









Spatiotemporal prediction of wildfire size extremes with Bayesian finite sample maxima

MAXWELL B. JOSEPH ^{1,3} MATTHEW W. ROSSI ¹ NATHAN P. MIETKIEWICZ ¹ ADAM L. MAHOOD,¹ MEGAN E. CATTAU ¹ LISE ANN ST. DENIS,¹ R. CHELSEA NAGY ¹ VIRGINIA IGLESIAS ¹ JOHN T. ABATZOGLOU ² AND JENNIFER K. BALCH ¹

¹Earth Lab, University of Colorado Boulder, 4001 Discovery Drive, Suite S348 611 UCB, Boulder, Colorado 80303 USA

²Department of Geography, University of Idaho, 875 Perimeter Drive, MS 3021, Moscow, Idaho 83844-3021 USA

Citation: M. B. Joseph, M. W. Rossi, N. P. Mietkiewicz, A. L. Mahood, M. E. Cattau, L. A. St. Denis, R. C. Nagy, V. Iglesias, J. T. Abatzoglou, and J. K. Balch. 2019. Spatiotemporal prediction of wildfire size extremes with Bayesian finite sample maxima. *Ecological Applications* 29(6):e01898. 10.1002/eap.1898

Abstract. Wildfires are becoming more frequent in parts of the globe, but predicting where and when wildfires occur remains difficult. To predict wildfire extremes across the contiguous United States, we integrate a 30-yr wildfire record with meteorological and housing data in spatiotemporal Bayesian statistical models with spatially varying nonlinear effects. We compared different distributions for the number and sizes of large fires to generate a posterior predictive distribution based on finite sample maxima for extreme events (the largest fires over bounded spatiotemporal domains). A zero-inflated negative binomial model for fire counts and a lognormal model for burned areas provided the best performance. This model attains 99% interval coverage for the number of fires and 93% coverage for fire sizes over a six year withheld data set. Dryness and air temperature strongly predict extreme wildfire probabilities. Housing density has a hump-shaped relationship with fire occurrence, with more fires occurring at intermediate housing densities. Statistically, these drivers affect the chance of an extreme wildfire in two ways: by altering fire size distributions, and by altering fire frequency, which influences sampling from the tails of fire size distributions. We conclude that recent extremes should not be surprising, and that the contiguous United States may be on the verge of even larger wildfire extremes.

Key words: Bayesian; climate; extremes; fire; spatiotemporal; wildfire.

INTRODUCTION

Wildfire frequency and burned area has increased over the past couple decades in the United States (Dennison et al. 2014, Westerling 2016), and elsewhere (Krawchuk et al. 2009, Pechony and Shindell 2010). In addition to the ecological and smoke impacts associated with increased burned area, there has been an increasing interest in extreme wildfires (Williams 2013) given their impact on human lives and infrastructure (Kochi et al. 2010, Diaz 2012). While case studies of particular extremes provide insight into what caused past events (Peterson et al. 2015, Nauslar et al. 2018), predictions of future extremes at a national level could inform disaster related resource allocation. The term “extreme” has multiple meanings with respect to wildfires (Tedim et al. 2018), and, in this paper, we consider an extreme wildfire to be a fire with the largest burned area over a bounded spatiotemporal domain, i.e., the block maximum within a spatial region and a temporal interval (Coles et al. 2001). For example,

the block maxima for wildfires across the contiguous United States can be defined on a yearly basis (Fig. 1).

Factors driving wildfire extremes vary in space and time (Barbero et al. 2014), but it is unclear how best to account for this in a predictive model. Previous efforts have used year- or recognition-specific models, aggregating over space or time (Bermudez et al. 2009), temporally or spatially explicit models (Mendes et al. 2010), and spatial models with year as a covariate (Díaz-Avalos et al. 2016). Recently, rich spatiotemporal models have been described with linear, spatially constant, covariate effects (Serra et al. 2014a, b). However, linear, spatially constant effects are suboptimal over large spatial domains with nonlinear drivers (Fosberg 1978, Goodrick 2002, Preisler et al. 2004; Preisler and Westerling 2007, Balshi et al. 2009, Krawchuk et al. 2009, Pechony and Shindell 2009, Vilar et al. 2010, Woolford et al. 2011, 2014). For example, global wildfire probability shows a hump-shaped relationship with temperature and moisture (Moritz et al. 2012). Interactions among drivers also impose nonlinearity, e.g., in hot and dry climates fires are fuel limited (McLaughlin and Bowers 1982), but in cold and wet climates fires are energy limited (Krawchuk and Moritz 2011).

Manuscript received 14 August 2018; revised 19 February 2019; accepted 19 March 2019. Corresponding Editor: David S. Schimel.

³E-mail: maxwell.b.joseph@colorado.edu




FIG. 1. Sizes of wildfires over 405 ha in the contiguous United States, from the Monitoring Trends in Burn Severity multi-agency program. Each point represents a fire event, and the largest fires for each year (the block maxima) are shown as solid black points.

Prediction is also complicated by uncertainty in which distribution(s) to use to assign probabilities to extreme events. The generalized Pareto distribution (GPD) has frequently been used (Bermudez et al. 2009, Jiang and Zhuang 2011), but the GPD requires a threshold to delineate extreme events (Davison and Smith 1990, Coles 2014). The utility and validity of a threshold for extremes in a heterogeneous region is debatable (Tedim et al. 2018). Recently proposed metastatistical extreme value (MEV) approaches do not require such a threshold, and are based on the statistical distribution of finite sample maxima, i.e., the probability distribution of the maximum value for a finite number of events (Marani and Ignaccolo 2015, Zorzetto et al. 2016). In the MEV framework, the occurrence and size of future events, and the parameters of their distributions are treated as random variables which together imply a distribution for extremes. This approach has roots in compound distributions (Dubey 1970, Wiitala 1999), doubly stochastic processes (Cox and Isham 1980), superstatistics (Beck and Cohen 2003), and the Bayesian posterior predictive distribution (Gelman et al. 2013). The link to Bayesian inference is particularly useful, as it provides an easy way to propagate uncertainty forward to to predictions of extremes (Coles et al. 2003).

Here, we extend the finite sample maximum approach to account for non-linear, spatially varying covariate effects with the goal of predicting extreme wildfire events from a statistical perspective across the contiguous United States. Specifically, we aim to predict occurrence (where and when), and magnitude (burned area) of large wildfires at a monthly time scale and regional spatial scale across the contiguous United States.

METHODS

Data description

We acquired wildfire event data for the contiguous United States from the Monitoring Trends in Burn

Severity (MTBS) program (Eidenshink et al. 2007), which includes spatiotemporal information on the occurrence of wildfires in the United States from 1984 to 2016. The MTBS data contain fires >1,000 acres (~405 ha) in the western U.S. and >500 acres (~202 ha) in the eastern U.S. For consistency across the United States, we discarded all records in the MTBS data <1,000 acres, retaining 10,736 fire events (Fig. 2A). Each event in the MTBS data has a discovery date, spatial point location, and final size.

To explain fire size and occurrence, we used a combination of meteorological variables including humidity, air temperature, precipitation, and wind speed. These variables were selected on the basis of previous work, and also with an aim to drive a predictive model with interpretable meteorological quantities. Meteorological layers were acquired from the gridMET data (Abatzoglou 2013) that blends monthly high-spatial resolution (~4-km) climate data from the Parameter-elevation Relationships on Independent Slopes Model (Daly et al. 2008) with high-temporal resolution (hourly) data from the National Land Data Assimilation System (NLDAS2) using climatologically aided interpolation. The resultant products are a suite of surface meteorological variables summarized at the daily time step and at a 4-km pixel resolution. Daily total precipitation, minimum relative humidity, mean wind speed, and maximum air temperature were averaged at a monthly time step for each of 84 Environmental Protection Agency level 3 (L3) ecoregions for each month from 1984 to 2016 (Omernik 1987, Omernik and Griffith 2014). We also computed cumulative monthly precipitation over the previous 12 months for each ecoregion-month combination. We chose to segment the United States with level 3 ecoregions as a compromise between the more numerous (computationally demanding) level 4 ecoregions, and the coarser level 2 ecoregions.

We used publicly available housing density estimates that were generated based on the U.S. 2000 decennial census as explanatory variables that may relate to

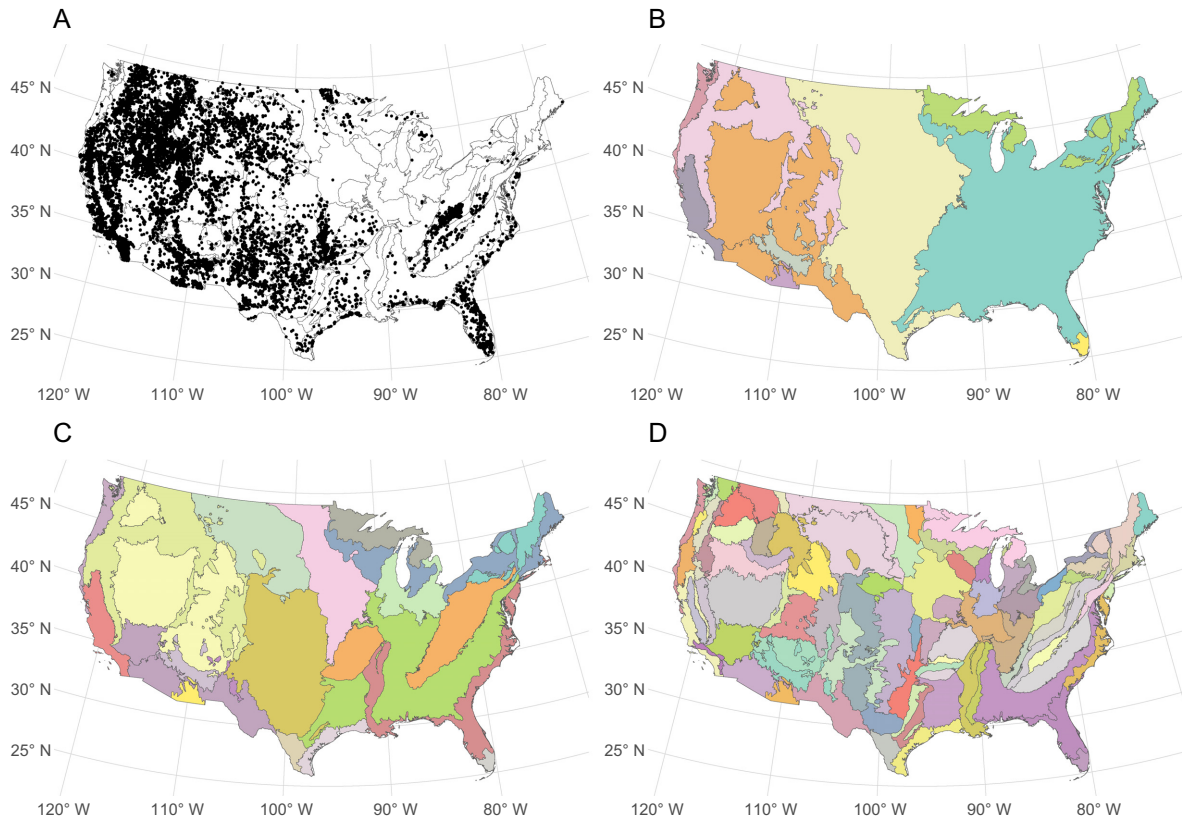


FIG. 2. (A) Large wildfire ignition locations are shown as points across the study region. Colors in panels (B), (C), and (D) show level 1, 2, and 3 ecoregions respectively.

human ignition pressure (Radeloff et al. 2010). These are provided at decadal time steps, and spatially at the level of census partial block groups. To generate approximate measures of housing density at monthly time intervals, we used a simple linear interpolation over time for each block group, then aggregated spatially across block groups to compute mean housing density for each ecoregion in each month.

Model development

We built two types of models: one describing the occurrence of fires within each L3 ecoregion over time (i.e., the total number of fires occurring in each ecoregion for each month from 1984 to 2016), and another describing the size of each wildfire in each ecoregion and month. For occurrence models, the response variable was a count (number of fires), and for burned area models, the response was a continuous positive quantity (size of each fire event). We used the period from 1984 to 2009 for training, withholding the period from 2010 to 2016 to evaluate predictive performance.

Fire occurrence.—We constructed four models for fire occurrence and compared their predictive performance based on test-set log likelihood and posterior predictive

checks for the proportion of zeros, maximum count, and total count. The models differed in the distributions used in the likelihood, representing counts as a Poisson, negative binomial, zero-inflated Poisson, or zero-inflated negative binomial random variable. The Poisson distribution is a common choice for counts, and the negative binomial distribution provides an alternative that can account for overdispersion. The zero-inflated versions of these distributions include a component to represent extra zeros, which might be expected to work well if there are independent processes that determine whether nonzero counts are possible (Lambert 1992).

For spatial units (ecoregions) $s = 1, \dots, S$ and time steps (months) $t = 1, \dots, T$, each model defines a probability mass function for $n_{s,t}$: the number of fires over 405 ha in ecoregion s and time step t . For each of the four count distributions under consideration, location parameters $\mu_{s,t}$ and (for zero-inflated models) structural zero inflation parameters $\pi_{s,t}$ were allowed to vary in space and time. We used a log link function to ensure that $\mu_{s,t} > 0$, and a logit link function to ensure that $\pi_{s,t} \in (0, 1)$. Concatenating over spatial and temporal units, so that $\boldsymbol{\mu} = (\mu_{s=1,t=1}, \mu_{s=2,t=1}, \dots, \mu_{s=S,t=1}, \mu_{s=S,t=2}, \dots, \mu_{s=S,t=T})$, and similarly for $\boldsymbol{\pi}$, we modeled location and (when applicable) zero inflation parameters as

$$\log(\boldsymbol{\mu}) = \alpha^{(\mu)} + \mathbf{X}\boldsymbol{\beta}^{(\mu)} + \boldsymbol{\phi}^{(\mu)} + \log(\mathbf{a})$$

$$\text{logit}(\boldsymbol{\pi}) = \alpha^{(\pi)} + \mathbf{X}\boldsymbol{\beta}^{(\pi)} + \boldsymbol{\phi}^{(\pi)}$$

where $\alpha^{(\mu)}$ and $\alpha^{(\pi)}$ are scalar intercept parameters, \mathbf{X} is a known $(S \times T) \times p$ design matrix, where p is the number of input features, $\boldsymbol{\beta}^{(\mu)}$ and $\boldsymbol{\beta}^{(\pi)}$ are column vector parameters of length p , $\boldsymbol{\phi}^{(\mu)}$ and $\boldsymbol{\phi}^{(\pi)}$ are column vector parameters of length $S \times T$ containing spatiotemporal adjustments, and \mathbf{a} is a known offset vector of areas for spatial unit $s = 1, 2, \dots, S$, repeated T times.

Burned area.—We developed five candidate models for fire size, each of which specified a different distribution for the size (burned area) of individual fire events (Reed and McKelvey 2002, Hernandez et al. 2015), including the generalized Pareto (Hosking and Wallis 1987), tapered Pareto (Schoenberg et al. 2003), lognormal, gamma, and Weibull distributions. We evaluated each model in terms of test set log likelihood, and posterior predictive checks for fire size extremes. We defined the response y_i as the number of hectares burned over 405 for the i th fire event, which occurred in spatial unit s_i and time step t_i .

Because each burned area distribution has a different parameterization, we included covariate effects in a distribution-specific way. For the generalized Pareto distribution (GPD), we assumed a positive shape parameter, leading to a Lomax distribution for exceedances (Bermudez et al. 2009). The GPD and Lomax shape parameters are related by $\kappa^{(\text{GPD})} = 1/\kappa^{(L)}$, and the GPD scale parameter is related to the Lomax scale and shape parameters by $\sigma^{(\text{GPD})} = \sigma^{(L)}/\kappa^{(L)}$. We introduced covariate dependence via the Lomax scale parameter using a log link. For event i , $\log(\sigma_i^{(L)}) = \alpha + \mathbf{X}_{(s_i, t_i)}\boldsymbol{\beta} + \phi_{s_i, t_i}$, where α is an intercept parameter, $\boldsymbol{\beta}$ is a length p vector of coefficients, $\mathbf{X}_{(s_i, t_i)}$ is a row vector from \mathbf{X} , and ϕ_{s_i, t_i} is a spatiotemporal adjustment for s_i and t_i . For the tapered Pareto model, we modeled the shape parameter as $\log(\kappa_i) = \alpha + \mathbf{X}_{(s_i, t_i)}\boldsymbol{\beta} + \phi_{s_i, t_i}$. The lognormal model included covariate dependence via the location parameter: $\mu_i = \alpha + \mathbf{X}_{(s_i, t_i)}\boldsymbol{\beta} + \phi_{s_i, t_i}$. The gamma model used a log link for the expected value: $\log(E(y_i)) = \alpha + \mathbf{X}_{(s_i, t_i)}\boldsymbol{\beta} + \phi_{s_i, t_i}$. Last, we modeled the

Weibull scale parameter as $\log(\sigma_i) = \alpha + \mathbf{X}_{(s_i, t_i)}\boldsymbol{\beta} + \phi_{s_i, t_i}$. More detail on the parameterization of each burned area distribution is provided in Appendix S2.

Accounting for nonlinear forcing.—The design matrix \mathbf{X} was constructed to allow for spatially varying nonlinear effects of housing density and meteorological drivers. We used B-splines to account for nonlinearity (Fig. 3) and allowed the coefficients for each basis vector to vary spatially (Wood 2017). First, we constructed univariate B-splines for log housing density, wind speed, same month precipitation, previous 12 month precipitation, air temperature, and humidity, with five degrees of freedom (including an intercept) for each variable. This step generated 30 basis vectors (five for each of six variables).

To allow for spatial variation in these nonlinear effects, we added interaction effects between each of the basis vectors and ecoregions (Brezger and Lang 2006, Kneib et al. 2009). The hierarchical nesting of ecoregion designations (Fig. 2B–D) lends itself to such interactions. Conceptually, coefficients in a level 3 ecoregion may be related to coefficients in the level 2 ecoregion containing the level 3 region, the level 1 ecoregion containing the level 2 region, and a global effect. The coefficient associated with a basis vector for any level 3 ecoregion is treated as a sum of a global effect, a level 1 ecoregion adjustment, a level 2 ecoregion adjustment, and a level 3 ecoregion adjustment. Thus, for every univariate basis vector, we included interaction effects with ecoregion at each of the three ecoregion levels. This allows borrowing of information across space (level 3 ecoregions in a level 2 ecoregion are often adjacent), and for regions that are ecologically similar. We also included adjustments on the global intercept for each level 1, 2, and 3 ecoregion to account for spatial variation that is unrelated to climate or housing density. This specification induces sparsity in \mathbf{X} that we exploit to increase the efficiency of computing $\boldsymbol{\mu}$ and $\boldsymbol{\pi}$. In total, \mathbf{X} has $p = 3,472$ columns, with 97% zero entries.

Prior specification

To avoid overfitting, we used a regularized horseshoe prior on the coefficients associated with the spatially




Fig. 3. Conceptual Figure to illustrate the use of B-splines to construct nonlinear functions. In the left panel, five B-spline vectors are shown, which map values of an input variable (on the x -axis) to a value on the y -axis. The middle panel shows the same B-spline vectors, but weighted (multiplied) by real numbers, with the weights illustrated as annotations. These weighted B-spline vectors are summed to produce the values of a nonlinear function (right panel).

varying nonlinear effects described above (Piironen et al. 2017). This prior places high probability close to zero, while retaining heavy enough tails that nonzero coefficients are not shrunk too strongly toward zero. This is consistent with our prior expectation that most of the coefficients associated with the columns in \mathbf{X} were close to zero. For the zero inflated count models, we used a multivariate horseshoe to allow information sharing between the zero inflated and distribution specific location parameters (Peltola et al. 2014). For the remaining count models and all burned area models, this was a univariate horseshoe prior. Spatiotemporal random effects were constructed using a temporally autoregressive, spatially intrinsically autoregressive formulation (Besag and Kooperberg 1995, Banerjee et al. 2014). Details of these priors and the resulting joint distributions are provided in Appendices S1 and S2, respectively.

Posterior predictive inference for finite sample maxima

We used the posterior predictive distribution to check each model and make inference on extremes. The posterior predictive distribution provides a distribution for replications of observed data (y^{rep}), and predictions of future data (Gelman et al. 2013). Conceptually, for a good model, y^{rep} should be similar to observed training data y , and future predictions should be similar to future data. Distributions over both quantities can be obtained by conditioning on y and marginalizing over model parameters θ , e.g., $[y^{\text{rep}}|y] = \int [y^{\text{rep}}|\theta][\theta|y]d\theta$.

Posterior predictive distributions facilitate model checks that compare predicted and observed test statistics (Gelman et al. 1996). To evaluate whether models captured tail behavior, we compared empirical maxima ($T(y) = \max(y)$) to the predicted distribution of maxima $T(y^{\text{rep}})$. We also include predictive checks for the proportion of zero counts, and totals for count and burned area models. Posterior predictive inference for finite sample maxima is similar in spirit to the MEV approach. Both obtain a distribution over maxima by marginalizing over unknowns including the number of events, size of each event, and parameters of their distributions (Marani and Ignaccolo 2015). However, a Bayesian approach explicitly conditions on the observed data to obtain a posterior distribution of parameters.

Seeing this connection is useful in the context of including priors and propagating uncertainty to derived parameters. For any ecoregion s and timestep t , if we define a particular maximum fire size conditional on a fire having occurred as $z_{s,t}$, and let $Z_{s,t}$ represent the random variable of maximum fire size, then the cumulative distribution function (CDF) for $z_{s,t}$ is given by $\Pr(Z_{s,t} \leq z_{s,t}) = F(y_{(s,t)})^{n_{s,t}}$, where $F(y_{(s,t)})$ is the CDF of fire size, and $n_{s,t}$ is the number of wildfire events. Thus, $\Pr(Z_{s,t} \leq z_{s,t})$ is the distribution function for the finite sample maximum. The CDF for $z_{s,t}$ can be inverted to produce a quantile function that permits computation of prediction intervals for maximum fire sizes, conditional

on fires having occurred. Given a collection of posterior draws from a burned area model that parameterize $F(y_{(s,t)})$, and a collection of posterior draws of $n_{s,t}^{\text{rep}}$ from the posterior predictive distribution of a wildfire count model, a posterior distribution for the CDF or quantile function of maximum fire size can be generated which combines the two models to facilitate inference on the distribution of extremes.

Parameter estimation

We used a combination of variational approximations and Hamiltonian Monte Carlo methods to sample from the posterior distributions of count and burned area models. A variational approximation (Kucukelbir et al. 2015) was used for count models to quickly identify a preferred model. The best performing count model and all burned area models were fit using the No-U-Turn Sampler (Hoffman and Gelman 2014). Models were fit in the Stan probabilistic programming language using the `rstan` package (Carpenter et al. 2016, Stan Development Team 2018). We ran four chains for 1,000 iterations each, discarding the first 500 iterations as warmup. Convergence was assessed using visual inspection of trace plots, with potential scale reduction statistic values $\hat{R} \geq 1.1$ as an indicator convergence failure (Brooks and Gelman 1998).

Implementation

All data processing, model fitting, and visualization were implemented with open source software, primarily in the R programming language (R Core Team 2017), and wrapped in a reproducible workflow via GNU Make and Docker (Stallman et al. 2004, Boettiger 2015). Data cleaning and transformation required the R packages `assertthat` (Wickham 2017a), `lubridate` (Grolemund and Wickham 2011), `Matrix` (Bates and Maechler 2018), `pbapply` (Solymos and Zawadzki 2018), `splines` (R Core Team 2018), `tidyverse` (Wickham 2017b), and `zoo` (Zeileis and Grothendieck 2005). Spatial data were processed with `raster` (Hijmans 2017), `rgdal` (Bivand et al. 2018), `sf` (Pebesma 2018), and `spdep` (Bivand and Piras 2015). Finally, we used `cowplot` (Wilke 2017), `ggrepel` (Slowikowski 2018), `ggthemes` (Arnold 2018), `patchwork` (Pedersen 2017), and `RColorBrewer` (Neuwirth 2014) for visualization. The manuscript was written in R Markdown (Allaire et al. 2018). Analyses were run on an Amazon Web Services `m5.2xlarge` EC2 instance with four physical cores and 32 GB of RAM, and the whole workflow requires ~ 72 h. All code to reproduce the analysis is available online (see *Data Availability*; Joseph 2019).

RESULTS

Wildfire occurrence

The zero-inflated negative binomial distribution performed best on the held-out test set (Table 1), and was

TABLE 1. Performance of count models on the test set in descending order.

Model	Holdout log likelihood
ZI negative binomial	-3,671 (70)
ZI Poisson	-4,093 (77)
Negative binomial	-4,298 (114)
Poisson	-4,572 (139)

Note: Posterior means are provided with standard deviations in parentheses. ZI, Zero-inflated.

able to recover the proportion of zeros, count maxima, and count totals in posterior predictive checks for both the training and test data (Fig. 4). All of the other count models that we considered exhibited lack of fit to at least one of these statistics in posterior predictive checks. Hereafter, we report results from the zero-inflated negative binomial model.

Minimum relative humidity and maximum air temperature had the strongest effects on both the zero-inflation

component and the expected value of the negative binomial component (Fig. 5, posterior median for ρ , 0.665, 95% credible interval [CI]: 0.319–0.861). The model uncovered unique effects of meteorological variables at level 1, 2, and 3 ecoregions (Fig. 6). For example, a positive interaction effect between the second air temperature basis vector and the L1 Great Plains ecoregions indicates that the expected number of wildfires in plains ecoregions with cold conditions is high relative to other ecoregions. The Ozark/Ouachita-Appalachian forest and Ozark Highlands were also identified as having region-specific temperature effects (Fig. 6). Twelve month total precipitation also had region specific effects in the Mississippi Alluvial and Southeast Coastal Plains ecoregion, where it was associated with lower expected fire counts (Fig. 6). In contrast, increasing cumulative twelve month precipitation was associated with higher counts in desert ecoregions (Fig. 5). Housing density showed a unimodal relationship to expected count (Fig. 5), with lower expected counts in sparsely populated ecoregions,

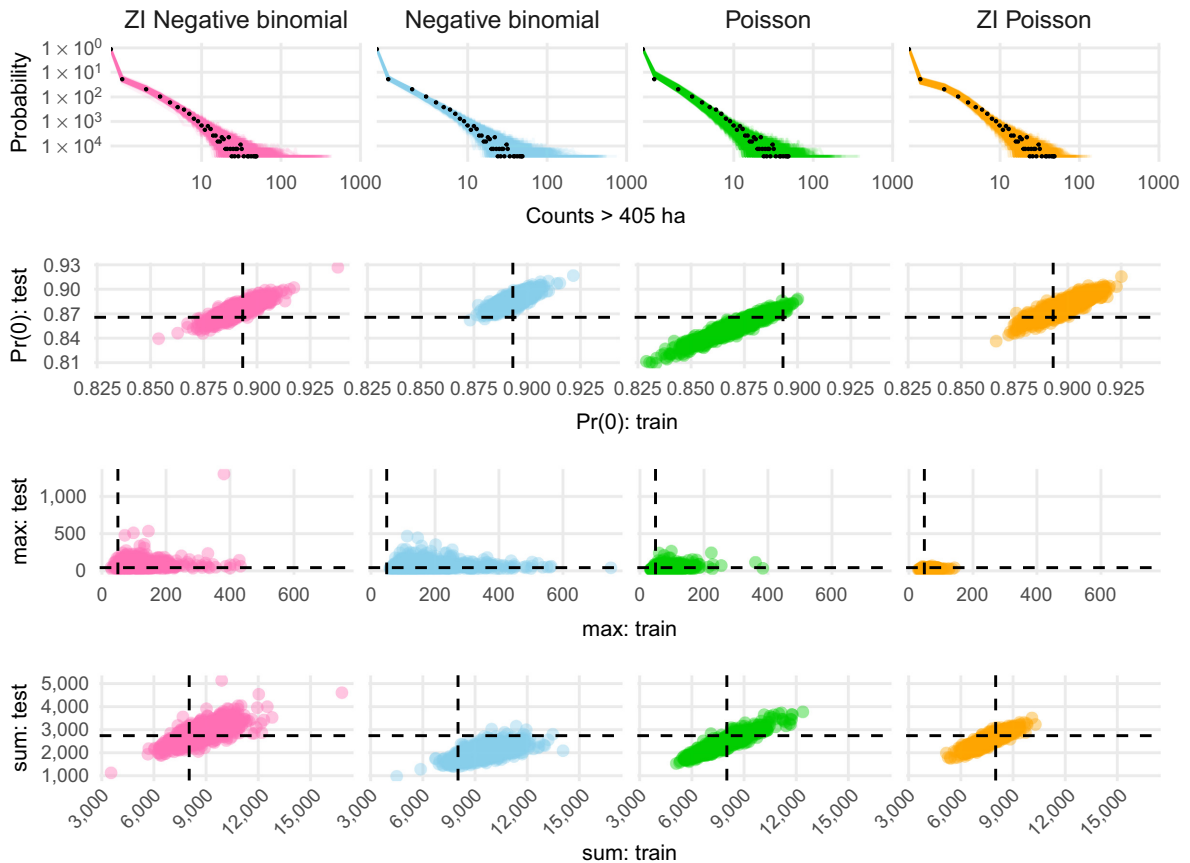


FIG. 4. Count predictive checks. Row one shows observed count frequencies as black points and predicted frequencies as lines. Rows two, three, and four show predicted proportions of zeros, maxima, and sums (respectively) in the training and test data, with empirical values as dashed lines. Rows two through four facilitate comparison of performance on training and test sets. Ideally, model predictions cluster around the dashed lines for both the training (x -axis direction) and test (y -axis direction) sets, leading to a tight cluster of points at the intersection of the dashed lines. ZI, zero inflation; max: test is the maximum value in the test data, sum: test is the sum of values in the test data, Pr(0): test is the probability of zero in the test data, while max: train, sum: train, and Pr(0): train represent the same quantities for the training data.




FIG. 5. Partial effects on the log-transformed negative binomial mean component of the zero-inflated negative binomial model for each level 3 ecoregion, colored by level 1 ecoregion. Lines are posterior medians. Results are similar for the zero-inflation component. Housing density was measured as units/km².

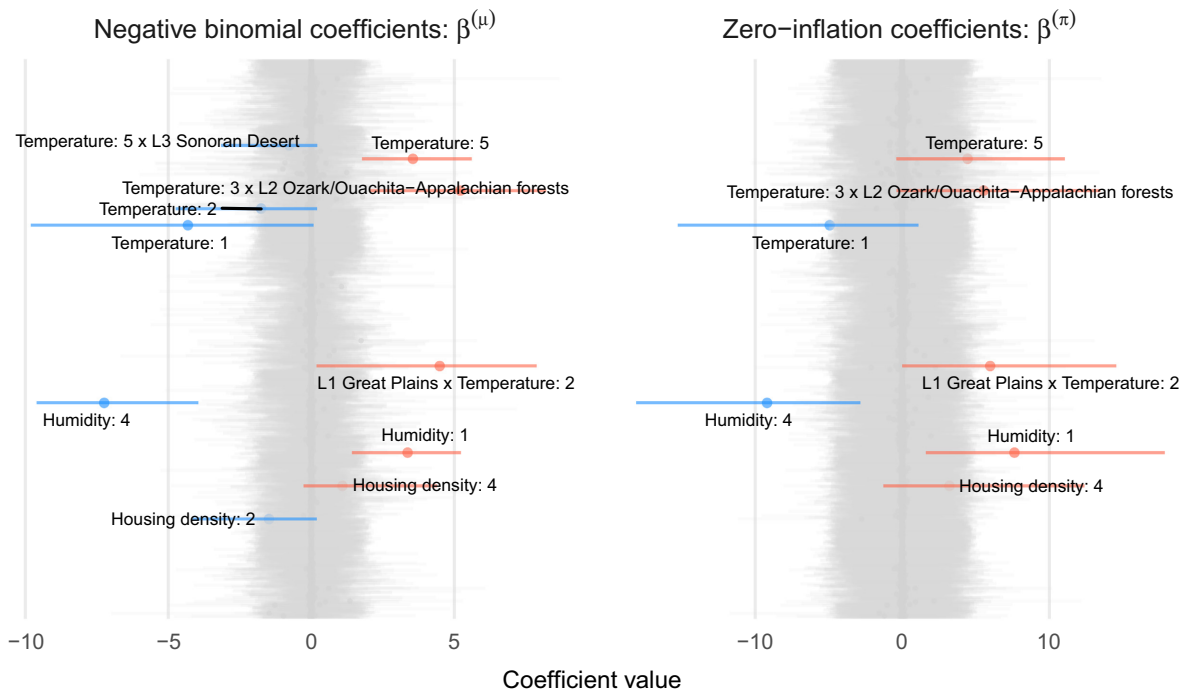


FIG. 6. Caterpillar plots of zero inflated negative binomial model coefficients, $\beta^{(\mu)}$ (left) and $\beta^{(\pi)}$ (right). Horizontal line segments denote 95% credible intervals. Gray segments indicate coefficients with a <87% posterior probability of being positive or negative, and colored segments indicate coefficients that are probably positive (red) or negative (blue). B-spline vectors are indicated by colons, e.g., Humidity : 1 indicates the first basis vector corresponding to humidity. Interactions between variables a and b are represented as a x b. Level 1 ecoregions are represented by L1 ecoregion name, and L2 and L3 indicate level 2 and 3 ecoregions.

and higher expected counts with moderately populated ecoregions.

Posterior 95% credible interval coverage for the number of fires over 405 ha in the test set was 98.8%. The

lowest test set interval coverage was 89.3%, in the Cross Timbers L3 ecoregion. When observed counts fell outside the 95% prediction interval, counts were larger than predicted 100% of the time. The largest difference

between observed numbers and predicted 97.5% posterior quantiles (the upper limit for the 95% credible interval) occurred for the Columbia Mountains/Northern Rockies L3 ecoregion in August 2015, when 36 fires over 405 ha occurred and at most 22 were predicted. For nearly half of the level 3 ecoregions (43 of 85), accounting for 39.7% of the land area of the contiguous United

States, the zero-inflated negative binomial model had 100% test set prediction interval coverage.

Wildfire burned areas

The lognormal distribution performed best on the test set (Table 2), and captured tail-behavior better than other burned area distributions (Fig. 7). The GPD model was too heavy-tailed to adequately capture the pattern in the empirical data, predicting fires far larger than those observed in the training and test sets (Fig. 7). The tapered Pareto distribution was too light-tailed (Fig. 7). The gamma and Weibull models performed very poorly overall on the test set (Table 2), apparently due to a lack of congruence between the shapes of these distributions and the actual burned area distribution. Despite a poor fit to the bulk of the wildfire burned area distribution, both performed adequately in the upper tails (Fig. 7). Hereafter we present results for the lognormal model, which had the highest test set log likelihood

TABLE 2. Performance of burned area models on the test set in descending order.

Model	Holdout log likelihood
Lognormal	-26,341 (43)
Generalized Pareto	-26,377 (45)
Tapered Pareto	-26,386 (49)
Weibull	-27,592 (236)
Gamma	-30,675 (993)

Note: Posterior means are provided with standard deviations in parentheses.




FIG. 7. Predictive checks for burned area models. The top row shows predicted density in color and empirical density for the training set in black, which reveals overall lack of fit for the gamma and Weibull models. Row two shows the complementary cumulative distribution function (CCDF) at the tails, with 95% and 50% prediction intervals shown in color and observed data as black points, which shows that the Generalized Pareto distribution predicts values that are too extreme. The third and fourth rows show checks for maximum and total burned areas in the training and test set, with observed values as dashed lines and posterior draws as colored points. These final two rows facilitate checks for summary statistics on both the training and test set, with the ideal model generating predictions (colored points) clustered close to where the dashed lines intersect.

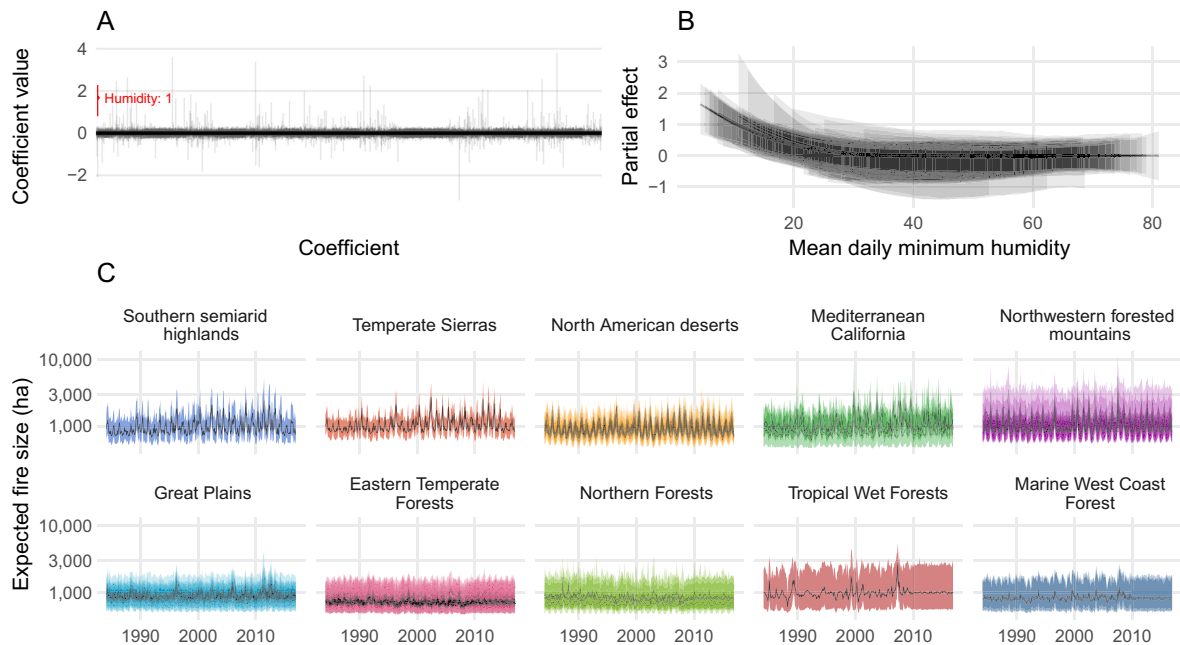


FIG. 8. (A) Estimated posterior medians and 95% credible intervals for each of the 3,473 coefficients associated with expected burned area. Only one coefficient, the first basis vector for humidity, had a 95% credible interval that excluded zero, shown in red. (B) Visualization of effect from panel A. Partial effects of mean daily minimum humidity for each level 3 ecoregion, with posterior medians drawn as lines, and the 95% credible intervals as ribbons. (C) Monthly time series of expected fire sizes for every level 3 ecoregion, faceted and colored by level 1 ecoregions sorted by mean humidity. Lines are posterior medians and ribbons are 95% credible intervals.

and captured tail behavior of the empirical fire size distribution.

Relative humidity was the primary driver of expected burned area for a fire event (Fig. 8A). The first basis vector for mean daily minimum relative humidity was the only coefficient with a 95% credible interval that did not include zero (posterior median: 1.68, 95% CI 0.8–2.29). This nonlinear effect can be observed in Fig. 8B as an increase in the expected burned area below 20% mean daily minimum humidity. This leads to a seasonality gradient among ecoregions of expected fire sizes, with little or no seasonal signal in typically humid ecoregions such as Marine West Coast Forests of the Pacific Northwest, and seasonal oscillations in ecoregions that have periodic fluctuations between dry and humid conditions such as the Temperate Sierras (Fig. 8C). There was not strong evidence that meteorological variables had spatially variable effects on expected wildfire burned area.

Overall, 95% posterior predictive interval coverage in the test set for burned areas was 93%. The lowest test set coverage was 0%, for the Eastern Great Lakes Lowlands L3 ecoregion, followed by 50%, for the Central California Valley L3 ecoregion, though these ecoregions had just one and two wildfire events in the test set. When observed fire sizes fell outside the 95% prediction interval, 24.9% of wildfires were smaller than predicted, and 75.1% of wildfires were larger than predicted. The largest discrepancy between the actual size of a wildfire and

the predicted 97.5% posterior quantile was observed with the Wallow Fire in 2011, which burned 228,107 ha, but the predicted upper limit for size was 20,756. We investigate this discrepancy further in the case study below. The lognormal burned area model achieved 100% interval coverage in 24 of 67 ecoregions that had wildfire events in the test set, accounting for 26% of the land area of the contiguous United States.

Inference on extremes

By combining the output of the event count and burned area models, we derived posterior prediction intervals for the size of the largest fire in a month for each region (the “burned area maximum”), integrating over uncertainty in the number of fires, as well as the lognormal mean and standard deviation for burned area. We evaluated the posterior distribution for the quantile function of the finite sample maximum of a lognormal distribution ($\exp(\mu + \sigma\sqrt{2}\text{erf}^{-1}(2P^{1/n} - 1))$, where n is the number of wildfire events, erf^{-1} is the inverse error function, and P is a probability) to generate prediction intervals for maximum fire sizes by month and ecoregion, conditional on one or more fires having occurred. In the holdout period from 2010 to 2016, a 99% prediction interval achieved 77.4% interval coverage, with 14.8% of the burned area maxima (140 fire events) being larger than predicted (Fig. 9). As an additional check,




FIG. 9. Posterior 99% (light red) and 95% (dark red) prediction intervals for the burned area of the largest fire event by month and level 3 ecoregion in the test set, shown for ecoregions with wildfires in more than 20 months. Empirical maxima are shown as black dots.

we used the posterior distribution of the model to predict the total area burned by wildfires in the test set. The model predicted the total area burned over the entire contiguous United States in test period from 2010 to 2016 to be 30,597,494 (95% CI 20,423,596–49,741,080) and the actual value was 30,440,173.

While fires over a million acres (~404,686 ha) in size have happened historically in the contiguous United States (Pernin 1971), no such fires were represented in in the training or test sets. If we extrapolate, the probability of at least one fire this large in the period from 2010 to 2016 was estimated to be between 0.191 and 0.651 (95% CI), with a posterior median of 0.348. The highest probability for such an event was 0.014 (posterior median), with a 95% CI of (0, 0.237) seen for the Southwestern Tablelands ecoregion in June 2011. The second highest probability was 0.004 (posterior median), with a 95% CI of (0, 0.056) seen for the Arizona/New Mexico Mountains ecoregion in June 2011. Aggregating spatially, we estimated monthly probabilities of a million acre (1 acre = 0.40 ha) wildfire. These probabilities show seasonal signals corresponding to peak fire seasons, with a shift toward higher and broader peaks beginning in the 21st century (Fig. 10).

Error analysis case study: the 2011 Wallow Fire

To better understand how well the model could or could not anticipate notable extreme events, and why, we used the largest fire in the test set as a case study. The Wallow Fire was accidentally ignited on 29 May 2011 by two campers in the L3 Arizona/New Mexico Mountains




FIG. 10. Estimated monthly posterior probabilities that one or more fire events exceed one million acres (404,686 ha). The line represents the posterior median, and shaded region represents an 80% credible interval. The training period up to 2010 is shown in black, and the test period for which data were withheld during parameter estimation is shown in red.

ecoregion. It burned through the month of June and into early July. The model underpredicted the total burned area of the Wallow Fire. Integrating over uncertainty in the predicted number of fires and expected fire size, the 99% credible interval for the maximum fire size for May 2011 was 730–107,419 ha, but the Wallow Fire is recorded as 228,103 ha.

We evaluated the contribution of each covariate to the linear predictor functions of the three model components (lognormal mean for burned areas, negative binomial mean for counts, and the logit probability of the

zero-inflation component) to understand why these predictions differed. We defined the contribution of a variable as the dot product of the elements in the design matrix \mathbf{X} corresponding to a particular driver variable (e.g., humidity), and the estimated coefficients in β corresponding to that variable. This provides a quantitative measure of how each input variable contributes to the linear predictor for an ecoregion, and incorporates the overall, level 1, level 2, and level 3 ecoregion adjustments on these effects. Humidity is the primary driver of variation in the model's predictions overall, and June 2011, the month after ignition, favored more large fires, with drier, hotter conditions (Fig. 11). The 99% credible interval for June 2011 was 4,258–428,765 ha, which contains the true value. Had the Wallow Fire ignited two days later, the true final size would have been contained in the prediction interval. Evidently, conditions in May that drove (under)predictions of maximum burned area were not representative of the conditions over most of the Wallow Fire's duration.

Temporal mismatch aside, meteorological conditions local to the Wallow Fire differed from the monthly regional means (Fig. 12). In particular, wind speeds in the Wallow Fire vicinity exceeded the regional monthly mean values on the date of ignition and in the weeks following ignition. Over the majority of the duration of the Wallow Fire (29 May to 8 July), local daily conditions were drier and hotter on average than regional mean monthly conditions in May, which were used to drive the statistical model. This local variability is not represented in the regional models developed here. The failure of the model to correctly

predict the size of the Wallow fire suggests potential avenues for improvement.

DISCUSSION

Extreme wildfires are often devastating, but perhaps they need not be surprising. By allowing the non-linear effects of weather and housing density to vary across space, this model achieves good predictive accuracy for fire extremes at a regional scale over a six year prediction window. This model predicts that extremely large wildfires, perhaps even over one million acres (404,686 ha), have a non-negligible probability of occurrence in the contiguous United States. Such predictions can support regional wildfire management and probabilistic hazard assessment.

Driving a model with meteorological features raises challenges related to predictive uncertainty and covariate shift: a change in the underlying distribution of forcing variables, potentially outside of the historic range. Ideally, this uncertainty would be propagated forward in a predictive model, possibly through stacking of predictive distributions that are generated from multiple models of future climate dynamics (Yao et al. 2017). But, even if one had a perfect forecast, novel conditions present a challenge for predictive modeling (Quionero-Candela et al. 2009). For example, the High Plains ecoregion had its highest mean monthly precipitation, lowest 12 month running precipitation, driest, hottest, and windiest conditions in the test set period. Extrapolating beyond the range of training inputs is generally difficult, but the hierarchical spatial effect specification used here




FIG. 11. Posterior median contribution of each input variable to the linear predictor function of model components for the Arizona/New Mexico Mountains level 3 ecoregion from 2010 to 2016. A dotted vertical line marks May 2011, when the Wallow Fire ignited. Vertical positions of colored lines show contributions to the linear predictor function of each model component.

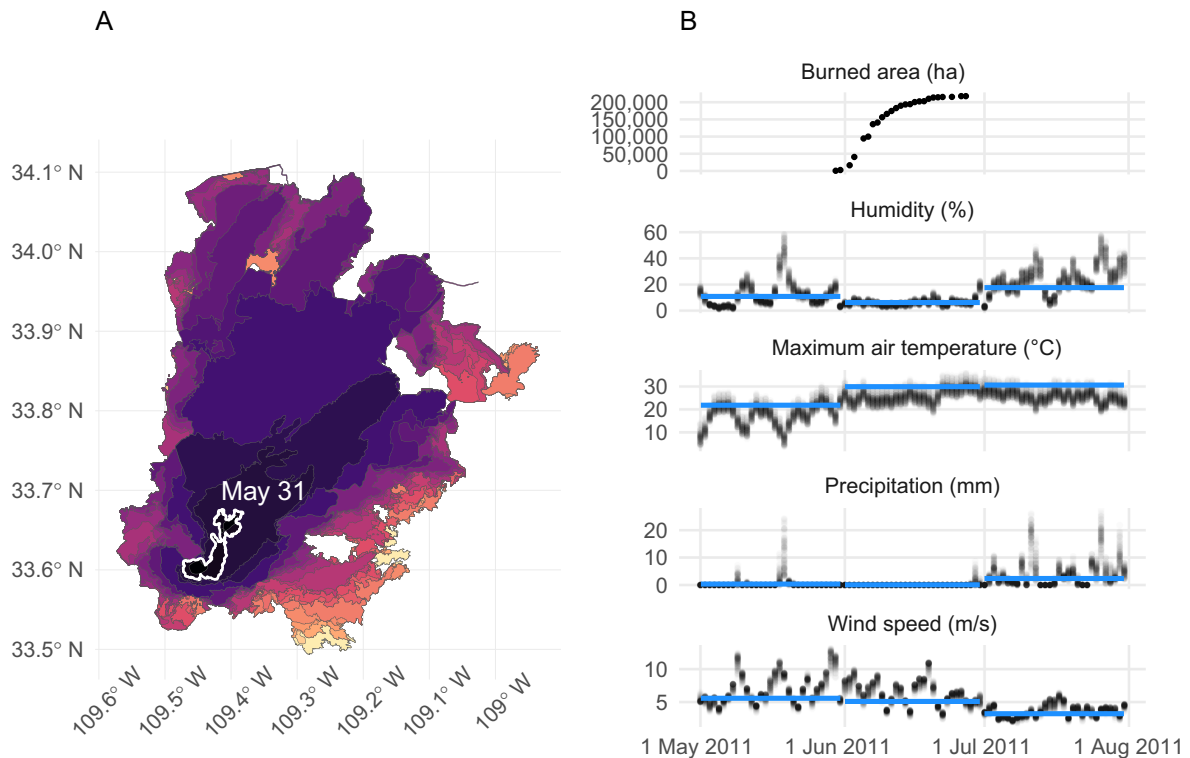


FIG. 12. (A) Progression of the Wallow Fire based on perimeter data collected from the GeoMAC database, spanning 30 May 2011 to 27 June 2011. The perimeter at the end of May is outlined in white, with brighter colors indicating later dates. (B) Daily local meteorological conditions and monthly regional mean conditions for the Wallow Fire, along with the associated burned area over time. The blue line represents monthly averages of meteorological quantities computed over the entire Arizona/New Mexico Mountains ecoregion, and black points represent values extracted for “local” 4-km grid cells contained within the final burned area perimeter of the Wallow Fire.

allows partial pooling among climatically similar ecoregions that can inform such predictions, unlike models fit separately to disjoint spatial regions.

Similar issues could arise when making predictions for observed but rare meteorological conditions. For example, mean daily minimum humidity values over 60% accounted for just 3.76% of the ecoregion-months in the training data, and 0 fires occurred in such months. As a consequence, there is relatively little data that can be used to inform the model for such conditions, and the prior distribution which shrinks coefficients toward zero may dominate the likelihood in the posterior distribution. In this case, the posterior distribution for the last basis coefficient for the partial effect of humidity is likely to be close to zero. This could explain why the estimated partial effect of humidity on the expected counts was less negative at the upper end of the observed humidity range, although previous work has found similarly nonlinear partial effects (Preisler et al. 2004). The count model performed extremely well in this range, with 100% interval coverage for the 299 ecoregion-months with mean daily minimum humidity values >60% in the withheld test data. The model nearly always predicted zero counts with high confidence when conditions were this humid: 298 of 299 predictions made for such

conditions were 95% credible intervals of (0, 0). The remaining prediction had a posterior median of zero, along with a 95% credible interval from 0 to 1. Monotonicity constraints could be incorporated into these models via monotonic spline bases (Ramsay et al. 1988), or an ordered prior distribution for basis coefficients (Brezger and Steiner 2008). In this case, the count model performs well under humid conditions without monotonicity constraints, and there seems to be little room for performance improvements that might result from such constraints.

Human-caused climate change is expected to increase fire activity in the western United States (Rogers et al. 2011, Westerling et al. 2011, Moritz et al. 2012, Abatzoglou and Williams 2016) and elsewhere (Flannigan et al. 2009), but the nonlinear effect of housing density could provide additional insight into future expectations. While housing density is increasing over time in most U.S. ecoregions, some of these ecoregions are in the range of values in which this increases the expected number of large fires, while others are so populated that further increases would reduce the chance of a large fire. The hump-shaped effect of human density on the expected number of large fires is likely driven by ignition pressure and fire suppression (Balch et al. 2017). As

human density increases from zero, ignition pressure increases, but eventually landscapes become so urbanized, fragmented, and/or fire-suppressed that wildfire risk decreases (Syphard et al. 2007, 2017, Bowman et al. 2011, Bistinas et al. 2013, Knorr et al. 2013, Mcwethy et al. 2013, Nagy et al. 2018). At intermediate density, wildfire regimes respond to human ignition and altered fuel distributions (Guyette et al. 2002), but these responses depend on environmental context and characteristics of the human population (Marlon et al. 2008, Li et al. 2009). This model indicates that the combination of moderate to high human density and dry conditions would nonlinearly increase the chance of an extreme fire event. Both human density and dryness are expected to increase in the future across large swaths of the United States (Radeloff et al. 2010, Stavros et al. 2014, Lloyd et al. 2017), with potential implications for human mortality, health risks from smoke and particulate emission, and the financial burden of wildfire management (Reid et al. 2016, Radeloff et al. 2018).

This work points to promising directions for future predictive efforts. Default choices such as Poisson and GPD distributions should be checked against alternative distributions. Further, the predictive skill of this model seems to suggest that ordinary events provide information on extremes, which would not be the case if the generative distribution of extremes was completely unique. Previous case studies have identified that extremes or anomalies in climatological drivers play a role in the evolution of extreme wildfires (Peterson et al. 2015), but for this work, monthly averages of climatological drivers over fairly large spatial regions were used, which may smooth over anomalous or extreme conditions. Enhancing the spatiotemporal resolution of predictive models could better represent climatic and social drivers and provide localized insights to inform decision-making. This raises computational challenges, but recent advances in distributed probabilistic computing (Tran et al. 2017), efficient construction of spatiotemporal point processes (Shirota and Banerjee 2018), and compact representations of nonlinear spatial interactions (Lee and Durbán 2011) may provide solutions.

The Wallow Fire case study reveals at least one limitation of increasing the spatiotemporal resolution. When the model predictions are driven by covariates that are summarized in space and time (e.g., a mean across an ecoregion in a month), summary values may not represent conditions that are most relevant to an event. With a discrete space-time segmentation, events can occur at the boundary of a spatiotemporal unit, e.g., if a fire spreads into an adjacent ecoregion or ignites on the last day of the month. Large wildfires can span months, and a model that only uses conditions upon ignition to predict total burned area can fail to account for conditions that change over the course of the event. Modeling ignitions as a point process in continuous space and time (Brillinger et al. 2003), and explicitly modeling subsequent fire duration and spread could better separate

conditions that ignite fires from those that affect propagation. Such an approach might be amenable to including information on fuel continuity, which is likely to limit the size of extremely large fires and did not factor into the current predictions (Hargrove et al. 2000, Rollins et al. 2002).

To the extent that a model reflects the generative process for extreme events, the decomposition of contributions to the model's predictions may provide insight into attribution for meteorological and anthropogenic drivers of extremes. However, a model trained to represent a region-wide distribution of fire sizes will inevitably fail to capture local factors that are relevant to specific events such as the Wallow Fire. If predicting the dynamics of particular fire events is a goal, process-based models designed to model fire spread are likely to be more appropriate than regional statistical models such as those developed here.

This paper presents and evaluates a statistical approach to explain and predict extreme wildfires that incorporates spatially varying non-linear effects. The model reveals considerable differences among ecoregions spanning the mountain west to the great plains, deserts, and eastern forests, and suggests a non-negligible chance of extreme wildfires larger than those seen in over the past 30 yr in the contiguous United States. Predictive approaches such as this can inform decision-making by placing probabilistic bounds on the number of wildfires and their sizes, while provide deeper insights into wildfire ecology.

ACKNOWLEDGMENTS

We thank Mitzi Morris, Kyle Foreman, Daniel Simpson, Bob Carpenter, and Andrew Gelman for contributing to the implementation of an intrinsic autoregressive spatial prior in Stan. We also thank two anonymous reviewers for their thoughtful comments, which greatly improved the paper.

LITERATURE CITED

- Abatzoglou, J. T. 2013. Development of gridded surface meteorological data for ecological applications and modelling. *International Journal of Climatology* 33:121–131.
- Abatzoglou, J. T., and A. P. Williams. 2016. Impact of anthropogenic climate change on wildfire across western US forests. *Proceedings of the National Academy of Sciences USA* 113:11770–11775.
- Allaire, J.J., Y. Xie, J. McPherson, J. Luraschi, K. Ushey, A. Atkins, H. Wickham, J. Cheng and W. Chang. 2018. Rmarkdown: dynamic documents for R. <https://CRAN.R-project.org/package=rmarkdown>
- Arnold, J. B. 2018. Ggthemes: extra themes, scales and geoms for 'Ggplot2'. <https://CRAN.R-project.org/package=ggthemes>
- Balch, J. K., B. A. Bradley, J. T. Abatzoglou, R. C. Nagy, E. J. Fusco, and A. L. Mahood. 2017. Human-started wildfires expand the fire niche across the United States. *Proceedings of the National Academy of Sciences USA* 114:2946–2951.
- Balshi, M. S., A. D. McGuire, P. Duffy, M. Flannigan, J. Walsh, and J. Melillo. 2009. Assessing the response of area burned to changing climate in western boreal North America using a

- multivariate adaptive regression splines (Mars) approach. *Global Change Biology* 15:578–600.
- Banerjee, S., B. P. Carlin, and A. E. Gelfand. 2014. Hierarchical modeling and analysis for spatial data. CRC Press, Boca Raton, Florida, USA.
- Barbero, R., J. T. Abatzoglou, E. A. Steel, and N. K. Larkin. 2014. Modeling very large-fire occurrences over the continental United States from weather and climate forcing. *Environmental Research Letters* 9:124009.
- Bates, D., and M. Maechler. 2018. Matrix: sparse and dense matrix classes and methods. <https://CRAN.R-project.org/package=Matrix>
- Beck, C., and E. G. D. Cohen. 2003. Superstatistics. *Physica A: Statistical Mechanics and Its Applications* 322:267–275.
- Bermudez, P. Z., J. Mendes, J. M. C. Pereira, K. F. Turkman, and M. J. P. Vasconcelos. 2009. Spatial and temporal extremes of wildfire sizes in Portugal (1984–2004). *International Journal of Wildland Fire* 18:983–991.
- Besag, J., and C. Kooperberg. 1995. On conditional and intrinsic autoregressions. *Biometrika* 82:733–746.
- Bistinas, I., D. Oom, A. C. L. Sá, S. P. Harrison, I. C. Prentice, and J. M. C. Pereira. 2013. Relationships between human population density and burned area at continental and global scales. *PLoS ONE*. <https://doi.org/10.1371/journal.pone.0081188>.
- Bivand, R., and G. Piras. 2015. Comparing implementations of estimation methods for spatial econometrics. *Journal of Statistical Software* 63:1–36. <https://doi.org/10.18637/jss.v063.i18>
- Bivand, R., T. Keitt, and B. Rowlingson. 2018. Rgdal: bindings for the ‘Geospatial’ data abstraction library. <https://CRAN.R-project.org/package=rgdal>
- Boettiger, C. 2015. An introduction to Docker for reproducible research. *ACM SIGOPS Operating Systems Review* 49:71–79.
- Bowman, D. M. J. S., et al. 2011. The human dimension of fire regimes on Earth. *Journal of Biogeography* 38:2223–2236.
- Brezger, A., and S. Lang. 2006. Generalized structured additive regression based on Bayesian P-Splines. *Computational Statistics & Data Analysis* 50:967–991.
- Brezger, A., and W. J. Steiner. 2008. Monotonic regression based on Bayesian P-Splines: an application to estimating price response functions from store-level scanner data. *Journal of Business & Economic Statistics* 26:90–104.
- Brillinger, D. R., H. K. Preisler, and J. W. Benoit. 2003. Risk assessment: a forest fire example. *Lecture Notes-Monograph Series* 30:177–196.
- Brooks, S. P., and A. Gelman. 1998. General methods for monitoring convergence of iterative simulations. *Journal of Computational and Graphical Statistics* 7:434–455.
- Carpenter, B., A. Gelman, M. Hoffman, D. Lee, B. Goodrich, M. Betancourt, M. A. Brubaker, J. Guo, P. Li, and A. Riddell. 2016. Stan: a probabilistic programming language. *Journal of Statistical Software* 20:1–37.
- Coles, S. 2014. An introduction to statistical modeling of extreme values. Springer. <https://books.google.com/books?id=G-D-sgEACAAJ>
- Coles, S., J. Bawa, L. Trenner, and P. Dorazio. 2001. An introduction to statistical modeling of extreme values. Volume 208. Springer, London, UK.
- Coles, S., L. R. Pericchi, and S. Sisson. 2003. A fully probabilistic approach to extreme rainfall modeling. *Journal of Hydrology* 273:35–50.
- Cox, D. R., and V. Isham. 1980. Point processes. Volume 12. CRC Press, Boca Raton, Florida, USA.
- Daly, C., M. Halbleib, J. I. Smith, W. P. Gibson, M. K. Doggett, G. H. Taylor, J. Curtis, and P. P. Pasteris. 2008. Physiographically sensitive mapping of climatological temperature and precipitation across the conterminous United States. *International Journal of Climatology* 28:2031–2064.
- Davison, A. C., and R. L. Smith. 1990. Models for exceedances over high thresholds. *Journal of the Royal Statistical Society. Series B (Methodological)* 52:393–442.
- Dennison, P. E., S. C. Brewer, J. D. Arnold, and M. A. Moritz. 2014. Large wildfire trends in the Western United States, 1984–2011. *Geophysical Research Letters* 41:2928–2933.
- Diaz, J. M. 2012. Economic impacts of wildfire. Joint Fire Science Program, Southern Fire Exchange 2012 Fact Sheet http://www.southernfireexchange.org/SFE_Publications/factsheets/2012-7.pdf
- Díaz-Avalos, C., P. Juan, and L. Serra-Saurina. 2016. Modeling fire size of wildfires in Castellon (Spain), using spatiotemporal marked point processes. *Forest Ecology and Management* 381:360–369.
- Dubey, S. D. 1970. Compound gamma, beta and F distributions. *Metrika* 16:27–31.
- Eidenshink, J., B. Schwind, K. Brewer, Z. L. Zhu, B. Quayle, and S. Howard. 2007. A project for monitoring trends in burn severity. *Nutrition and Cancer* 58:28–34.
- Flannigan, M. D., M. A. Krawchuk, W. J. de Groot, B. M. Wotton, and L. M. Gowman. 2009. Implications of changing climate for global wildland fire. *International Journal of Wildland Fire* 18:483–507.
- Fosberg, M. A. (1978) Weather in wildland fire management: the Fire Weather Index. United States Forest Service Reprints of Articles by Forest Service Employees, Washington, D.C., USA.
- Gelman, A., X.-L. Meng, and Hal Stern. 1996. Posterior predictive assessment of model fitness via realized discrepancies. *Statistica Sinica* 6:733–760.
- Gelman, A., J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin. 2013. Bayesian data analysis. Third edition. Chapman & Hall/CRC Texts in Statistical Science. Taylor & Francis, London, UK. <https://books.google.com/books?id=ZXL6AQAAQBAJ>
- Goodrick, S. L. 2002. Modification of the Fosberg fire weather index to include drought. *International Journal of Wildland Fire* 11:205–211.
- Grolemund, G., and H. Wickham. 2011. Dates and times made easy with lubridate. *Journal of Statistical Software* 40:1–25.
- Guyette, R. P., R. M. Muzika, and D. C. Dey. 2002. Dynamics of an anthropogenic fire regime. *Ecosystems* 5:472–486.
- Hargrove, W. W., R. H. Gardner, M. G. Turner, W. H. Romme, and D. G. Despain. 2000. Simulating fire patterns in heterogeneous landscapes. *Ecological Modelling* 135:243–263.
- Hernandez, C., C. Kerbin, P. Drobinski, and S. Turquety. 2015. Statistical modelling of wildfire size and intensity: a step toward meteorological forecasting of summer extreme fire risk. *Annales Geophysicae* 33:1495–1506.
- Hijmans, R. J. 2017. Raster: Geographic data analysis and modeling. <https://CRAN.R-project.org/package=raster>
- Hoffman, M. D., and A. Gelman. 2014. The No-U-Turn Sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research* 15:1593–1623.
- Hosking, J. R. M., and J. R. Wallis. 1987. Parameter and quantile estimation for the generalized Pareto distribution. *Technometrics* 29:339–349.
- Jiang, Y., and Q. Zhuang. 2011. Extreme value analysis of wildfires in Canadian Boreal Forest Ecosystems. *Canadian Journal of Forest Research* 41:1836–1851.
- Joseph, M. (2019) GitHub repository [mjoseph/wildfire-extremes](https://github.com/mjoseph/wildfire-extremes): code and data to reproduce the analysis.
- Kneib, T., T. Hothorn, and G. Tutz. 2009. Variable selection and model choice in geospatial regression models. *Biometrics* 65:626–634.

- Knorr, W., T. Kaminski, A. Arneth, and U. Weber. 2013. Impact of human population density on fire frequency at the global scale. *Impact of human population density on fire frequency at the global scale*. *Impact of human population density on fire frequency at the global scale*. *Biogeosciences Discuss* 10:15735–15778.
- Kochi, I., G. H. Donovan, P. A. Champ, and J. B. Loomis. 2010. The economic cost of adverse health effects from wildfire-smoke exposure: a review. *International Journal of Wildland Fire* 19:803–817.
- Krawchuk, M. A., and M. A. Moritz. 2011. Constraints on global fire activity vary across a resource gradient. *Ecology* 92:121–132.
- Krawchuk, M. A., M. A. Moritz, M.-A. Parisien, J. Van Dorn, and K. Hayhoe. 2009. Global pyrogeography: the current and future distribution of wildfire. *PLoS ONE* 4:e5102.
- Kucukelbir, A., R. Ranganath, A. Gelman, and D. Blei. 2015. Automatic variational inference in Stan. Pages 568–576 in *Proceedings of the 28th International Conference on Neural Information Processing Systems*, MIT Press, Cambridge, Massachusetts, USA.
- Lambert, D. 1992. Zero-inflated Poisson regression, with an application to defects in manufacturing. *Technometrics* 34:1–14.
- Lee, D.-J., and M. Durbán. 2011. P-Spline Anova-type interaction models for spatio-temporal smoothing. *Statistical Modelling* 11:49–69.
- Li, L.-M., W.-G. Song, J. Ma, and K. Satoh. 2009. Artificial neural network approach for modeling the impact of population density and weather parameters on forest fire risk. *International Journal of Wildland Fire* 18:640.
- Lloyd, C. T., A. Sorichetta, and A. J. Tatem. 2017. High resolution global gridded data for use in population studies. *Scientific Data* 4:170001.
- Marani, M., and M. Ignaccolo. 2015. A metastatistical approach to rainfall extremes. *Advances in Water Resources* 79:121–126.
- Marlon, J. R., P. J. Bartlein, C. Carcaillet, D. G. Gavin, S. P. Harrison, P. E. Higuera, F. Joos, M. J. Power, and I. C. Prentice. 2008. Climate and human influences on global biomass burning over the past two millennia. *Nature Geoscience*. <https://doi.org/10.1038/ngeo313>.
- McLaughlin, S. P., and J. E. Bowers. 1982. Effects of wildfire on a Sonoran Desert Plant Community. *Ecology*. JSTOR: 246–248.
- Mcwethy, D. B., et al. 2013. A conceptual framework for predicting temperate ecosystem sensitivity to human impacts on fire regimes. *Global Ecology and Biogeography* 22:900–912.
- Mendes, J. M., P. C. Z. Bermudez, J. Pereira, K. F. Turkman, and M. J. P. Vasconcelos. 2010. Spatial extremes of wildfire sizes: Bayesian hierarchical models for extremes. *Environmental and Ecological Statistics* 17:1–28.
- Moritz, M. A., M.-A. Parisien, E. Batllori, M. A. Krawchuk, J. Van Dorn, D. J. Ganz, and K. Hayhoe. 2012. Climate change and disruptions to global fire activity. *Ecosphere* 3:1–22.
- Nagy, R., E. Fusco, B. Bradley, J. T. Abatzoglou, and J. Balch. 2018. Human-related ignitions increase the number of large wildfires across US ecoregions. *Fire* 1:4.
- Nauslar, N., J. Abatzoglou, and P. Marsh. 2018. The 2017 North Bay and Southern California fires: a case study. *Fire* 1:18.
- Neuwirth, E. 2014. RColorBrewer: ColorBrewer Palettes. <https://CRAN.R-project.org/package=RColorBrewer>
- Omernik, J. M. 1987. Ecoregions of the Conterminous United States. *Annals of the Association of American Geographers* 77:118–125.
- Omernik, J. M., and G. E. Griffith. 2014. Ecoregions of the conterminous United States: evolution of a hierarchical spatial framework. *Environmental Management* 54:1249–1266.
- Pebesma, E. 2018. Sf: Simple Features for R. <https://CRAN.R-project.org/package=sf>
- Pechony, O., and D. T. Shindell. 2009. Fire parameterization on a global scale. *Journal of Geophysical Research Atmospheres* 114: D16115. <https://doi.org/10.1029/2009JD011927>
- Pechony, O., and D. T. Shindell. 2010. Driving forces of global wildfires over the past millennium and the forthcoming century. *Proceedings of the National Academy of Sciences USA* 107:19167–19170.
- Pedersen, T. L. 2017. Patchwork: the composer of Ggplots. <https://github.com/thomasp85/patchwork>
- Peltola, T., A. S. Havulinna, V. Salomaa, and A. Vehtari. 2014. Hierarchical Bayesian survival analysis and projective covariate selection in cardiovascular event risk prediction. Pages 79–88 in *Proceedings of the Eleventh UAI Conference on Bayesian Modeling Applications Workshop-Volume 1218*. CEUR-WS.org
- Pernin, P. 1971. The Great Peshtigo fire: an eyewitness account. *Wisconsin Magazine of History* 54:246–272.
- Peterson, D. A., E. J. Hyer, J. R. Campbell, M. D. Fromm, J. W. Hair, C. F. Butler, and M. A. Fenn. 2015. The 2013 Rim Fire: implications for predicting extreme fire spread, pyroconvection, and smoke emissions. *Bulletin of the American Meteorological Society* 96:229–247.
- Piironen, J., et al. 2017. Sparsity information and regularization in the horseshoe and other shrinkage priors. *Electronic Journal of Statistics* 11:5018–5051.
- Preisler, H. K., and A. L. Westerling. 2007. Statistical model for forecasting monthly large wildfire events in Western United States. *Journal of Applied Meteorology and Climatology* 46:1020–1030.
- Preisler, H. K., D. R. Brillinger, R. E. Burgan, and J. W. Benoit. 2004. Probability based models for estimation of wildfire risk. *International Journal of Wildland Fire* 13:133–142.
- Quionero-Candela, J., M. Sugiyama, A. Schwaighofer, and N. D. Lawrence. 2009. *Dataset shift in machine learning*. The MIT Press, Cambridge, Massachusetts, USA.
- R Core Team. 2017. R: a language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org/>
- R Core Team. 2018. R: a language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org/>
- Radeloff, V. C., S. I. Stewart, T. J. Hawbaker, U. Gimmi, A. M. Pidgeon, C. H. Flather, R. B. Hammer, and D. P. Helmers. 2010. Housing growth in and near United States protected areas limits their conservation value. *Proceedings of the National Academy of Sciences* 107:940–945.
- Radeloff, V. C., et al. 2018. Rapid growth of the US Wildland-urban interface raises wildfire risk. *Proceedings of the National Academy of Sciences* 115:3314–3319.
- Ramsay, J. O., et al. 1988. Monotone regression splines in action. *Statistical Science* 3:425–441.
- Reed, W. J., and K. S. McKelvey. 2002. Power-law behaviour and parametric models for the size-distribution of forest fires. *Ecological Modelling* 150:239–254.
- Reid, C. E., M. Brauer, F. H. Johnston, M. Jerrett, J. R. Balmes, and C. T. Elliott. 2016. Critical review of health impacts of wildfire smoke exposure. *Environmental Health Perspectives* 124:1334.
- Rogers, B. M., R. P. Neilson, R. Drapek, J. M. Lenihan, J. R. Wells, D. Bachelet, and B. E. Law. 2011. Impacts of climate change on fire regimes and carbon stocks of the U.S. Pacific Northwest. *Journal of Geophysical Research Biogeosciences* 116. <https://doi.org/10.1029/2011JG001695>.
- Rollins, M. G., P. Morgan, and T. Swetnam. 2002. Landscape-scale controls over 20th century fire occurrence in two large

- Rocky Mountain (USA) wilderness areas. *Landscape Ecology* 17:539–557.
- Schoenberg, F. P., R. Peng, and J. Woods. 2003. On the distribution of wildfire sizes. *Environmetrics* 14:583–592.
- Serra, L., M. Saez, P. Juan, D. Varga, and J. Mateu. 2014a. A spatio-temporal Poisson hurdle point process to model wildfires. *Stochastic Environmental Research and Risk Assessment* 28:1671–1684.
- Serra, L., M. Saez, J. Mateu, D. Varga, P. Juan, C. Díaz-Ávalos, and H. Rue. 2014b. Spatio-temporal log-Gaussian Cox processes for modelling wildfire occurrence: the case of Catalonia, 1994–2008. *Environmental and Ecological Statistics* 21:531–563.
- Shirota, S. and S. Banerjee. (2018) Scalable inference for space-time Gaussian Cox processes. arXiv Preprint arXiv:1802.06151
- Slowikowski, K. 2018. Ggrepel: Automatically position non-overlapping text labels with ‘Ggplot2’. <https://CRAN.R-project.org/package=ggrepel>
- Solymos, P., and Z. Zawadzki. 2018. Pbapply: adding progress bar to ‘*Apply’ functions. <https://CRAN.R-project.org/package=pbapply>
- Stallman, R. M., R. McGrath, and P. D. Smith. 2004. GNU Make: a program for directing recompilation, for Version 3.81. Free Software Foundation, Boston, Massachusetts, USA.
- Stan Development Team. 2018. RStan: the R interface to Stan. <http://mc-stan.org/>
- Stavros, E. N., J. T. Abatzoglou, D. McKenzie, and N. K. Larkin. 2014. Regional projections of the likelihood of very large wildland fires under a changing climate in the contiguous Western United States. *Climatic Change* 126:455–468.
- Syphard, A. D., V. C. Radeloff, J. E. Keeley, T. J. Hawbaker, M. K. Clayton, S. I. Stewart, and R. B. Hammer. 2007. Human influence on California fire regimes. *Ecological Applications* 17:1388–1402.
- Syphard, A. D., J. E. Keeley, A. H. Pfaff, and K. Ferschweiler. 2017. Human presence diminishes the importance of climate in driving fire activity across the United States. *Proceedings of the National Academy of Sciences USA*. <https://doi.org/10.1073/pnas.1713885114>.
- Tedim, F., et al. 2018. Defining extreme wildfire events: difficulties, challenges, and impacts. *Fire* 1:9.
- Tran, D, M.D. Hoffman, R.A. Saurous, E. Brevdo, K. Murphy and D. M. Blei. (2017) Deep probabilistic programming. arXiv Preprint arXiv:1701.03757
- Vilar, L., D. G. Woolford, D. L. Martell, and M. P. Martín. 2010. A model for predicting human-caused wildfire occurrence in the region of Madrid, Spain. *International Journal of Wildland Fire* 19:325–337.
- Westerling, A. L. 2016. Increasing Western US Forest Wildfire Activity: sensitivity to changes in the timing of spring. *Philosophical Transactions of the Royal Society B* 371:20150178.
- Westerling, A. L., B. P. Bryant, H. K. Preisler, T. P. Holmes, H. G. Hidalgo, T. Das, and S. R. Shrestha. 2011. Climate change and growth scenarios for California wildfire. *Climatic Change* 109:445–463.
- Wickham, H. 2017a. Assertthat: easy pre and post assertions. <https://CRAN.R-project.org/package=assertthat>
- Wickham, H. 2017b. Tidyverse: easily install and load the ‘Tidyverse’. <https://CRAN.R-project.org/package=tidyverse>
- Wittala, M. R. 1999. Assessing the risk of cumulative burned acreage using the Poisson probability model. *Fire Economics, Planning, and Policy: Bottom Lines* 51:51–57.
- Wilke, C. O. 2017. Cowplot: streamlined plot theme and plot annotations for ‘Ggplot2’. <https://CRAN.R-project.org/package=cowplot>
- Williams, J. 2013. Exploring the onset of high-impact mega-fires through a forest land management prism. *Forest Ecology and Management* 294:4–10.
- Wood, S. N. 2017. Generalized additive models: an introduction with R. Second edition. Chapman Hall/CRC, London, UK.
- Woolford, D. G., D. R. Bellhouse, W. J. Braun, Ch B Dean, D. L. Martell, and J. Sun. 2011. A spatio-temporal model for people-caused forest fire occurrence in the Romeo Malette Forest. *Journal of Environmental Statistics* 2:2–16.
- Woolford, D. G., C. B. Dean, D. L. Martell, J. Cao, and B. M. Wotton. 2014. Lightning-caused forest fire risk in northwestern Ontario, Canada, Is increasing and associated with anomalies in fire weather. *Environmetrics* 25:406–416.
- Yao, Y., et al. 2017. Using stacking to average bayesian predictive distributions. *Bayesian Analysis*. *International Society for Bayesian Analysis* 13:917–1007.
- Zeileis, A., and G. Grothendieck. 2005. Zoo: S3 infrastructure for regular and irregular time series. *Journal of Statistical Software* 14:1–27.
- Zorretto, E., G. Botter, and M. Marani. 2016. On the emergence of rainfall extremes from ordinary events. *Geophysical Research Letters* 43:8076–8082.

SUPPORTING INFORMATION

Additional supporting information may be found online at: <http://onlinelibrary.wiley.com/doi/10.1002/eap.1898/full>

DATA AVAILABILITY

Data and code are available on Zenodo: <https://doi.org/10.5281/zenodo.2605495>