an objective function \mathcal{O} . More precisely, let $\hat{\theta}_i^{(j)}$ be the boosting estimate for the i -th observation at the j -th iteration. We add a tree f_j at each iteration to minimize

$$\mathcal{O}^{(j)} = \sum_{i=1}^n \mathcal{L}\{y_i, \hat{\theta}_i^{(j-1)} + f_j(\mathbf{x}_i)\} + \Omega(f_j),$$

where $\Omega(f_j) = \eta L^{(j)} + \lambda \|\mathbf{c}\|^2/2$, \mathcal{L} is a differentiable loss function, $L^{(j)}$ is the number of leaves in the tree f_j and $\mathbf{c} \in \mathbb{R}^{L^{(j)}}$ is the corresponding vector of scores. The regularization term Ω is added to penalize the complexity of each tree, and the positive parameters η and λ control the penalization. The form of Ω is simple enough to allow parallel computation (Chen and Guestrin 2016).

Using a second-order approximation for the objective (Friedman et al. 2000) gives

$$\mathcal{O}^{(j)} \simeq \sum_{i=1}^n \left\{ \mathcal{L}(y_i, \hat{\theta}_i^{(j-1)}) + g_i f_j(\mathbf{x}_i) + h_i f_j^2(\mathbf{x}_i) \right\} + \Omega(f_j), \tag{4}$$

where $g_i = \partial \mathcal{L}(y_i, \hat{\theta}_i^{(j-1)})/\partial \hat{\theta}_i^{(j-1)}$ and $h_i = \partial^2 \mathcal{L}(y_i, z_i)/\partial z_i^2|_{z_i=\hat{\theta}_i^{(j-1)}}$. We then minimize (4) with respect to the tree structure and weight vector \mathbf{c} .

Let $I_l = \{i : \mathbf{x}_i \in A_l\}$ denote the *instance set* of leaf l . For a fixed tree structure with regions $A_1, \dots, A_{L^{(j)}}$, the optimal weights \mathbf{c}^* can easily be found and have components

$$c_l^* = -\frac{\sum_{i \in I_l} g_i}{\sum_{i \in I_l} h_i + \lambda}, \quad l = 1, \dots, L^{(j)}.$$

Plugging the weights \mathbf{c}^* into (4) and removing the term that does not depend on f_j gives

$$\tilde{\mathcal{O}}^{(j)} = -\frac{1}{2} \sum_{l=1}^{L^{(j)}} \frac{(\sum_{i \in I_l} g_i)^2}{\sum_{i \in I_l} h_i + \lambda} + \eta L^{(j)}, \tag{5}$$

which can be used as a scoring function to measure the quality of the tree structure, a role similar to the impurity score in Breiman et al. (1984).

Assume that a split has been performed, and let I_L and I_R denote the instance sets of the left and right leaves from this split. Define $I = I_L \cup I_R$. The loss reduction from this split is

$$G = \frac{1}{2} \left\{ \frac{(\sum_{i \in I_L} g_i)^2}{\sum_{i \in I_L} h_i + \lambda} + \frac{(\sum_{i \in I_R} g_i)^2}{\sum_{i \in I_R} h_i + \lambda} - \frac{(\sum_{i \in I} g_i)^2}{\sum_{i \in I} h_i + \lambda} \right\} - \eta, \tag{6}$$

and (6) is used for evaluating candidate split variables and points.

As it is impossible to enumerate all possible tree structures, most existing tree boosting implementations, such as in `scikit-learn` (Pedregosa et al. 2011) and `gbm` (Greenwell et al. 2020), use greedy algorithms that start from a single leaf and iteratively add branches to the tree based on (6), until the gain for the best split is negative. Here we use the greedy algorithm implemented in `xgboost` (called the approximate algorithm with weighted quantile sketch Chen and Guestrin 2016, Appendix A), that further reduces computational cost

and parallelizes computations when the data do not fit into memory. This algorithm proposes candidate splitting points from the empirical quantiles of each covariate, instead of considering all possible splitting points for each variable. We also subsample a proportion $s \in [0, 1]$ of the covariates at each iteration, like in the random forest algorithm (Breiman 2001), to further prevent overfitting and to accelerate parallel computation. For the full algorithmic details, we refer the reader to Chen and Guestrin (2016).

Convexity of \mathcal{L} is a desirable feature that would guarantee that a unique global minimum exists; if this does not hold in practice and one is concerned about the algorithm being stuck at local minimums, then a potential solution is to rerun the algorithm multiple times from different initial boosting estimates $\hat{\theta}_i^{(0)}, i = 1, \dots, n$, and select the run that gives the lowest loss. If \mathcal{L} is convex and there are enough iterations, then the choice of the initial boosting estimate will have minimal effect on the overall prediction. A natural choice is to set $\hat{\theta}_i^{(0)} = \hat{\theta}_j^{(0)} = \hat{\theta}$, for all $i \neq j$, where $\hat{\theta} = \arg \min_{\theta} \sum_{i=1}^n \mathcal{L}(y_i, \theta)$, i.e., this is the common parameter value that minimises the loss across all observations, unconditionally on the covariates.

As gradient tree boosting is an ensemble method combining predictions from many trees, the exact relationship between covariates (or their interactions) and the response is difficult to determine. Suppose we treat the covariates as random, and let $\mathbf{X}_{\mathcal{S}}$ denote the random subvector (of size $l < p$) of the full covariate vector $\mathbf{X} = (X_1, \dots, X_p)^{\top}$, indexed by the set $\mathcal{S} \subset \{1, \dots, p\}$. Let \mathcal{C} be the set complementary to \mathcal{S} . For a predictive function f at a fixed point $\mathbf{x} \in \mathbb{R}^l$, the partial dependence function (Friedman 2001),

$$f_{\mathcal{S}}(\mathbf{x}) = \mathbb{E}_{\mathbf{X}_{\mathcal{C}}} f\{\mathbf{x}, \mathbf{X}_{\mathcal{C}}^{\top}\},$$

can be estimated by Monte Carlo as

$$\hat{f}_{\mathcal{S}}(\mathbf{x}) = \frac{1}{n} \sum_{i=1}^n f\{\mathbf{x}, \mathbf{x}_{i\mathcal{C}}^{\top}\}, \quad (7)$$

where $\mathbf{x}_{1\mathcal{C}}, \dots, \mathbf{x}_{n\mathcal{C}}$ are realisations of $\mathbf{X}_{\mathcal{C}}$ from the training data.

Metrics used to rank covariates in terms of their importance include the *coverage* for a chosen covariate, which is the sum of the second order gradients h_i from (5), in each node which uses this covariate, standardized by dividing by the sum of the metrics for all other covariates (so the resulting metric is a proportion). The *gain* metric represents the fractional contribution of a chosen covariate to the model based on the total gain of all the splits involving this covariate, measured by G in (6); it is the total improvement of the model in terms of the objective function, from the branches that include the covariate. In both cases a higher proportion implies a more important predictive variable.

The loss function \mathcal{L} in (4) strongly governs the type of models that we fit, and we discuss choices for this next.

3.3 Loss functions

3.3.1 For wildfire counts

The most common statistical model for count data is the Poisson distribution, and many existing boosting implementations use a scaled version of the simplified Poisson loss as their default loss function for counts. Given n monthly wildfire counts in a grid cell, y_1, \dots, y_n , this loss is

$$\mathcal{L}_{\text{Pois}}(y_i, \hat{\theta}_i) = y_i \log\{y_i / \exp(\hat{\theta}_i)\} - y_i + \exp(\hat{\theta}_i), \tag{8}$$

where Stirling’s approximation $\log(y_i!) \approx y_i \log(y_i) - y_i$ is used and the boosting estimate $\hat{\theta}_i$ models the log mean of the i -th Poisson count. The terms that do not depend on $\hat{\theta}_i$ can be dropped when minimising this loss in an optimisation procedure. Although (8) is also the most common choice for modelling wildfire counts in the literature, the zero-inflation and potentially heavy tails of count distributions suggest that it may be unsuitable.

Instead, if we let $\alpha = 1/\xi > 0$ and $\lambda = \sigma\alpha$ in (2), we motivate a new loss function for counts from extreme-value theory, the discrete generalized Pareto (dGPD) loss

$$\mathcal{L}_{\text{dGPD}}(y_i, \hat{\theta}_i) = \{1 + \exp(\hat{\theta}_i)y_i\}^{-\alpha} - \{1 + \exp(\hat{\theta}_i)(y_i + 1)\}^{-\alpha}. \tag{9}$$

The boosting estimate $\hat{\theta}_i$ models the logarithm of the rescaled scale parameter λ_i . If $\alpha > 1$, the predicted mean of the i -th count is

$$\hat{m}_i = \sum_k^{\infty} 1 / \{1 + \exp(\hat{\theta}_i)k\}^{\alpha}, \tag{10}$$

and otherwise the mean does not exist.

The first and second derivatives of (9) with respect to the boosting estimate, i.e.,

$$\begin{aligned} g_i^{\text{dGPD}} &= -\alpha \{1 + \exp(\hat{\theta}_i)y_i\}^{-\alpha-1} \{\exp(\hat{\theta}_i)y_i\} \\ &\quad + \alpha \{1 + \exp(\hat{\theta}_i)(y_i + 1)\}^{-\alpha-1} \{\exp(\hat{\theta}_i)(y_i + 1)\}, \\ h_i^{\text{dGPD}} &= -\alpha(-\alpha - 1) \{1 + \exp(\hat{\theta}_i)y_i\}^{-\alpha-2} \{\exp(\hat{\theta}_i)y_i\}^2 \\ &\quad - \alpha \{1 + \exp(\hat{\theta}_i)y_i\}^{-\alpha-1} \{\exp(\hat{\theta}_i)y_i\} \\ &\quad + \alpha(-\alpha - 1) \{1 + \exp(\hat{\theta}_i)(y_i + 1)\}^{-\alpha-2} \{\exp(\hat{\theta}_i)(y_i + 1)\}^2 \\ &\quad + \alpha \{1 + \exp(\hat{\theta}_i)(y_i + 1)\}^{-\alpha-1} \{\exp(\hat{\theta}_i)(y_i + 1)\}, \end{aligned}$$

are used in the second-order approximation of the objective in (4) and are essential for determining split variable and split point candidates when building trees with (6).

3.3.2 For wildfire sizes

Assuming normality of the response conditional on the covariates may be inappropriate if the distributional tail decays more slowly than exponential. Moreover, burned areas cannot be negative. Modelling the log-transformed burned areas with a normal distribution addresses the latter issue, though doing so still excludes conditional distributions with Pareto-like tails.

Another approach to modelling the full distribution of wildfire sizes is to use a mixture, by first choosing an appropriately high threshold and then fitting different parametric distributions to burned areas below and above that threshold using different loss functions. To additionally handle the zero-inflation of wildfire sizes, we can left truncate the sizes below the high threshold at zero. This mixture approach models burned areas by splitting the distribution into three groups: zero, intermediate and extreme sizes.

To model the monthly burned area in a grid cell, y_i , below a chosen threshold $u > 0$, the negative log-loss likelihood of a truncated distribution could be used. The probability density function of a right truncated gamma distribution is

$$f(x) = \begin{cases} \frac{(\mu/k)^{-k} x^{k-1} \exp(-xk/\mu)}{\gamma(k, ku/\mu)} & , \quad 0 < x \leq u, \\ 0 & , \quad x > u, \end{cases}$$

where $u > 0$ is the right truncation, $\mu > 0$ is the rescaled scale parameter, $k > 0$ is the shape parameter and $\gamma(k, s) = \int_0^s t^{k-1} \exp(-t) dt$, $s > 0$, is the lower incomplete gamma function. Modelling $\log(\mu)$ with the boosting estimate and dropping the terms that do not depend on $\hat{\theta}_i$ gives

$$\begin{aligned} \mathcal{L}_{\text{trGamma}}(y_i, \hat{\theta}_i) &= k\hat{\theta}_i + y_i k / \exp(\hat{\theta}_i) + \log \gamma \{k, ku / \exp(\hat{\theta}_i)\}, \\ g_i^{\text{trGamma}} &= \exp(\hat{\theta}_i) \{k / \exp(\hat{\theta}_i) - y_i k / \exp(\hat{\theta}_i)^2 \\ &\quad + \gamma' \{k, ku / \exp(\hat{\theta}_i)\} / \gamma \{k, ku / \exp(\hat{\theta}_i)\}\}, \\ h_i^{\text{trGamma}} &= \exp(\hat{\theta}_i)^2 \left\{ -k / \exp(\hat{\theta}_i)^2 + 2y_i k / \exp(\hat{\theta}_i)^3 \right. \\ &\quad \left. + \frac{\gamma \{k, ku / \exp(\hat{\theta}_i)\} \gamma'' \{k, ku / \exp(\hat{\theta}_i)\} - \gamma' \{k, ku / \exp(\hat{\theta}_i)\}^2}{\gamma \{k, ku / \exp(\hat{\theta}_i)\}^2} \right\} \\ &\quad + g_i^{\text{trGamma}}, \end{aligned}$$

where $\gamma' \{k, ku / \exp(\hat{\theta}_i)\}$ and $\gamma'' \{k, ku / \exp(\hat{\theta}_i)\}$ are the first and second derivatives of $\gamma \{k, ku / \exp(\hat{\theta}_i)\}$ with respect to $\hat{\theta}_i$, and have closed forms (see Supplement Sect. 6).

To model only the excesses above a threshold u , we can use the GPD in (1), and assume that $\xi > 0$, since burned areas tend to be heavy-tailed. If we reparameterize and model the logarithmic $\kappa \in [0, 1]$ quantile of the excesses with the

boosting estimate, i.e., $\hat{\theta}_i = \log\{(1 - \kappa)^{-\xi} - 1\}\sigma_i/\xi$, we can define the generalized Pareto (GPD) loss and obtain its derivatives

$$\begin{aligned} \mathcal{L}_{\text{GPD}}(y_i, \hat{\theta}_i) &= \frac{\xi + 1}{\xi} \log \left[1 + \frac{y_i\{(1 - \kappa)^{-\xi} - 1\}}{\exp(\hat{\theta}_i)} \right] + \log \left[\frac{\xi \exp(\hat{\theta}_i)}{\{(1 - \kappa)^\xi - 1\}} \right], \\ g_i^{\text{GPD}} &= -\frac{f'\{y_i, \exp(\hat{\theta}_i), \xi\}}{f\{y_i, \exp(\hat{\theta}_i), \xi\}}, \\ h_i^{\text{GPD}} &= -\frac{f\{y_i, \exp(\hat{\theta}_i), \xi\}f''\{y_i, \exp(\hat{\theta}_i), \xi\} - f'\{y_i, \exp(\hat{\theta}_i), \xi\}^2}{f\{y_i, \exp(\hat{\theta}_i), \xi\}^2} \end{aligned}$$

where f' and f'' are the first and second derivatives of the reparameterized probability density function $f\{y_i, \exp(\hat{\theta}_i), \xi\}$ given in Supplement Sect. 6.

3.3.3 For wildfire size classification

Adopting the mixture modelling approach to wildfire sizes requires us to model the probability that a fire belongs to each size component $1, \dots, C$, where C is the number of components. Let $y_i = (y_{i,1}, \dots, y_{i,C})$ denote the vector of wildfire size component indicators, where $y_{i,c} = 1$ if the i -th fire size is in component c , and otherwise $y_{i,c} = 0$. We can model the probability of each class with the boosting estimate using the softmax function $\sigma : \mathbb{R}^C \rightarrow [0, 1]^C$ defined by

$$\sigma(z)_i = \frac{\exp(z_i)}{\sum_{j=1}^C \exp(z_j)}, \quad i = 1, \dots, C, \quad z = (z_1, \dots, z_C) \in \mathbb{R}^C. \tag{11}$$

The generalization of the logistic (Cox 1958) loss to multiple classes is the cross-entropy loss, which can be reweighted to give

$$\mathcal{L}_{\text{CE}}(y_i, \hat{\theta}_i) = -w_i \sum_{c=1}^C y_{i,c} \log \left\{ \exp(\hat{\theta}_{i,c}) / \sum_{d=1}^C \exp(\hat{\theta}_{i,d}) \right\}, \tag{12}$$

where the vector of component probabilities is modelled by applying (11) to the boosting estimate $\hat{\theta}_i = (\hat{\theta}_{i,1}, \dots, \hat{\theta}_{i,C})^\top$, and the weights w_1, \dots, w_n could be chosen to improve predictions in unbalanced classification tasks.

3.4 A spatiotemporal cross-validation scheme

The use of k -fold cross-validation generally presupposes independent replicates, so it would produce optimistic predictive performance estimates in our setting because data points that are geographically closer will have stronger dependencies. To address this, we first study the spatiotemporal process leading to grid cells being masked, which we call the *masking process*. Figure 3 hints at either a common or two inter-correlated spatially dependent latent processes governing the observed masking processes for CNT and BA, which we model with a

common latent Gaussian process (Rasmussen and Williams 2005). We then fit a Bernoulli model to observations arising from the masking process, and simulate observations from the model to generate cross-validation folds.

We consider only the months m with masked observations, and let M denote the number of those months and D denote the number of grid cells. Let $R_{d,m}^{\text{CNT}}$ and $R_{d,m}^{\text{BA}}$ denote the binary 0-1 observations indicating whether the grid cell $d \in \{1, \dots, D\}$ (with centroid s_d) at month $m \in \{1, \dots, M\}$ was masked for the CNT and BA responses, respectively. Our hierarchical model for the masking processes is

$$R_{d,m}^{\text{CNT}} \mid \mu_{dm}^{\text{CNT}} \sim \text{Bernoulli}\{\text{expit}(\mu_{dm}^{\text{CNT}})\},$$

$$R_{d,m}^{\text{BA}} \mid \mu_{dm}^{\text{BA}} \sim \text{Bernoulli}\{\text{expit}(\mu_{dm}^{\text{BA}})\};$$

$$\mu_{dm}^{\text{CNT}} = \beta_0^{\text{CNT}} + g_m(s_d) + \epsilon_m^{\text{CNT}},$$

$$\mu_{dm}^{\text{BA}} = \beta_0^{\text{BA}} + \beta g_m(s_d) + \epsilon_m^{\text{BA}};$$

$$g_m(\bullet) \sim \mathcal{GP}(\zeta),$$

$$\epsilon_m^{\text{CNT}}, \epsilon_m^{\text{BA}} \sim \mathcal{N}(0, \phi),$$

$$\beta \sim \mathcal{N}(0, \omega);$$

$$\zeta = \{\beta_0, \beta, \zeta, \phi, \omega\} \sim \text{Priors},$$

where $\text{expit}(x) = \{1 + \exp(-x)\}^{-1}$ is the inverse logit function.

We fit this model using the integrated nested Laplace approximation (INLA, Rue et al. 2009), which is an approximate Bayesian inference technique well-suited for latent Gaussian models. The parameter β governs the degree of latent sharing between the two masking processes and we use a flat and independent zero-centered Gaussian hyperprior for it. Similar frameworks were used by Koh et al. (2023) for the joint modelling of different wildfire risk components and by Diggle et al. (2010) and Pati et al. (2011) to model preferential sampling. The spatial process g_t is independently replicated in time and each replicate has a Gaussian process prior \mathcal{GP} with a Matérn covariance structure governed by the parameter vector ζ . We represent these Gaussian processes via a numerically convenient Gauss–Markov random field approximation, constructed by solving a stochastic partial differential equation (Lindgren et al. 2011). Supplement Sect. 6 details the full procedure.

We generate samples from this Bayesian model by first sampling parameters from the posterior distribution, and then generating observations from the Bernoulli model with the sampled parameters. We do this for all months, including in those where observations were masked; if a location was already part of the test set, i.e., if it was already masked, then we removed it from the validation set. Five samples were generated to obtain five folds for our cross-validation scheme. Figure 4 shows two samples from this model for March 1993. The degree of spatial and inter-variable dependencies resemble those of the masking processes in Fig. 3, and the numbers of grid cells masked and chosen for validation in each month are also similar. The triplet (2.5%

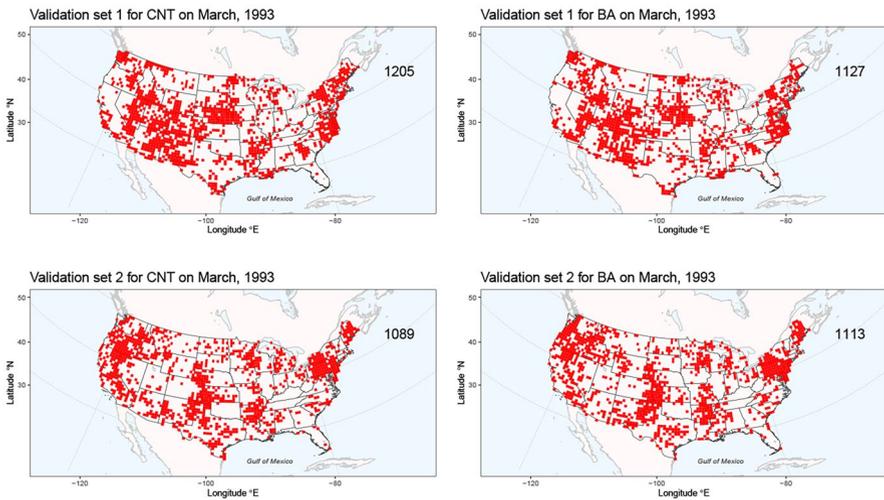


Fig. 4 The first (top) and second (bottom) validation folds (in red) for CNT (left) and BA (right) for March, 1993, from our spatiotemporal cross-validation scheme, generated from the Bayesian spatial model. The number at the top right indicates the sum of grid cells chosen

posterior quantile, posterior mean, 97.5% posterior quantile) for the scaling parameter β is (0.28, 0.42, 0.58).

4 Models

4.1 Fitting procedure

We fit our gradient tree boosting models with the approach outlined in Sect. 3.2, minimizing the loss functions described in Sect. 3.3.

We use the Poisson and dGPD loss functions to fit models on the full CNT distribution. We experimented with different high thresholds u but achieved the best prediction when we modelled the full distribution with the dGPD, i.e. $u = 0$. For all observations with a masked count response but an unmasked zero valued burnt area, we artificially set their corresponding counts to zero, and use these observations to fit the CNT models.

For wildfire sizes, we first fit models assuming that the log-transformed burned areas are conditionally normal given covariates, and call the best fitted model from this category ‘Log-Normal’; in practice, we fit these models on $\log(1+BA_i)$, $i = 1, \dots, n$, using the squared loss to predict the conditional mean of the transformed burned areas. We also consider mixture models which require split modelling of the distribution. For these, we first choose a sufficiently high (95% empirical quantile) threshold at 200ac. We then use the fire sizes above the threshold to fit a model with the GPD loss, and the log-transformed positive sizes below the threshold with the truncated gamma loss. Lastly, we fit a multi-class classifier to the wildfire size component indicators

$y_i = (y_{i,1}, \dots, y_{i,3})$, defined in Sect. 3.3.3, using the cross-entropy loss from (12); here, $y_{i,1} = 1$ if we observe no fire, $y_{i,2} = 1$ if BA_i is a medium fire (between 0 to 200ac), and $y_{i,3} = 1$ if it is a large fire ($> 200ac$). For the classifier, we also used all the observations with a masked burned area but an unmasked zero valued count, i.e., we artificially set the corresponding burned area to be zero for these observations.

Given the covariates at the i -th observation x_i , we combine the three model components to get the prediction of the cumulative conditional probability for each observation

$$\begin{aligned} \hat{\Pr}(BA_i \leq b \mid x_i) &= \hat{\Pr}(y_{i,1} = 1 \mid x_i) + \hat{\Pr}(y_{i,2} = 1 \mid x_i) \hat{\Pr}(BA_i \leq b \mid x_i, y_{i,2} = 1) \\ &\quad + \hat{\Pr}(y_{i,3} = 1 \mid x_i) \hat{\Pr}(BA_i \leq b \mid x_i, y_{i,3} = 1), \quad b \geq 0. \end{aligned}$$

We also engineered new covariates to improve our predictions. To incorporate more spatial information from our covariates (other than the longitude and latitude coordinates), we took the average value of the covariate across neighbouring grid cells for each month; this smooths the climatic variables across space. We also allow land-use covariates at neighbouring grid cells to help predictions at each grid cell.

The relationship between CNT and BA is positive, though a high CNT does not imply a high BA. Nevertheless, it is still natural to consider the other response as a covariate when modelling a given response, or at least to use this information whenever possible. As the test grid cells for BA and CNT are correlated, there are instances where the BA response was masked but the CNT wasn't, and vice versa; for 39% of masked BA observations, their corresponding CNT observations were unmasked. Using the CNT/BA covariate to predict BA/CNT thus raises the question of how best to impute its value if it was masked for a given observation. The default way to handle a missing covariate is to impute its mean across all observations, such as in algorithms from `xgboost` or `gbm`, though this will be sub-optimal for predictions on a spatially heterogeneous dataset. Instead, we use an imputation method which fits a model for the covariate and then imputes the best estimate from this model.

When modelling BA, we first fit a gradient boosting model with the dGPD loss on the CNT response, without using BA as a covariate so as to prevent data leakage, and then use the estimated parameters to find the estimated mean CNT for each observation using (10); we then impute this estimate whenever CNT was masked. When modelling CNT, we first fit a gradient boosting model, without using CNT as a covariate, with the cross-entropy loss on the wildfire size component indicators y_i , $i = 1, \dots, n$, and then impute the estimates of the probabilities $\Pr(y_{i,1} = 1)$, $\Pr(y_{i,2} = 1)$ and $\Pr(y_{i,3} = 1)$ from the fitted model, whenever BA was masked; when BA wasn't masked, we use the observed indicators y_i as a covariate.

To assess models using these covariates, a cross-validation scheme should also reflect this imputation procedure; thus, it becomes even more important to appropriately model the inter-variable dependence between CNT and BA of the masking processes, i.e., the parameter β in our spatiotemporal cross-validation scheme described in Sect. 3.4.

Our models have hyperparameters that must be tuned by cross-validation. They include the regularization parameters λ and η in (5), the proportion s of covariates subsampled at each iteration, the maximum number of leaves for each tree

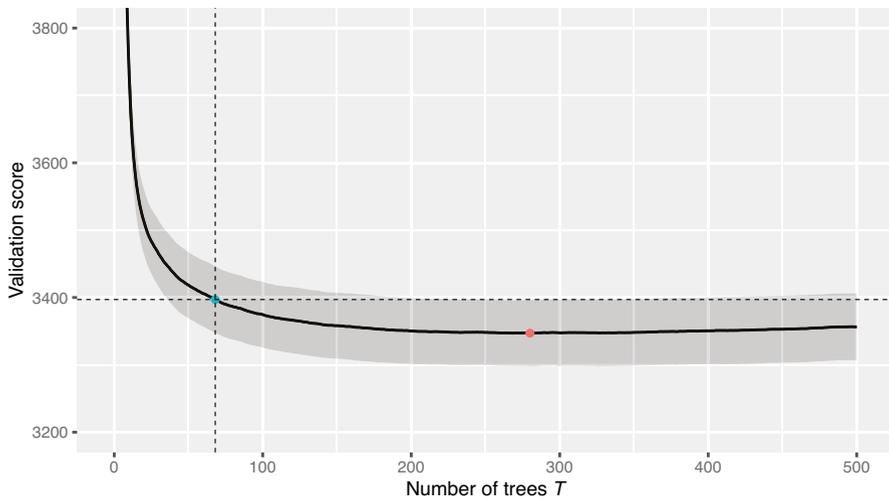


Fig. 5 The average rescaled evaluation score across all five folds of our spatiotemporal cross-validation scheme for a mixture model, as a function of the number of trees T . The shaded region shows the point-wise one-standard-error bound. The red point shows the minimum average validation score and the blue point shows the T chosen by the one-standard-error rule

L , and the number of trees T (see Sect. 3.2). Other hyperparameters from the loss functions include k , ξ , and α , which govern the shape and tails of the fitted conditional distribution, and weights w_i ($i = 1, \dots, n$) which govern the importance of each observation in the cross entropy loss. Some of our models assume common shape parameters ξ and α governing the tails of wildfire sizes and counts across the whole sample space; our preliminary analysis suggests that this assumption is reasonable in space, as the frequentist estimates of the shape parameters from pooled data are relatively homogeneous across the wildfire coordination regions.

We use the Bayesian optimization procedure outlined in Snoek et al. (2012) to choose the parameters (excluding the number of trees T) minimizing the average evaluation metric, calculated on the five cross-validation folds generated in Sect. 3.4. This procedure treats the objective function as random and first places a Gaussian process prior on it. After gathering function evaluations, the prior is updated to form the posterior distribution over the objective function, which is then used to construct an infill sampling criterion. For the mixture model, we implement separate Bayesian optimization procedures for each of the three model components.

Given the other parameters, we choose T with the one-standard-error rule (Hastie et al. 2009); i.e., we select the largest T within one standard error of the parameter that achieves the minimum in terms of the evaluation metric. Figure 5 shows the evaluation metric on the validation folds as a function of T for the wildfire size classifier in a mixture model.

We fitted our models by combining our own R routines with the `xgboost` package. We implemented the described Bayesian optimisation procedure with the `rBayesian-Optimization` package.

Table 1 The averaged rescaled evaluation score for all models, according to the data-driven cross-validation (CV) scheme outlined in Sect 3.4, the 5-fold CV scheme with random partitioning, and the true score. The bold figures highlight the best model chosen by our cross-validation approach. The asterisks indicate models not using the CNT/BA covariate

	Model	Data-driven CV	5-fold CV	Truth
Counts	Benchmark*	5172	5235	5565
	Poisson*	3413	3213	3302
	Poisson	3283	3102	3131
	dGPD*	3304	3092	3194
	dGPD	3215	2896	3068
Sizes	Benchmark*	3923	3834	4244
	Log-Normal*	3553	3450	3771
	Log-Normal	3473	3202	3501
	Mixture*	3480	3403	3582
	Mixture	3364	3089	3446

4.2 Results

The benchmark models are described in Opitz (2021). That for CNT corresponds to a generalized linear model (GLM, Nelder and Wedderburn 1972) with a Poisson response distribution and log link linear predictor using all the original covariates. The benchmark for BA first fits a generalized linear model with Gaussian response and log-link using all of the original covariates. The probability predictions, $\Pr(\text{BA}_i \leq u_{\text{BA}})$, are obtained by combining the log-Gaussian BA model with the estimated probability that $\text{CNT}_i = 0$, obtained from the benchmark Poisson model for CNT.

We relied on our cross-validation scheme devised in Sect. 3.4, using the evaluation metrics used from the competition outlined in Sect. 2, to choose which model predictions to use for the competition. After the competition, we had access to the truth and could calculate how the predictions of every model would have performed on the test set.

Table 1 shows that incorporating the engineered CNT and BA covariates improves the scores of all models by up to 7%. According to our cross-validation scheme, the best model for wildfire sizes is the mixture model, and for the counts it is the dGPD model. The best mixture model and dGPD models from the Bayesian optimization procedure have $\alpha = 52$ and $\xi = 0.8$. This implies a fat tail for the size distribution, but a thinner (Gumbel-like) tail for the wildfire count distribution, though the parameter α in the dGPD loss provides additional flexibility to the model that gives slightly better predictions than the Poisson model. All our gradient boosting models outperform the benchmark by around 10–50%.

Our cross-validation scheme tends to perform better than the 5-fold cross-validation scheme as a proxy for the true test set performance; the scores from the 5-fold cross-validation scheme are generally too optimistic compared to the true test error, especially for the wildfire size models. This optimism is especially pronounced when evaluating models that use the engineered CNT and BA covariates. Our scheme is better able to capture the inter-variable dependence between the CNT and BA masking processes,

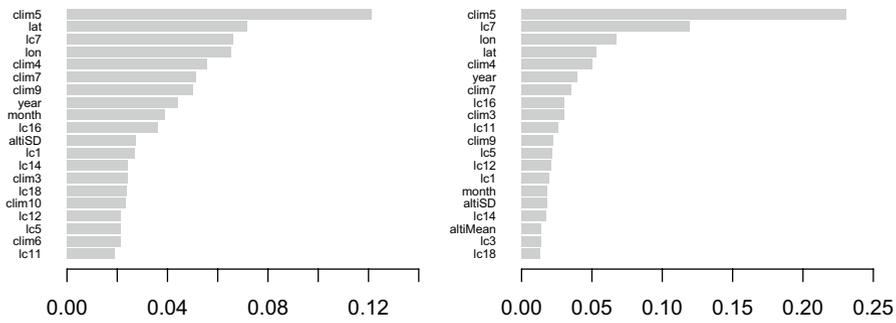


Fig. 6 The coverage (left) and gain (right) metrics for the top 20 covariates in the wildfire size classifier submodel of the best mixture model for BA (without using the engineered covariates). More information about the covariates is given in [Opitz \(2021\)](#)

giving a better reflection of how the models that incorporate the engineered covariates would perform when predicting responses on the test set.

Figure 6 shows that the gain and coverage metrics introduced in Sect. 3.2 give similar orderings for the importance of covariates when predicting the probability of being in a given wildfire size component with the best mixture model. As hypothesized in Sect. 2, clim5, lc7 and the spatial covariates of longitude (lon) and latitude (lat) are among the five most important variables for both metrics. The other variables (e.g., clim4, clim7, clim9, year, lc16, etc.) are relevant, but each is less than half as important as clim5, the most important covariate.

To evaluate the marginal effect of clim5 on the CNT response in the best dGPD model, we transformed the partial dependence estimate (7) by (10) to get the predicted mean count \hat{m}_i , and evaluated and plotted it with clim5 in the set of interest S in (7). As it is computationally infeasible to evaluate all data points x_{iC} in our setting with over 500,000 observations, we subsampled 10,000 observations to obtain our Monte Carlo estimates.

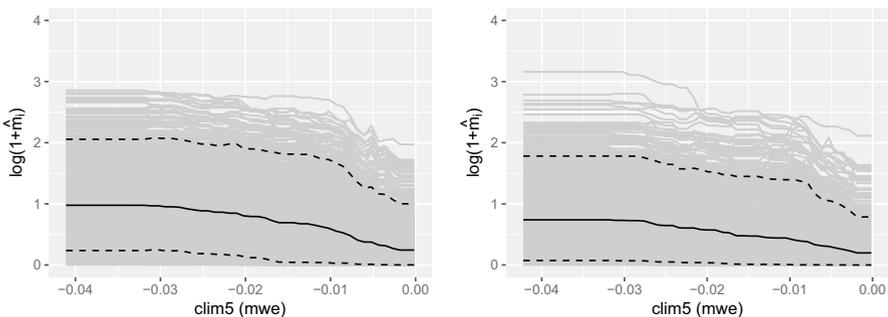


Fig. 7 The partial dependence plot with clim5 in the set of interest S , transformed so the y-axis shows the predicted mean CNT from (10), for the observations in the Rocky Mountain Area (left) and Great Basin (right) regions. The empirical 95% and 5% pointwise quantiles from the Monte Carlo estimates are shown by the dotted lines

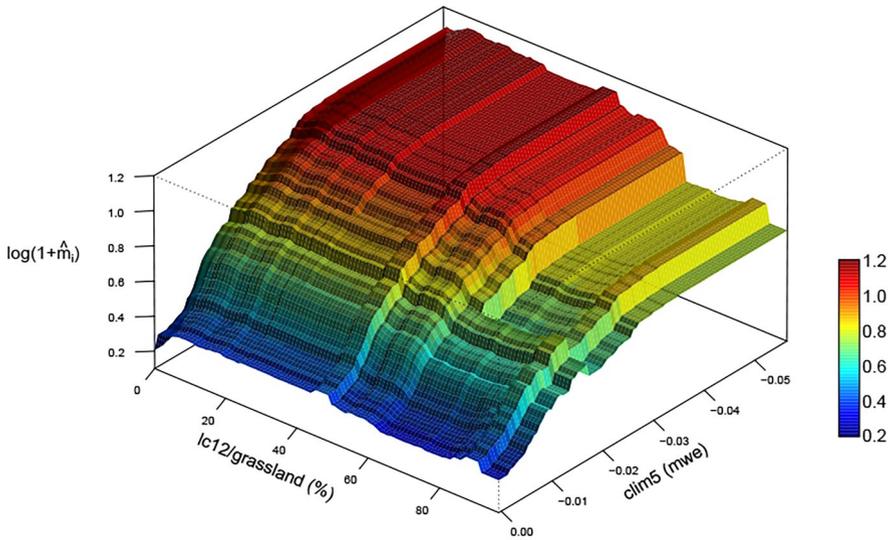


Fig. 8 The three-dimensional partial dependence function with clim5 and lc12 in the set of interest S , transformed as in Fig. 7 and plotted for all observations

Figure 7 shows that the marginal effect of clim5 on \hat{m}_i tends to be negative, especially above -0.03mwe . Figure 8 displays the joint marginal effects of clim5 and land cover covariate 12 (lc12; grassland, in %) on the predicted mean count, i.e., with clim5 and lc12 in the set of interest S in (7). The figure hints at interaction between the two covariates; increasing lc12 tends to decrease the response CNT slightly if clim5 is low, but not when clim5 is high. Although the partial dependence plot is useful for showing the overall marginal trend of a covariate on the response across all observations considered, it is important to be honest about the uncertainty associated with the Monte Carlo estimate in (7). The estimates in Figs. 7 and 8 are associated with high uncertainty throughout (not shown for the latter); our dataset is very heterogeneous and it is not possible to quantify the marginal effects of covariates with less uncertainty.

Our chosen models perform competitively when compared to the other teams' submission in the data challenge (Opitz 2021), placing second out of 28 teams in the final ranking; the other top three teams used other popular prediction techniques such as random forests, hierarchical Bayesian modelling and ANN models with adapted loss functions (Opitz 2022).

5 Discussion

We have implemented novel gradient boosting models for wildfire activity that are trained with loss functions motivated by extreme-value theory. Compared to models trained on the Poisson loss, our chosen model for wildfire counts has an additional

parameter α that governs the tail of the count distribution, which, after tuning by cross-validation, enables the model to give better predictions. Our chosen model for burned areas has specific components for extreme fire sizes. According to the given score criteria which put more weight on large fire sizes, this model improves on the models that do not discriminate between extreme and non-extreme fire sizes.

As the use of other data sources, e.g., covariates not included in the provided dataset, was strictly prohibited for the competition, we were not able to leverage other spatial information, such as the Geographic Area Coordination Center of a grid cell. Each center follows its own governing jurisdictions that could affect its fire mitigation and suppression strategies. Including this information, either by incorporating additional covariates, or by having a separate gradient boosting model for each coordination center, would improve predictions.

Our mixture model has the threshold fixed at $u = 200ac$. However, u could have also been allowed to be an additional parameter chosen by cross-validation, or as a spatio-temporally varying threshold $u_{d,m}$, $d \in \{1, \dots, D\}$, $m \in \{1, \dots, M\}$, that could be either independent (Opitz et al. 2018) or dependent (Velthoen et al. 2021) on covariates. However, implementing these approaches would significantly raise the computational cost here, especially if the model for the threshold also has hyperparameters that need to be tuned, as the optimization of all of our model components would need to be done jointly because they all rely on the threshold. Future work could investigate whether these potential approaches that increase complexity do indeed improve wildfire prediction.

Our best chosen dGPD model from the Bayesian optimization procedure has $\alpha = 52$, which implies a thin tail for the CNT distribution that is not too different from the Poisson distribution. This explains the slight improvement of the model using the dGPD compared to the Poisson loss. Had the tail of the CNT distribution been fatter, as in the case of another data application (e.g., insurance claim counts), we would have also noticed a larger improvement in predictions by the dGPD model.

We have implemented a spatial cross-validation scheme for our context which partly fixes the optimism when using traditional k -fold cross-validation to evaluate complex models with engineered covariates over a spatially heterogeneous but dependent dataset. One should always have the real-world prediction scenario in mind when choosing a cross-validation scheme, and we appealed to tools from spatial statistics to aid the validation of model predictions on our test data.

Apart from cross-validation approaches, model comparison using other predictive scores, e.g., Continuous Ranked Probability Scores (CRPS, Matheson and Winkler 1976) or tail-weighted CRPS (Gneiting and Ranjan 2011), could be used to compare simulated predictive distributions of burned areas or counts. These scores, along with graphical summaries for validating models, are an important future topic of research, and have already been explored in the extreme wildfire prediction context by Joseph et al. (2019), Pimont et al. (2020) and Koh et al. (2023). Due to the time constraints of the competition however, this is out of the scope of this paper.

Our gradient boosting models have hyperparameters from the loss functions that govern the tails of the predictive distributions: ξ , and α , which, once chosen by cross-validation, are fixed in the model. To allow more flexible modelling of the distributional tails, one could incorporate them as an additional boosting estimate in Sect. 3.2

which would allow these parameters to depend on the covariates. The boosting estimate $\hat{\theta}_i$, gradient g_i and hessian h_i in (4) would then be two-dimensional vectors. The described approach is similar to the recent work by Velthoen et al. (2021). Another avenue for future work is to assess how the the tail indices affect the cross-validation score when the other hyperparameters are fixed.

Apart from better predictions, our models improve decision support in wildfire management. The partial dependence plots in Figs. 7 and 8 allow marginal and interaction effects of covariates to be assessed, though one should be aware of the large uncertainty associated with these estimates. Importance metrics like the gain and coverage in Fig. 6 could be used for covariate selection and could prompt national wildfire predictive services to rethink designs of fire danger warning systems (e.g., indices) across the contiguous US.

6 Supplement

6.1 Terms in the gradients and Hessians

6.1.1 For the truncated gamma loss

The lower incomplete gamma function is $\gamma(k, s) = \int_0^s t^{k-1} \exp(-t) dt$, $s > 0$. Set $s = ku / \exp(\hat{\theta}_i)$. By applying the fundamental theorem of calculus and the chain rule, taking derivatives with respect to $\hat{\theta}_i$ gives

$$\gamma' \{k, ku / \exp(\hat{\theta}_i)\} = \exp(\hat{\theta}_i) \left\{ - \left(\frac{ku}{\exp(\hat{\theta}_i)} \right)^{k-1} \exp \left(- \frac{ku}{\exp(\hat{\theta}_i)} \right) \frac{ku}{\exp(\hat{\theta}_i)^2} \right\},$$

$u > 0, \hat{\theta}_i \in \mathbb{R}$,

and

$$\begin{aligned} \gamma'' \{k, ku / \exp(\hat{\theta}_i)\} = & \left[\left\{ \left(\frac{ku}{\exp(\hat{\theta}_i)^2} \right)^2 (k-1) \left(\frac{ku}{\exp(\hat{\theta}_i)} \right)^{k-2} \right. \right. \\ & + 2 \left(\frac{ku}{\exp(\hat{\theta}_i)} \right)^{k-1} \frac{ku}{\exp(\hat{\theta}_i)^3} \left. \right\} \exp \left(- \frac{ku}{\exp(\hat{\theta}_i)} \right) \\ & - \left(\frac{ku}{\exp(\hat{\theta}_i)} \right)^{k-1} \left(\frac{ku}{\exp(\hat{\theta}_i)^2} \right)^2 \exp \left(\frac{-ku}{\exp(\hat{\theta}_i)} \right) \left. \right] \exp(\hat{\theta}_i)^2 \\ & + \gamma' \{k, ku / \exp(\hat{\theta}_i)\}. \end{aligned}$$

6.1.2 For the GPD loss

Let $\kappa \in (0, 1)$. We first reparameterize the GPD probability density function as

$$f\{y_i, \exp(\hat{\theta}_i), \xi\} = \frac{\{(1 - \kappa)^{-\xi} - 1\}}{\xi \exp(\hat{\theta}_i)} \left\{ 1 + \frac{y_i \{(1 - \kappa)^{-\xi} - 1\}}{\exp(\hat{\theta}_i)} \right\}^{-(\xi+1)/\xi}, \quad (13)$$

$\xi > 0, \hat{\theta}_i \in \mathbb{R}.$

The functions $f'\{y_i, \exp(\hat{\theta}_i), \xi\}$ and $f''\{y_i, \exp(\hat{\theta}_i), \xi\}$ are obtained by differentiating (13) with respect to $\hat{\theta}_i$. Write $A = \{(1 - \kappa)^{-\xi} - 1\} / \exp(\hat{\theta}_i)$, and notice that $\partial A / \partial \hat{\theta}_i = A' = -A$. Then

$$f'\{y_i, \exp(\hat{\theta}_i), \xi\} = \frac{A^2 y_i (\xi + 1)}{\xi^2} (1 + A y_i)^{-(2\xi+1)/\xi} - f\{y_i, \exp(\hat{\theta}_i), \xi\},$$

and the expression for $f''\{y_i, \exp(\hat{\theta}_i), \xi\}$ can be easily obtained in a similarly straightforward manner.

6.2 Priors and SPDE triangulation

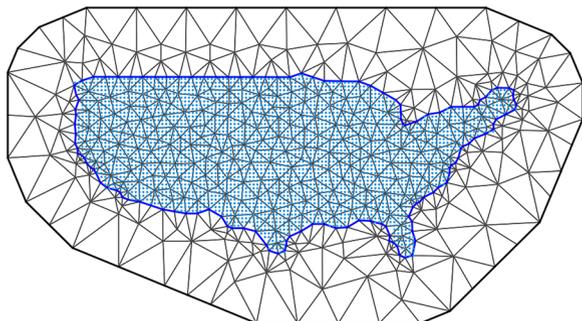
This section details the prior and SPDE triangulation specifications of the spatiotemporal Bernoulli model for the masking processes.

The fixed effect coefficient β_0^{CNT} and β_0^{BA} in our model was assigned a flat Gaussian prior with zero mean and precision 0.001. The prior for the scaling parameter β is a zero-centered Gaussian distribution with precision $\omega = 1/20$. Lastly, we assigned a log-gamma hyperprior with mean unity and precision 0.0005 to the variance hyperparameter ϕ .

The spatial Gaussian random effects g_m and their conditional distributions must be tractable in our setting with many pixels. We use the Matérn covariance function for random effects, given as follows for two points s_1 and s_2 :

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

Fig. 9 Triangulation mesh of the spatial study region (blue contours) for the SPDE approach. Neumann boundary conditions are set on the exterior (black) boundary to obtain a unique solution. The blue points represent the centroids of the grid cells. The finite element solution defines a Gauss–Markov random vector with one variable in each node



with Euclidean distance $\|\cdot\|$, gamma function Γ , the 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 driven by Gaussian noise (SPDE, Lindgren et al. 2011), where we fix the smoothness ν at unity.

The parameter vector ζ consists of r and σ parameters which were assigned penalized complexity priors (Fuglstad et al. 2018). Such priors penalize the distance of the prior of a model component towards a simpler baseline at a constant rate (Simpson et al. 2017).

The discretization points when triangulating the spatial domain are chosen as the nodes of a finite element representation which enables efficient inference for random effects representing spatial variation. Our spatial triangulation mesh in Fig. 9 has 508 nodes. It is sparser in the extended zone around the study area to ensure that the SPDE boundary conditions have negligible influence on the study area.

For computational reasons, we only use observations of the masking processes from the first ten months, and subsample 30% of the data when fitting our model.

Acknowledgements The author is grateful to the two anonymous referees and the Guest Editor for their constructive comments that improved the quality of this paper. Further thanks go to Anthony C. Davison for helpful comments and discussions, Thomas Opitz for organising the EVA 2021 data challenge, and the Swiss National Science Foundation for financial support.

Funding Open access funding provided by University of Bern.

Data availability statement The data that support the findings of this study are from the EVA 2021 data challenge (Opitz 2022), and are available on request from the the organiser.

Declarations

Conflicts of interest The author declares that he has no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Breiman, L.: Random forests. *Mach. Learn.* **45**(1), 5–32 (2001)
- Breiman, L., Friedman, J.H., Olshen, R.A., Stone, C.J.: *Classification and Regression Trees*. Wadsworth and Brooks, Monterey, CA (1984)
- Brillinger, D.R., Preisler, H.K., Benoit, J.W.: Probabilistic risk assessment for wildfires. *Environmetrics* **17**(6), 623–633 (2006)

- Bühlmann, P., Hothorn, T.: Boosting algorithms: regularization, prediction and model fitting. *Stat. Sci.* **22**(4), 477–505 (2007)
- Chen, T., Guestrin, C.: XGBoost: a scalable tree boosting system. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16, pp. 785–794. ACM, New York, NY, USA (2016)
- Cox, D.R.: The regression analysis of binary sequences (with discussion). *J. Roy. Stat. Soc.: Ser. B (Methodol.)* **20**(2), 215–232 (1958)
- Cui, W., Perera, A.H.: What do we know about forest fire size distribution, and why is this knowledge useful for forest management? *Int. J. Wildland Fire* **17**(2), 234–244 (2008)
- Cumming, S.: A parametric model of the fire-size distribution. *Can. J. For. Res.* **31**(8), 1297–1303 (2001)
- Davison, A.C., Smith, R.L.: Models for exceedances over high thresholds (with discussion). *J. Roy. Stat. Soc.: Ser. B (Methodol.)* **52**(3), 393–442 (1990)
- De Angelis, A., Ricotta, C., Conedera, M., Pezzatti, G.B.: Modelling the meteorological forest fire niche in heterogeneous pyrologic conditions. *PLoS ONE* **10**(2), 1–17 (2015)
- De Zea Bermudez, P., Mendes, J., Pereira, J.M., Turkman, K.F., Vasconcelos, M.J.: Spatial and temporal extremes of wildfire sizes in Portugal (1984–2004). *Int. J. Wildland Fire* **18**(8), 983–991 (2009)
- Diggle, P.J., Menezes, R., Su, T.-L.: Geostatistical inference under preferential sampling (with discussion). *J. Roy. Stat. Soc.: Ser. C (Appl. Stat.)* **59**(2), 191–232 (2010)
- Dutta, R., Aryal, J., Das, A., Kirkpatrick, J.B.: Deep cognitive imaging systems enable estimation of continental-scale fire incidence from climate data. *Sci. Rep.* **3**(1), 3188 (2013)
- Friedman, J.H.: Greedy function approximation: a gradient boosting machine. *Ann. Stat.* **29**, 1189–1232 (2001)
- Friedman, J., Hastie, T., Tibshirani, R.: Additive logistic regression: a statistical view of boosting (with discussion and a rejoinder by the authors). *Ann. Stat.* **28**(2), 337–407 (2000)
- Fuglstad, G.-A., Simpson, D., Lindgren, F., Rue, H.: Constructing priors that penalize the complexity of Gaussian random fields. *J. Am. Stat. Assoc.* **114**(525), 445–452 (2018)
- Genton, M.G., Butry, D.T., Gumpertz, M.L., Prestemon, J.P.: Spatio-temporal analysis of wildfire ignitions in the St Johns River water management district, Florida. *Int. J. Wildland Fire* **15**(1), 87–97 (2006)
- Gneiting, T., Ranjan, R.: Comparing density forecasts using threshold- and quantile-weighted scoring rules. *J. Bus. Econ. Stat.* **29**(3), 411–422 (2011)
- Greenwell, B., Boehmke, B., Cunningham, J., Developers, G.: GBM: generalized boosted regression models. *R Package Version* **2**(1), 8 (2020)
- Hastie, T., Tibshirani, R., Friedman, J.: *The Elements of Statistical Learning: Data Mining, Inference and Prediction*, 2nd edn. Springer (2009)
- Hitz, A., Davis, R., Samorodnitsky, G.: 'Discrete extremes'. Preprint <https://arxiv.org/abs/1707.05033> (2017)
- Jain, P., Coogan, S.C., Subramanian, S.G., Crowley, M., Taylor, S., Flannigan, M.D.: A review of machine learning applications in wildfire science and management. *Environ. Rev.* **28**(4), 478–505 (2020)
- 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., Balch, J.K.: Spatiotemporal prediction of wildfire size extremes with Bayesian finite sample maxima. *Ecol. Appl.* **29**(6), e01898 (2019)
- Koh, J., Pimont, F., Dupuy, J.-L., Opitz, T.: Spatiotemporal wildfire modeling through point processes with moderate and extreme marks. *Ann Appl Stat.* **17**(1), 560–582 (2023)
- Liang, H., Zhang, M., Wang, H.: A neural network model for wildfire scale prediction using meteorological factors. *IEEE Access* **7**, 176746–176755 (2019)
- Lindgren, F., Rue, H., Lindström, J.: An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach (with discussion). *J. R. Stat. Soc. Series B. Stat. Methodol.* **73**(4), 423–498 (2011)
- Matheson, J.E., Winkler, R.L.: Scoring rules for continuous probability distributions. *Manage. Sci.* **22**(10), 1087–1096 (1976)
- Mitsopoulos, I., Mallinis, G.: A data-driven approach to assess large fire size generation in Greece. *Nat. Hazards* **88**(3), 1591–1607 (2017)
- National Interagency Fire Center: 'Total wildfires and acres' (2021). Data retrieved from <https://www.predictiveservices.nifc.gov/intelligence/intelligence.htm>. Accessed 17 Jun 2021
- Nelder, J.A., Wedderburn, R.W.M.: 'Generalized linear models'. *J. R. Stat. Soc. Ser. A Stat. Soc.* **135**(3), 370–384 (1972)
- Opitz, T.: 'EVA 2021 data challenge' (2021). <https://www.maths.ed.ac.uk/school-of-mathematics/eva-2021/competitions/data-challenge>. Accessed 17 Jun 2021

- Opitz, T.: 'Editorial: EVA 2021 data competition on spatio-temporal prediction of wildfire activity in the United States'. *Extremes to appear* (2022)
- Opitz, T., Bonneau, F., Gabriel, E.: Point-process based modeling of space-time structures of forest fire occurrences in Mediterranean France. *Spat. Stat.* **40**, 100429 (2020)
- Opitz, T., Huser, R., Bakka, H., Rue, H.: INLA goes extreme: Bayesian tail regression for the estimation of high spatio-temporal quantiles. *Extremes* **21**(3), 441–462 (2018)
- Pati, D., Reich, B.J., Dunson, D.B.: Bayesian geostatistical modelling with informative sampling locations. *Biometrika* **98**(1), 35–48 (2011)
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, Édouard.: Scikit-learn: machine learning in Python. *J. Mach. Learn. Res.* **12**(85), 2825–2830 (2011)
- Peng, R.D., Schoenberg, F.P., Woods, J.A.: A space-time conditional intensity model for evaluating a wildfire hazard index. *J. Am. Stat. Assoc.* **100**(469), 26–35 (2005)
- Pereira, J.M.C., Turkman, K.F.: Statistical models of vegetation fires: spatial and temporal patterns. In: *Handbook of Environmental and Ecological Statistics*, pp. 401–420. Chapman and Hall/CRC (2019)
- Pimont, F., Fargeon, H., Opitz, T., Ruffault, J., Barbero, R., Martin-StPaul, N., Rigolot, E. I., Rivière, M., Dupuy, J.-L.: Prediction of regional wildfire activity in the probabilistic Bayesian framework of Firelihood. *Ecol. Appl.* e02316 (2021)
- Pohjankukka, J., Pahikkala, T., Nevalainen, P., Heikkonen, J.: Estimating the prediction performance of spatial models via spatial k-fold cross validation. *Int. J. Geogr. Inf. Sci.* **31**(10), 2001–2019 (2017)
- Preisler, H.K., Brillinger, D.R., Burgan, R.E., Benoit, J.: Probability based models for estimation of wildfire risk. *Int. J. Wildland Fire* **13**(2), 133–142 (2004)
- Prieto, F., Gómez-Déniz, E., Sarabia, J.M.: Modelling road accident blackspots data with the discrete generalized Pareto distribution. *Accid. Anal. Prev.* **71**, 38–49 (2014)
- Rasmussen, C.E., Williams, C.K.I.: *Gaussian Processes for Machine Learning*. The MIT Press (2005)
- Roberts, D.R., Bahn, V., Ciuti, S., Boyce, M.S., Elith, J., Guillera-Aroita, G., Hauenstein, S., Lahoz-Monfort, J.J., Schröder, B., Thuiller, W., Warton, D.I., Wintle, B.A., Hartig, F., Dormann, C.F.: Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure. *Ecography* **40**(8), 913–929 (2017)
- Rue, H., Martino, S., Chopin, N.: Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations (with discussion). *J. R. Stat. Soc. Series B Stat. Methodol.* **71**(2), 319–392 (2009)
- Sakr, G.E., Elhadj, I.H., Mitri, G., Wejinya, U.C.: Artificial intelligence for forest fire prediction. In: 2010 IEEE/ASME International Conference on Advanced Intelligent Mechatronics, pp. 1311–1316. (2010)
- Shidik, G.F., Mustofa, K.: Predicting size of forest fire using hybrid model. In: Linawati, Mahendra, M.S., Neuhold, E.J., Tjoa, A.M., You, I. (eds.) *Information and Communication Technology*, pp. 316–327. Springer Berlin Heidelberg, Berlin, Heidelberg (2014)
- Shimura, T.: Discretization of distributions in the maximum domain of attraction. *Extremes* **15**(3), 299–317 (2012)
- Simpson, D., Rue, H., Riebler, A., Martins, T.G., Sørbye, S.H., et al.: Penalising model component complexity: A principled, practical approach to constructing priors. *Stat. Sci.* **32**(1), 1–28 (2017)
- Snoek, J., Larochelle, H., Adams, R.P.: Practical Bayesian optimization of machine learning algorithms. In: Pereira, F., Burges, C.J.C., Bottou, L., Weinberger, K.Q. (eds.) *Advances in Neural Information Processing Systems*, vol. 25. Curran Associates, Inc (2012)
- Stewart, S.I., Radeloff, V.C., Hammer, R.B., Hawbaker, T.J.: Defining the Wildland-Urban Interface. *J. Forest.* **105**(4), 201–207 (2007)
- Taylor, S.W., Woolford, D.G., Dean, C.B., Martell, D.L.: Wildfire prediction to inform fire management: Statistical science challenges. *Stat. Sci.* **28**(4), 586–615 (2013)
- Tonini, M., Pereira, M.G., Parente, J., Orozco, C.V.: Evolution of forest fires in Portugal: from spatio-temporal point events to smoothed density maps. *Nat. Hazards* **85**(3), 1489–1510 (2017)
- Turkman, K.F., Amaral Turkman, M.A., Pereira, J.M.: Asymptotic models and inference for extremes of spatio-temporal data. *Extremes* **13**(4), 375–397 (2010)
- van Wagner, C.: Conditions for the start and spread of crown fire. *Can. J. For. Res.* **7**(1), 23–34 (1977)
- Velthoen, J., Dombry, C., Cai, J.-J., Engelke, S.: Gradient boosting for extreme quantile regression. Preprint at <https://arxiv.org/abs/2103.00808> (2021)
- Vilar, L., Woolford, D.G., Martell, D.L., Martín, M.P.: Spatio-temporal analysis of wildfire ignitions in the St Johns River water management district, Florida. *Int. J. Wildland Fire* **19**(3), 325–337 (2010)
- Wood, S.: *Generalized Additive Models: an Introduction with R*, 2nd edn. Chapman and Hall/CRC (2017)

- Woolford, D.G., Bellhouse, D.R., Braun, W.J., Dean, C.B., Martell, D.L., Sun, J.: A spatio-temporal model for people-caused forest fire occurrence in the Romeo Malette forest. *J. Environ. Stat.* **2**(1), 1–26 (2011)
- Xi, D.D., Taylor, S.W., Woolford, D.G., Dean, C.: Statistical models of key components of wildfire risk. *Annu. Rev. Stat. Appl.* **6**, 197–222 (2019)
- Xie, Y., Peng, M.: Forest fire forecasting using ensemble learning approaches. *Neural Comput. Appl.* **31**(9), 4541–4550 (2019)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.