

A Space-Time Conditional Intensity Model for Evaluating a Wildfire Hazard Index

Roger D. Peng

Frederic Paik Schoenberg

James Woods

Author's footnote: Roger D. Peng (rpeng@stat.ucla.edu) is a Graduate Student, Department of Statistics, University of California, Los Angeles CA 90095; Frederic Paik Schoenberg is Assistant Professor, Department of Statistics, University of California, Los Angeles CA 90095; and James Woods is GIS Lab Manager and Instructor, Department of Geography, California State University, Long Beach, CA 90840. This work is part of the first author's Ph.D. dissertation from the University of California, Los Angeles. The material is based upon work supported by the National Science Foundation under Grant No. 9978318. The authors thank Larry Bradshaw at the USDA Forest Service for providing the weather station data as well as LADPW and LACFD (esp. Mike Takeshita and Frank Vidales) for generously sharing their data and expertise.

Abstract

Numerical indices are commonly used as tools for assisting in wildfire management and hazard assessment. While the usage of such indices is widespread, assessment of these indices in their respective regions of application is rare. We evaluate the effectiveness of the Burning Index (BI) for predicting wildfire occurrences in Los Angeles County, California using space-time point process models. The models are based on an additive decomposition of the conditional intensity, with separate terms to describe spatial and seasonal variability as well as contributions from the BI. The models are fit to wildfire and BI data from the years 1976–2000 using a combination of nonparametric kernel smoothing methods and parametric maximum likelihood. In addition to using AIC to compare competing models, new multi-dimensional residual methods based on approximate random thinning and rescaling are employed to detect departures from the models and to ascertain the precise contribution of the BI to predicting wildfire occurrence. We find that while the BI appears to have a positive impact on wildfire prediction, the contribution is relatively small after taking into account natural seasonal and spatial variation. In particular, the BI does not appear to take into account increased activity during the years 1979–1981 and can overpredict during the early months of the year.

Keywords: Point process residual analysis; Random thinning; Random rescaling; Conditional intensity model; Model evaluation; Wildfire risk

1 Introduction

Fire departments all over the world often use numerical indices to aid in wildfire management. These indices are designed to summarize local meteorological and fuel information and provide an estimate of the current risk of fire. The Burning Index (BI) is part of the U.S. National Fire-Danger Rating System, a collection of numerical indices designed to be used for fire planning and management. In Los Angeles County, California, the Fire Department uses the BI for creating short-term wildfire hazard maps of the County which help managers make decisions involving the allocation of resources and the coordination of presuppression activities.

While the BI is already in common use by Los Angeles and other fire departments, there have been relatively few attempts to assess the Index’s performance in predicting wildfires. In the general area of index evaluation, there has been some work in evaluating elements of the U.S. system (e.g. Haines et al., 1983), various national (non-U.S.) systems (Viegas et al., 1999), and in using indices for prediction (Westerling et al., 2000). However, the BI’s ability to adapt to particular regions such as Los Angeles County has yet to be fully scrutinized. Mandallaz and Ye (1997) have noted that in general, wildfire hazard indices are developed on the basis of experience in a given area. Therefore, one must take caution when adapting indices to other areas.

The aim of this paper is to evaluate the performance of the BI in predicting wildfires in Los Angeles County. Our approach is to evaluate the best-fitting conditional intensity model both with and without the BI and other information, in order to determine not only the optimal use of the BI in point process prediction, but also to assess the increase in prediction performance using the BI as compared to other information. The various conditional intensity models are compared using the Akaike Information Criterion (AIC) as well as multi-dimensional residual analysis methods based on approximate random thinning and rescaling. While the AIC proves to be useful for finding the best model in a set of possibilities, residual analysis can identify specific areas where the performance is poor and suggest directions for improvement.

In the Sections to follow we briefly describe the U.S. National Fire-Danger Rating System and provide a summary of the data used for this analysis. We then outline the point process methodology used for evaluating the performance of the BI. Finally, the results of applying these methods to the wildfire data from Los Angeles County, California are discussed.

2 A Brief Summary of the National Fire-Danger Rating System

The U.S. National Fire-Danger Rating System (NFDRS) was developed by the U.S. Department of Agriculture Forest Service in 1972 (Deeming et al., 1972) and was revised in 1978 (Deeming et al., 1977; Bradshaw et al., 1983). Since then there have been some adjustments (see e.g. Burgan,

1988). The NFDRS actually consists of multiple components which can be combined to form three different indices, of which the BI is one. Although this is a “national” system, there are many parameters which can be calibrated to adapt the system to local environments. In particular, a fire manager must choose a *fuel model* (from a set of 20 available models) which corresponds to the available fuel in the region. The fuel model is then incorporated into the index computations to produce an index for a specific region (Bradshaw et al., 1983).

2.1 Computing the Burning Index

The BI is computed from the Spread Component (SC) and the Energy Release Component (ERC) of the NFDRS. Both the SC and the ERC are computed using meteorological and fuel data gathered by Remote Automatic Weather Stations (RAWS). The SC is simply the unmodified fire spread model of Rothermel (1972) which is a function of wind, slope, and various fuel properties. The ERC is a function of the loading-weighted reaction intensity and the surface area to volume ratio of the fuel bed. Given values for the SC and the ERC, the BI itself is computed via the relation $BI = 10 \times 0.45 \times [(SC/60)(25 \times ERC)]^{0.46}$ (Bradshaw et al., 1983). It is important to emphasize that the BI is to be interpreted as summarizing and integrating information from a variety of meteorological and fuel variables. That is, one would expect the index to reflect the current fuel conditions (such as fuel age or fuel moisture) in addition to the usual meteorological conditions.

3 Data

3.1 Wildfire Data

The wildfire data analyzed here were collected and compiled by various agencies, including the Los Angeles County Fire Department (LACFD), the Los Angeles County Department of Public Works, the Santa Monica Mountains National Recreation Area, the Ventura County Fire Department, and the California Department of Forestry and Fire Protection. The full dataset consists of origin dates and polygons mapping the areas burned by wildfires between January 1878 and September 2000.

LACFD officials have indicated that only fires burning greater than 100 acres were consistently mapped before 1950, though since 1950 the Department has mapped some fires as small as 1 acre. Fires prior to 1976 were excluded from the present analysis (due to the unavailability of BI data) and based on Fire Department guidance, a lower threshold of 10 acres was chosen for inclusion of post-1976 fires.

Figure 1(a) shows the spatial distribution of wildfires larger than 10 acres in Los Angeles County for 1976–2000. The locations are represented with a (scaled) state-plane coordinate system using the NAD 83 datum so that one spatial unit corresponds to approximately 18.9 miles. Much of the wildfire activity occurs in the Angeles National Forest and parts of the Los Padres National Forest (eastern and northern areas) as well as the Santa Monica Mountains area (the protrusion in the western part of the County). Figure 2 shows the times and areas burned (in acres) for each of the fires in the dataset. In the years 1979–1981, there appears to be some intense temporal clustering of points, especially for fires in the 50–500 acre range. In addition, there is a decreased level of activity around the years 1990–1991. We will return to these particular features of the data in Section 5.

3.2 Meteorological and Burning Index Data

Daily meteorological observations for eight RAWS around Los Angeles County were obtained from the USDA Forest Service. The locations for each of the RAWS are shown in Figure 1. The RAWS collect data on precipitation, wind direction, wind speed, air temperature, fuel temperature, and relative humidity (Warren and Vance, 1981). Collection of the data occurs at approximately 1:00 PM when conditions for fire are considered to be most severe.

For each of the 8 stations, daily values of the BI were then computed using the FireFamily Plus software (freely available from the Forest Service). Not all of the stations contained data covering the entire 25 year span from 1976 to 2000. Of the eight stations, only Stations 3 and 4 had data going back to 1976. However, each of the stations had at least five years of daily data.

The data from each of the RAWS exhibit the natural seasonal patterns for weather in Los

Angeles County. Figure 3 shows the average computed BI value (averaged over all available years) for each day in the year. There is a general increase in BI from July through September, followed by a decrease from October through March. Station 1 did not have any observations for the months of January, February, and March. However, in the entire 25 year interval of interest, only 10 fires ever occurred in the months of January, February, and March, representing less than 2% of the total number of fires. Therefore, for Station 1, the BI values were set to zero during that three month span. The other stations also contained days with missing weather records. In this situation we filled in a missing BI value on a given day with the average of that day across all the other available years. The percentages of missing data are shown in Table 1 for the off-season (January–April), the fire season (May–December) and overall. In Section 6 we discuss the possible impacts of missing data on the analysis.

4 Methodology

In evaluating the BI our approach considers the times and centroids of each fire as points of a space-time point process. A space-time point process N is a σ -finite counting measure on the spatial-temporal domain $\mathcal{S} \times \mathbb{R}^+$. Given a Borel set $B \subset \mathcal{S} \times \mathbb{R}^+$, $N(B)$ is the number of points in B . Let \mathcal{F}_t be a filtration, that is, an increasing family of σ -algebras, and take $N(\mathcal{S} \times [0, t])$ to be \mathcal{F}_t -adapted for each Borel set $S \in \mathcal{S}$. In applications, \mathcal{S} is usually taken to be a subset of \mathbb{R}^2 .

The conditional intensity function of a point process is defined as a non-negative \mathcal{F}_t -predictable process $\lambda(t, x, y)$ such that for each Borel set $S \in \mathcal{S}$

$$N(S \times [0, t]) - \iint_S \int_0^t \lambda(t, x, y) dt dx dy$$

is an \mathcal{F}_t -martingale. It is well known that for point processes with simple ground processes (i.e. with no two points at exactly the same time), the conditional intensity (when it exists) uniquely characterizes all of the finite-dimensional distributions of the point process. For a more thorough treatment of conditional intensities we refer the reader to Jacod (1975), Brémaud (1981), and Daley and Vere-Jones (2003). Intuitively, one may consider the conditional intensity function as describing

the conditional expected rate of occurrence of points at time t and location (x, y) . In particular, a version of the conditional intensity may be given by the process

$$\lambda(t, x, y) = \lim_{u \uparrow t} \lim_{\Delta u \Delta x \Delta y \downarrow 0} \frac{\mathbb{E}[N((u, u + \Delta u) \times (x, x + \Delta x) \times (y, y + \Delta y)) | \mathcal{F}_t]}{\Delta u \Delta x \Delta y}$$

for $(t, x, y) \in S \times \mathbb{R}^+$, provided the limit exists (Schoenberg et al., 2002).

In prescribing a model for λ , we use a spatial background component $m(x, y)$ which takes into account the spatial inhomogeneity of the wildfire occurrences. This component can be thought of as incorporating previous knowledge about where wildfires are more or less likely to occur. For example, it would be undesirable for the model to predict a wildfire occurrence in downtown Los Angeles. For this component a simple two-dimensional kernel smoother is used,

$$m(x, y) = \frac{1}{n_0} \sum_{j=1}^{n_0} \frac{1}{\phi_x \phi_y} K\left(\frac{x - x_{0j}}{\phi_x}\right) K\left(\frac{y - y_{0j}}{\phi_y}\right).$$

where K is a suitable kernel function. In estimating the spatial background $m(x, y)$, the smoother is not computed using the 1976–2000 data. Rather, the spatial locations of the wildfires occurring before 1976 are smoothed, guaranteeing that the estimate of the conditional intensity at time t is based strictly on information from before time t . Here, n_0 is the number of wildfires in the full dataset occurring before 1976 and (x_{0j}, y_{0j}) represents the spatial coordinates of the j th fire in that subset. Of legitimate concern is the similarity between the spatial configurations of the wildfires before and after 1976. Figure 1(b) shows the locations of wildfires occurring before 1976. One can see that the spatial distribution is quite similar between the two eras. A notable difference is the presence of about 25 wildfires in the far northeast corner of the county which appear after 1976 but not before. While this difference reflects a slight change in the spatial distribution over time, it is unlikely to reflect a major shift in the overall wildfire regime. Other authors have commented on the lack of evidence to suggest any major change in the Southern California regime over the past century (see Keeley et al., 1999; Keeley and Fotheringham, 2001). In addition, since this difference affects all of the models under consideration here, the effect on any comparison between models should be minimal.

A seasonal component $S(t)$ is used to describe the overall seasonal variation of the wildfire activity. Here we smooth the times within each year of the pre-1976 fires,

$$S(t) = \frac{1}{n_0} \sum_{j=1}^{n_0} \frac{1}{\phi_{\text{seas}}} K\left(\frac{t^* - t_{0j}^*}{\phi_{\text{seas}}}\right)$$

where t^* indicates the time since the beginning of the year and t_{0j}^* is the time since the beginning of the year for the j th wildfire occurrence before 1976.

Finally, a BI component $B(t, x, y)$ describes, for an arbitrary point (t, x, y) , the contribution to the conditional intensity from the values of the BI at each station. Since the BI values are only observed at fixed locations in the County, some form of interpolation is required to compute values at other locations. We consider a BI component of the form

$$B(t, x, y) = \frac{1}{C} \sum_{s \in S_t} \left\{ \gamma_s K\left(\frac{x - x_s}{\beta_s}\right) K\left(\frac{y - y_s}{\beta_s}\right) \text{BI}(t, s) \right\}$$

where $\text{BI}(t, s)$ is the BI value recorded at time t from the s th station, (x_s, y_s) represents the location of the s th station, and S_t is the index set of stations in use at time t . The values γ_s are scaling coefficients for the BI values and are in units of events / (spatial unit² × day); the BI itself is dimensionless. The normalization constant C is simply the sum of the kernel weights. For each of the three components, the kernel function used is the standard normal density function.

One may consider as a basis of comparison simple baseline models such as a homogeneous Poisson model

$$\lambda^H(t, x, y) = \mu \tag{1}$$

and the following Spatial+Seasonal model,

$$\lambda^0(t, x, y) = \nu m(x, y) + \alpha S(t) \tag{2}$$

where μ , ν and α are parameters to be estimated. The model in (2) does not include any information from the BI and serves as a model against which we can compare models incorporating BI information from each station. We take

$$\lambda(t, x, y) = \nu m(x, y) + \alpha S(t) + B(t, x, y) \tag{3}$$

as our “BI model” and inspect the usefulness of the BI by comparing the performance of both λ and λ^0 .

4.1 Parameter Estimation

All parameters for each of the models were estimated by maximizing the log-likelihood function

$$\ell(\boldsymbol{\theta}) = \sum_{i=1}^n \log \lambda(t_i, x_i, y_i; \boldsymbol{\theta}) - \int_{T_1}^{T_2} \iint_{\mathcal{S}} \lambda(t, x, y; \boldsymbol{\theta}) dx dy dt$$

where $\boldsymbol{\theta}$ is a vector of free parameters and n is the total number of events (t_i, x_i, y_i) in the dataset, observed in the time interval $[T_1, T_2]$ over the area \mathcal{S} . The parameters to be estimated are ν , α , ϕ_x , ϕ_y , and ϕ_{seas} , as well as the BI parameters γ_s and β_s ($s = 1, \dots, 8$). Under fairly general conditions, maximum likelihood estimates (MLE) have been shown to be consistent and asymptotically normal (Ogata, 1978; Rathbun and Cressie, 1994; Rathbun, 1996).

When optimizing the log-likelihood some restrictions must be placed on the parameters in order to maintain a positive conditional intensity function and numerical stability of the optimization procedure. We restricted the parameters in each of the models to be positive. In addition, the bandwidth parameters in the spatial and seasonal components were bounded away from zero. The inclusion of the β_s parameters in the BI component increased the complexity of the likelihood surface considerably and created some difficulty with the numerical optimization. Ogata et al. (1982) and Ogata and Akaike (1982) handled a similar problem with a single bandwidth parameter by restricting that parameter to a finite grid and repeating the maximum likelihood procedure for each value of the bandwidth parameter on the grid. Unfortunately for our situation, with 8 separate parameters (one for each weather station), constructing a reasonable grid over which to optimize the log-likelihood was computationally infeasible. Rather, we chose to restrict the β_s parameters to be less than 3.0 spatial units (about 56 miles). This upper limit seemed reasonable in the sense that a particular weather station should not have influence over points more than 50 miles from the station (see e.g. Haines et al., 1983). An alternative modelling approach which could circumvent some of the problems mentioned above would be to model $\log \lambda(t, x, y)$ and use pseudo-likelihood

methods, though this approach would not likely solve the problems associated with the inclusion of the BI parameters.

4.2 Residual Analysis and Approximate Random Thinning

Various methods for constructing a multi-dimensional residual point process have been proposed based on random rescaling (Merzbach and Nualart, 1986; Nair, 1990; Schoenberg, 1999) and random thinning (Schoenberg, 2003). One-dimensional residual analysis via the rescaling method has been successfully applied in a wide variety of applications (e.g. Berman, 1983; Ogata, 1988; Diggle, 1990; Rathbun, 1993; Brown et al., 2001). However, rescaling can be awkward to use for multi-dimensional residual analysis. In practice, when the points are rescaled the domain of observation is also rescaled and can become uninterpretable or irregular (see Schoenberg, 1997, for examples). Schoenberg (2003) proposed a method based on approximate random thinning of the observed points which has the advantage that the resulting residual process lies in the same domain as the observed point process.

The algorithm for approximate random thinning is straightforward and easy to implement:

1. Choose a positive integer K such that $K < n$, where n is the total number of points observed.
2. For $i = 1, \dots, n$, compute $p_i = \lambda(t_i, x_i, y_i; \hat{\theta})^{-1} / \sum_{j=1}^n \lambda(t_j, x_j, y_j; \hat{\theta})^{-1}$ where $\lambda(t, x, y; \hat{\theta})$ is the estimated conditional intensity function.
3. Using probability weights p_1, \dots, p_n , take a subsample of size K from the original points $\{(t_i, x_i, y_i)\}$ to produce the residual process $\{(t_j^*, x_j^*, y_j^*)\}$ ($j = 1, \dots, K$).

The algorithm attempts to “thin out” points in areas with high intensity and retain points in areas of low intensity. Although points are deleted from the original dataset, one can repeat the algorithm many times to produce multiple random realizations of approximate thinned residuals.

If K is chosen to be relatively small compared to n and the fitted conditional intensity approximates closely the true conditional intensity governing the point process, then the residual process

should resemble a homogeneous Poisson process with rate $K/(\|\mathcal{S}\|(T_2 - T_1))$ over the original domain of the process. The primary advantage of generating a residual process is that the problem of evaluating the fit of a possibly complex model is reduced to examining whether the residual process is similar to a homogeneous Poisson process, a task for which there are many tests and diagnostics. Once the residuals have been produced one can simply display them in residual plots or compute summary statistics. For example, one may wish to test for residual clustering or inhibition via a statistic such as the K -function (Ripley, 1981). In Section 5.1.1 we use a spatial-temporal version of the K -function to quantify the clustering in the approximate thinned residual processes.

5 Application to Wildfire and Burning Index Data

Each of the models in (1)–(3) were fit to the Los Angeles County data by maximum likelihood. The MLE’s for the parameters in the spatial and seasonal components of each of the models are shown in Table 2. The estimate for $\hat{\mu}$ in the homogeneous Poisson model was 0.0041 events/(spatial unit² \times day), indicating an average of 24 fires per year in the County. The bandwidth parameters for the spatial background component of the Spatial+Seasonal model were estimated as $\hat{\phi}_x = 0.0440$ spatial units (0.83 miles) in the x direction (east-west) and $\hat{\phi}_y = 0.0259$ spatial units (0.49 miles) in the y direction (north-south) for the Spatial+Seasonal model. For the best-fitting BI model, the estimates for the spatial bandwidth parameters were similar, although somewhat smaller. The estimate for the bandwidth parameter of the seasonal component was approximately 8 days for both models. The estimates of the station multiplier coefficients ($\gamma_s, s = 1, \dots, 8$) and the β_s values in the BI component are shown in Table 3. Many of the coefficients are estimated to be zero, likely resulting from the high correlation of the BI values between different stations. The stations receiving non-zero weight are Stations 1, 4, 5, and 7.

To compare the overall fit of each of the models we used the Akaike Information Criterion (Akaike, 1973), or AIC, defined as $-2\ell(\hat{\boldsymbol{\theta}}) + 2p$, where $\ell(\hat{\boldsymbol{\theta}})$ is the log-likelihood of the model evaluated at the MLE $\hat{\boldsymbol{\theta}}$ and p is the number of free parameters in the model. Not unexpectedly, there is a

dramatic decrease in AIC from 7693.9 for the homogeneous Poisson model to 6741.1 for the Spatial+Seasonal model. The addition of the weather stations in the BI model decreases the AIC to 6704.6. The decrease in AIC between the Spatial+Seasonal and BI model indicates that the BI component is in fact improving the fit of the model, even with the addition of 16 parameters (2 for each station). However, the relative decrease in AIC is considerably smaller than the decrease between the homogeneous Poisson and the Spatial+Seasonal models, suggesting that the impact of the BI component is very subtle. One can obtain a sense of the individual contributions of the components by examining a BI-only model (no seasonal or spatial components), a Seasonal+BI model (no spatial component), and a Spatial+BI model (no seasonal component). The AIC values for these models were 7358.2, 7062.7, and 6857.9, respectively. While the BI component alone captures some seasonal and spatial variability, much can be gained by adding separate seasonal and spatial components. Furthermore, none of the smaller BI models fits as well as the Spatial+Seasonal model.

One can test the significance of the added BI component with a likelihood ratio test. Since the additive BI and Spatial+Seasonal models are nested, the test statistic has an asymptotic χ_k^2 distribution under the null hypothesis that the BI model provides no improvement (Ogata, 1988). For these two models, the difference in number of parameters is 16 and the log-likelihood ratio test statistic has a value of 35.2, which is significant at the 5% level (p -value of 0.0037).

For comparison, we also fit models where the different components are multiplied instead of added. The results were similar in that the model incorporating BI did not appear to provide a dramatic improvement in fit to the data. In this case, the BI model produced a conditional intensity that was quite smooth and did not accurately represent the spatial-temporal clustering in the data. The values of AIC for the Spatial \times Seasonal model and the full (multiplicative) BI model were 6831.6 and 6680.5, respectively. It should be noted that while the multiplicative full BI model has a slightly lower AIC than the additive version, differences in AIC between non-nested models should be interpreted with caution. In particular, those differences can be subject to substantial sampling fluctuations (Ripley, 1996; Stone, 1977).

Figure 4 shows the estimated conditional intensity function for the additive BI model on the 15th of each month in 1999. The year 1999 is a typical year in the dataset, containing a total of 21 fires. For this year the intensity reaches its lowest point around March and increases through August. The conditional intensity is generally high in the northwest region of the County where much of the wildfire activity takes place.

5.1 Residual Analysis

While AIC is useful for determining the relative improvement of fit for competing models, one may be interested in a more refined analysis of a particular model. Residual analysis of the Spatial+Seasonal and BI models was conducted using both approximate random thinning and the rescaling method in order to identify possible departures of the models from the data.

5.1.1 Approximate Thinned Residuals

For the approximate random thinning procedure we chose a subsample size of $K = 50$ for each thinning and generated 1000 thinnings from each model. In each realization the residuals appear to be spread uniformly across the County. Recall that in Figure 1(a) the data were highly clustered in the northwest region near weather stations 1 and 2 and there were relatively few points in the northeast corner. This clustering is not apparent in the thinned residuals, indicating adequate treatment of this clustering effect in the BI model.

Figure 5 shows two realizations of thinned residuals for the time and y -coordinates. Here we see that there is a cluster of points around the years 1979–1981 (other realizations demonstrated a similar pattern). This residual non-stationarity indicates that the BI model is not adequately taking into account the non-stationarity in the data around this specific time period. One possible explanation for this lack of fit is that the increased activity was not due to purely meteorological phenomena or changes in fuel properties. Indeed, the BI model for those years appears close to the Spatial+Seasonal model (this is discussed further in Section 5.1.2). Some of the wildfires observed during this period occur in the far northeast region mentioned in Section 4. However, in the three

year period of 1979–1981, only 8 wildfires occur in this area, while the total numbers of fires for these years are 56, 57, and 44 — double the average number of fires per year. Therefore, it is difficult to attribute the residual non-stationarity to a change in spatial distribution alone. The thinned residuals appear to indicate a failure in the BI to detect conditions associated with the prevalence of fires between 1979 and 1981.

While visual inspection of the residuals can be a useful method of model evaluation, it may be desirable to have a more systematic test available. Existing second-order methods for analyzing point patterns are largely two-dimensional, although there have been some extensions (e.g. Baddeley et al., 1993; Diggle et al., 1995). In order to test the homogeneity of the residual process, we used a simple space-time version of the K -function. The general K -function evaluated at distance h is the expected number of pairs of points per unit area that are within distance h of each other, i.e.

$$K(h) = \frac{1}{n} \frac{\|\mathcal{S}\|}{n} \sum_{i \neq j} \mathbf{1}\{d(x_i, x_j) < h\}$$

where x_i, x_j are points of the process, \mathcal{S} is the domain of observation, and d is a distance function. In order to evaluate the K -function, a distance function must be specified, which in purely spatial settings is typically Euclidean distance. For our application, we chose the following distance function, which is defined for two points (t_1, x_1, y_1) and (t_2, x_2, y_2) as

$$d\{(t_1, x_1, y_1), (t_2, x_2, y_2)\} = \sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2} + \delta|t_1 - t_2|. \quad (4)$$

Given a point x_i and distance h , we count the number of points x_j in the cone of radius h and height h (centered at x_i) and then average over all points in the pattern.

The value of δ in (4) was chosen so that the temporal and spatial scales were commensurate. Here this corresponds to $\delta = 1/5475$ days, which sets a spatial distance of 5 miles roughly equivalent to a temporal distance of 4 years. Previous research in Los Angeles County has suggested that the occurrence of a wildfire tends to inhibit the occurrence of another wildfire in the same location. The risk of fire then slowly increases for approximately 20 years, after which the risk of fire is nearly constant (Peng and Schoenberg, 2002; Schoenberg et al., 2003). Therefore, fires occurring within

4-5 years of each other (in the same location) would be considered “nearby” in a similar sense that fires occurring within 5 miles of each other (at the same time) would be considered “nearby”. The precise choice of δ does not affect the actual computation of the K -function, which is invariant to rescalings of the data. Finally, rather than plot the raw K -function, we use a normalized version (sometimes called the L -function) which is centered around zero for a homogeneous Poisson process.

Figure 6 shows the mean estimated K -functions for the 1000 approximate thinned residuals from both the Spatial+Seasonal model and the BI model. The K -function does not enter the negative range in this case, indicating a lack of inhibition in the residuals. Therefore, to show greater detail, the figure omits the negative range of the K -function. The estimated K -function for residuals of the BI model appears to decrease to zero faster than that of the Spatial+Seasonal model. However, there is significant clustering for smaller distances from 0.25 to 1.0, which corresponds to a range of 4–18.9 miles in the spatial domain and 4–15 years in the temporal domain. As shown in the following Section, the clustering we observed in the residuals is most likely due to a multi-year period of increased wildfire activity that is not being captured by the BI model.

5.1.2 Rescaled Temporal Residuals

The fit of the BI model in the temporal domain can be assessed using the rescaling method (Meyer, 1971) to create a residual process on the line. Each original event time t_i ($i = 1, \dots, n$) is mapped to a new time

$$\tau_i = \int_{T_1}^{t_i} \iint_{\mathcal{S}} \lambda(t, x, y; \hat{\boldsymbol{\theta}}) dx dy dt.$$

where T_1 is equal to January 1, 1976 and \mathcal{S} is the spatial observation window. We can then check whether the residuals $\tau_1 < \dots < \tau_n$ appear as a homogeneous Poisson process of rate 1 on the line. A histogram of the rescaled residuals (not shown) appeared to have a larger than expected number of points in the period 50–150 in transformed time. On the original time scale the interval 50–150 corresponds approximately to the years 1978–1982. This result along with the clustering in the thinned residuals (observed in Figure 5) confirms that the increased wildfire activity during

the years 1979–1981 is not captured by the BI model.

One can further examine the original data using the temporal intensity function,

$$r(t) = \iint_S \lambda(t, x, y; \hat{\theta}) dx dy.$$

Figure 7 shows $r(t)$ over two different time intervals for the Spatial+Seasonal and BI model. Figure 7(a) shows $r(t)$ for the years 1979–1981, the period in which we observe an unusually high level of wildfire activity. One might expect the BI model to have a higher intensity during this time period, given the dependence of the BI on local weather and fuel properties. However, the BI model appears to be quite close to the Spatial+Seasonal model and significantly underestimates the rate of activity during these three years. Figure 7(b) shows the period 1990–1991, a period which contained a larger than expected residual interevent time observed in a plot of the empirical log-survivor function. The two vertical dotted lines indicate the two events which generated the large residual interevent time. These two events correspond to the last fire in 1990 and the first fire in 1991. The first fire in 1991 comes on August 23rd, which is much later in the year than is typical for the first fire. Figure 7(b) shows that the conditional intensity of the BI model is much lower than that of the Spatial+Seasonal model in the months between May and September and is possibly reflecting a local change in weather or fuel conditions which is different from the usual seasonal pattern. However, it would appear that the BI model is not compensating enough, thus creating the larger than expected residual interevent time.

5.2 Model Predictions

Given a model for the conditional intensity of a point process, the process can be simulated via the random thinning algorithm of Lewis and Shedler (1979). We simulated one year’s worth of events to see if features of the simulations matched those of the observed events. The BI model was re-fit using the wildfire and BI data from 1976 through 1998 and the Lewis-Shedler algorithm was applied to generate random realizations of wildfire events for 1999. The spatial distribution of the simulations appeared to match the configuration of the observed fires fairly well. However,

we found that the BI model tends to predict more fires during the period between January and April than were actually observed. The first observed fire of 1999 was on January 3rd followed by a fire on April 23rd. However, in each simulation the BI model predicted, on average, 7 fires in the intervening months. In the entire 23 year period of 1976–1998, the average number of fires between January and April was less than 1. Figure 3 shows that on average many of the stations do not reach their lowest point until the middle of March or even April. Therefore, the BI model will produce a higher intensity because of the high BI values. This indicates a failure in the BI to characterize adequately the low risk of fire associated with the meteorological and vegetative conditions during the winter and early spring.

6 Summary and Discussion

In this article we have developed an approach for evaluating a wildfire hazard index using space-time conditional intensity models. This approach has allowed for a detailed analysis of the performance of the Burning Index in predicting wildfire occurrence in Los Angeles County. Our conclusions about the BI are based on an assessment of conditional intensity models which incorporate spatial, seasonal, and BI information. We find that the best-fitting model that incorporates BI information does not perform substantially better than a simple model which only takes into account natural spatial and seasonal variation.

Two point process residual analysis techniques were employed to supplement a standard likelihood based model evaluation criterion (AIC). The random thinning method enabled us to check for residual space-time clustering on the same temporal and spatial scales as the data, while the rescaling method allowed for the closer inspection of temporal clustering in the residuals. Together, the methods provided greater insight into precisely where the BI model fit poorly and where it was making some (minimal) improvement.

It is important to note the possible biases that may result from the missing data and the procedure used to fill in missing BI values. In Section 3 we replaced a missing value on a given day

with the average of the non-missing values for that day across all years. If the non-missing values do not accurately represent the missing data, then the resulting estimated conditional intensity could be biased. In our initial examination of the BI data we found that stations with relatively low percentages of missing data had very regular seasonal patterns from year to year. While one would expect some natural variation between stations, we see no reason why the other stations should not exhibit the same strong seasonal patterns. Therefore, the biases resulting from the missing data are likely to be small. Determining the optimal use of station data, including developing methods for imputing missing values, is an important subject for future work.

Another area for future investigation is the examination of the performance of other hazard indices in Los Angeles County, having already identified some specific deficiencies with the BI. Also, it may be necessary to incorporate other variables into index computations or reexamine the fuel models used to adapt the system to different locations.

While the incorporation of the BI into the models presented here is consistent with the way it is used in practice by Fire Department officials, the usage here is in some disagreement with the original motivation for the development of the BI, which was to predict flame length. With that in mind, it is perhaps not surprising that the utility of the BI in predicting wildfire occurrence is severely limited. One difficulty which applies to the development and usage of any index in Los Angeles County is that wildfire occurrences exhibit very strong seasonal and spatial patterns which already explain much of the variation in the data. Therefore, it will likely be challenging to develop an index for Los Angeles County which does considerably better than a model which contains only spatial and seasonal components.

References

Akaike, H. (1973), “Information Theory and an Extension of the Maximum Likelihood Principle,” in *Second International Symposium on Information Theory*, eds. Petrov, B. N. and Càski, F., Akademiai Kiado, Budapest, pp. 267–281.

- Baddeley, A. J., Moyeed, R. A., Howard, C. V., and Boyde, A. (1993), “Analysis of a Three-dimensional Point Pattern with Replication,” *Applied Statistics*, 42, 641–668.
- Berman, M. (1983), “Comment on ‘Likelihood Analysis of Point Processes and Its Applications to Seismological Data,’ by Y. Ogata,” *Bulletin of the International Statistical Institute*, 50, 412–418.
- Bradshaw, L. S., Deeming, J. E., Burgan, R. E., and Cohen, J. D. (1983), “The 1978 National Fire-Danger Rating System: Technical Documentation,” Tech. Rep. INT-169, USDA Forest Service, Intermountain Forest and Range Experiment Station.
- Brémaud, P. M. (1981), *Point Processes and Queues: Martingale Dynamics*, Springer-Verlag, New York.
- Brown, E. N., Barbieri, R., Ventura, V., Kass, R. E., and Frank, L. M. (2001), “The Time-Rescaling Theorem and its Application to Neural Spike Train Data Analysis,” *Neural Computation*, 14, 325–346.
- Burgan, R. E. (1988), “Revisions to the 1978 National Fire-Danger Rating System,” Tech. Rep. SE-273, USDA Forest Service, Southeastern Forest Experiment Station.
- Daley, D. J. and Vere-Jones, D. (2003), *An Introduction to the Theory of Point Processes: Volume I*, Springer, New York, 2nd ed.
- Deeming, J. E., Burgan, R. E., and Cohen, J. D. (1977), “The National Fire-Danger Rating System — 1978,” Tech. Rep. INT-39, USDA Forest Service, Intermountain Forest and Range Experiment Station.
- Deeming, J. E., Lancaster, J. W., Fosberg, M. A., Furman, R. W., and Schroeder, M. J. (1972), “The National Fire-Danger Rating System,” Tech. Rep. RM-84, USDA Forest Service, Rocky Mountain Forest and Range Experiment Station.
- Diggle, P. J. (1990), “A Point Process Modelling Approach to Raised Incidence of a Rare Phe-

- nomenon in the Vicinity of a Prespecified Point,” *Journal of the Royal Statistical Society, Series A*, 153, 349–362.
- Diggle, P. J., Chetwynd, A. G., Häggkvist, R., and Morris, S. E. (1995), “Second-order Analysis of Space-Time Clustering,” *Statistical Methods in Medical Research*, 4, 124–136.
- Haines, D. A., Main, W. A., Frost, J. S., and Simard, A. J. (1983), “Fire-Danger Rating and Wildfire Occurrence in the Northeastern United States,” *Forest Science*, 29, 679–696.
- Jacod, J. (1975), “Multivariate Point Processes: Predictable Projection, Radon-Nikodym Derivatives, Representation of Martingales,” *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 31, 235–253.
- Keeley, J. E. and Fotheringham, C. J. (2001), “Historic Fire Regime in Southern California Shrublands,” *Conservation Biology*, 15, 1536–1548.
- Keeley, J. E., Fotheringham, C. J., and Morais, M. (1999), “Reexamining fire suppression impacts on brushland fire regimes,” *Science*, 284, 1829–1832.
- Lewis, P. A. W. and Shedler, G. S. (1979), “Simulation of Nonhomogeneous Poisson Processes by Thinning,” *Naval Research Logistics Quarterly*, 26, 403–413.
- Mandallaz, D. and Ye, R. (1997), “Prediction of Forest Fires with Poisson Models,” *Canadian Journal of Forest Research*, 27, 1685–1674.
- Merzbach, E. and Nualart, D. (1986), “A Characterization of the Spatial Poisson Process and Changing Time,” *Annals of Probability*, 14, 1380–1390.
- Meyer, P. (1971), “Demonstration Simplifiée d’un Théorème de Knight,” in *Séminaire de Probabilités V*, Université Strasbourg, vol. 191 of *Lecture Notes in Mathematics*, pp. 191–195.
- Nair, M. G. (1990), “Random Space Change for Multiparameter Point Processes,” *Annals of Probability*, 18, 1222–1231.

- Ogata, Y. (1978), “The Asymptotic Behavior of Maximum Likelihood Estimators for Stationary Point Processes,” *Annals of the Institute of Statistical Mathematics*, 30, 243–261.
- (1988), “Statistical Models for Earthquake Occurrences and Residual Analysis for Point Processes,” *Journal of the American Statistical Association*, 83, 9–27.
- Ogata, Y. and Akaike, H. (1982), “On Linear Intensity Models for Mixed Doubly Stochastic Poisson and Self-exciting Point Processes,” *Journal of the Royal Statistical Society, Series B*, 44, 102–107.
- Ogata, Y., Akaike, H., and Katsura, K. (1982), “The Application of Linear Intensity Models to the Investigation of Causal Relations between a Point Process and Another Stochastic Process,” *Annals of the Institute of Statistical Mathematics*, 34, 373–387.
- Peng, R. D. and Schoenberg, F. P. (2002), “Estimation of the Fire Interval Distribution for Los Angeles County, California,” Tech. Rep. 337, UCLA Department of Statistics.
- Rathbun, S. L. (1993), “Modeling Marked Spatio-temporal Point Patterns,” *Bulletin of the International Statistics Institute*, 55, 379–396.
- (1996), “Asymptotic Properties of the Maximum Likelihood Estimator for Spatio-temporal Point Processes,” *Journal of Statistical Planning and Inference*, 51, 55–74.
- Rathbun, S. L. and Cressie, N. (1994), “Asymptotic Properties of Estimators for the Parameters of Spatial Inhomogeneous Poisson Point Processes,” *Advances in Applied Probability*, 26, 122–154.
- Ripley, B. D. (1981), *Spatial Statistics*, Wiley, New York.
- (1996), *Pattern Recognition and Neural Networks*, Cambridge University Press.
- Rothermel, R. C. (1972), “A Mathematical Model for Predicting Fire Spread in Wildland Fuels,” Tech. Rep. INT-115, USDA Forest Service.
- Schoenberg, F. (1997), “Assessment of Multi-dimensional Point Process Models,” Ph.D. thesis, University of California, Berkeley.

- (1999), “Transforming Spatial Point Processes into Poisson Processes,” *Stochastic Processes and their Applications*, 81(2), 155–164.
- Schoenberg, F. P. (2003), “Multi-dimensional Residual Analysis of Point Process Models for Earthquake Occurrences,” *Journal of the American Statistical Association*, in press.
- Schoenberg, F. P., Brillinger, D. R., and Guttorp, P. M. (2002), “Point Processes, Spatial-Temporal,” in *Encyclopedia of Environmetrics*, eds. El-Shaarawi, A. and Piegorsch, W., Wiley, NY, vol. 3, pp. 1573–1577.
- Schoenberg, F. P., Peng, R., Huang, Z., and Rundel, P. (2003), “Detection of Nonlinearities in the Dependence of Burn Area on Fuel Age and Climatic Variables,” *International Journal of Wildland Fire*, 12, 1–6.
- Stone, M. (1977), “An Asymptotic Equivalence of Choice of Model by Cross-Validation and Akaike’s Criterion,” *Journal of the Royal Statistical Society, Series B*, 39, 44–47.
- Viegas, D. X., Bovio, G., Ferreira, A., Nosenzo, A., and Sol, B. (1999), “Comparative Study of Various Methods of Fire Danger Evaluation in Southern Europe,” *International Journal of Wildland Fire*, 9, 235–246.
- Warren, J. R. and Vance, D. L. (1981), “Remote Automatic Weather Station for Resource and Fire Management Agencies,” Tech. Rep. INT-116, USDA Forest Service, Intermountain Forest and Range Experiment Station.
- Westerling, A. L., Cayan, D. R., Gershunov, A., Dettinger, M. D., and Brown, T. (2000), “Western Wildfire Season Forecast,” http://meteora.ucsd.edu/cap/fire_forecast.html.

A Tables

Station #	1	2	3	4	5	6	7	8
Jan–Apr	99.8	49.2	34.3	60.8	33.3	32.5	23.3	26.1
May–Dec	46.4	36.3	17.8	32.0	14.5	18.6	6.7	14.2
Overall	64.0	40.5	23.2	41.5	20.7	23.2	12.2	18.1

Table 1: Percentage of missing values for each weather station during the off-season (Jan–Apr), the fire season (May–Dec), and overall.

Model	$\hat{\mu}$	$\hat{\nu}$	$\hat{\phi}_x$	$\hat{\phi}_y$	$\hat{\alpha}$	$\hat{\phi}_{\text{seas}}$
H. Poisson	0.0041					
Spatial+Seasonal		0.0348	0.0440	0.0259	0.7053	8.44
BI model		0.0293	0.0339	0.0200	0.5570	8.27

Table 2: Maximum likelihood estimates of parameters for non-BI components. “H. Poisson” refers to the homogeneous Poisson model in (1).

	$s = 1$	2	3	4	5	6	7	8
$\hat{\gamma}_s (\times 10^{-5})$	2.2	0.0	0.0	17.4	3.1	0.0	25.6	0.0
$\hat{\beta}_s$	0.001	3.000	3.000	0.387	0.001	0.442	0.520	0.816

Table 3: Maximum likelihood estimates for parameters in the BI component.

B Figure Captions

1. (a) Spatial distribution of wildfires larger than 10 acres in Los Angeles County (1976–2000).
Locations of the 8 Remote Automatic Weather Stations in Los Angeles County are indicated by the numbers above the solid circles. One spatial unit is approximately 18.9 miles; (b) Locations of wildfires occurring in the dataset before 1976.
2. Dates of occurrence and areas burned for fires larger than 10 acres (1976–2000).
3. Average yearly BI pattern.
4. Estimated conditional intensity function for the BI model on the 15th of each month in 1999.
The units are in events / (spatial unit² × day) where one spatial unit is approximately 18.9 miles.
5. Two random realizations of residuals for the BI model (time and y -coordinate).
6. Normalized K -function for the Spatial+Seasonal (solid gray line) and BI (solid black line) models. The dotted line is the upper 95th percentile (pointwise) for the K -function applied to 1000 realizations of a homogeneous Poisson process.
7. Temporal intensity functions $r(t)$ for the Spatial+Seasonal (black line) and BI model (gray line) for (a) 1979–1981; and (b) 1990–1991.

C Figures

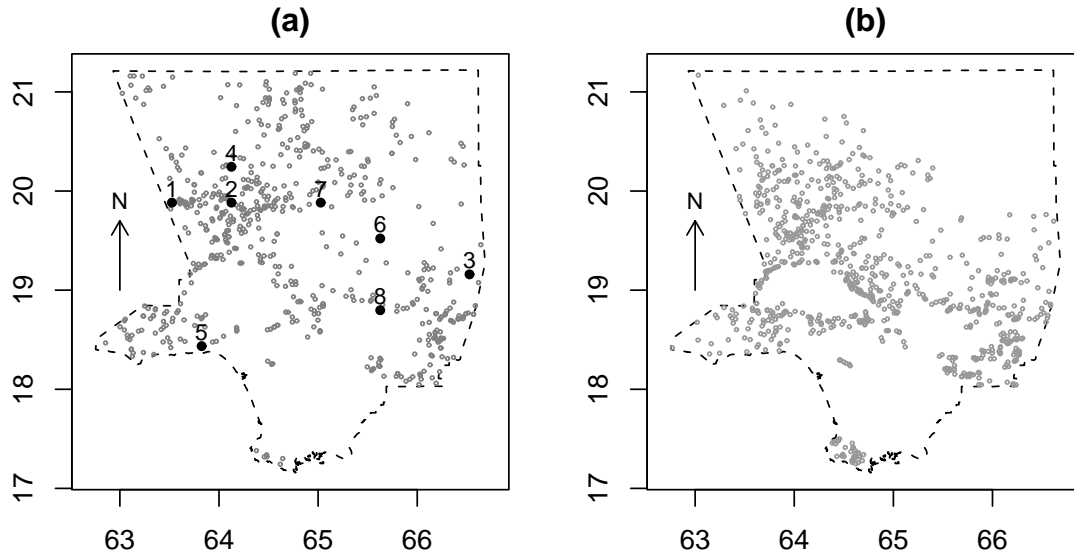


Figure 1: (a) Spatial distribution of wildfires larger than 10 acres in Los Angeles County (1976–2000). Locations of the 8 Remote Automatic Weather Stations in Los Angeles County are indicated by the numbers above the solid circles. One spatial unit is approximately 18.9 miles; (b) Locations of wildfires occurring in the dataset before 1976.




Figure 2: Dates of occurrence and areas burned for fires larger than 10 acres (1976–2000).




Figure 3: Average yearly BI pattern.




Figure 4: Estimated conditional intensity function for the BI model on the 15th of each month in 1999. The units are in events / (spatial unit² × day) where one spatial unit is approximately 18.9 miles.




Figure 5: Residuals for the BI model (time and y -coordinate).




Figure 6: Normalized K -function for the Spatial+Seasonal (solid gray line) and BI (solid black line) models. The dotted line is the upper 95th percentile (pointwise) for the K -function applied to 1000 realizations of a homogeneous Poisson process.




Figure 7: Temporal intensity functions $r(t)$ for the Spatial+Seasonal (black line) and BI model (gray line) for (a) 1979–1981; and (b) 1990–1991.