

ORIGINAL ARTICLE

A non-parametric Hawkes process model of primary and secondary accidents on a UK smart motorway

Kieran Kalair¹ | Colm Connaughton^{2,3} | Pierfrancesco Alaimo Di Loro⁴

¹Centre for Complexity Science, University of Warwick, Coventry, UK

²London Mathematical Laboratory, London, UK

³Mathematics Institute, University of Warwick, Coventry, UK

⁴Department of Statistics, University of Rome "La Sapienza", Rome, Italy

Correspondence

Kieran Kalair, Centre for Complexity Science, Zeeman Building, University of Warwick, Coventry, CV4 7AL, UK.
Email: k.kalair@warwick.ac.uk

Funding information

EPSRC, Grant/Award Number: EP/L015374/1

Abstract

A self-exciting spatiotemporal point process is fitted to incident data from the UK National Traffic Information Service to model the rates of primary and secondary accidents on the M25 motorway in a 12-month period during 2017–2018. This process uses a background component to represent primary accidents, and a self-exciting component to represent secondary accidents. The background consists of periodic daily and weekly components, a spatial component and a long-term trend. The self-exciting components are decaying, unidirectional functions of space and time. These components are determined via kernel smoothing and likelihood estimation. Temporally, the background is stable across seasons with a daily double peak structure reflecting commuting patterns. Spatially, there are two peaks in intensity, one of which becomes more pronounced during the study period. Self-excitation accounts for 6–7% of the data with associated time and length scales around 100 min and 1 km, respectively. In-sample and out-of-sample validation are performed to assess the model fit. When we restrict the data to incidents that resulted in large speed drops on the network, the results remain coherent.

KEYWORDS

Hawkes processes, non-parametric, smart motorways, traffic incidents

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2020 The Authors. Journal of the Royal Statistical Society: Series C (Applied Statistics) published by John Wiley & Sons Ltd on behalf of Royal Statistical Society

1 | INTRODUCTION

The United Kingdom has one of the lowest per capita death rates from traffic accidents in the world, estimated by the World Health Organization (2018) at 3.1 per 100,000 of population in 2016. Nevertheless, 1782 deaths and 25,484 serious injuries resulted from accidents on UK roads in 2018 Department for Transport (2019). Apart from the direct human cost of serious accidents, indirect economic costs result even from relatively minor incidents. This is because crashes, collisions and breakdowns can cause severe congestion leading to significant drops in the efficiency of the road transport network. For these reasons, there is an imperative to further reduce the accident rate on UK roads. However traffic accidents are rare in absolute terms and are not distributed uniformly across the network. Further rate reductions are therefore likely to require targeted interventions.

Targeted interventions might try to improve safety at specific locations where the accident risk is known to be high compared to the baseline or might try to mitigate against particular mechanisms that are known to account for a significant proportion of accidents. Infrastructure modifications to improve the safety of accident-prone junctions are an example of the first type of intervention. The deployment of the Motorway Incident Detection and Automatic Signalling (MIDAS) system to reduce the number of secondary accidents on motorways is an example of the second. Secondary accidents occur when a driver fails to react appropriately to the disruption caused by an existing accident leading to a subsequent accident upstream of the first one. MIDAS uses a network of induction sensors, known as loops, embedded in the road surface to automatically detect queues and then warn upstream drivers of the danger ahead via roadside signage, and is one component of the UK ‘smart motorways’ infrastructure, see Highways England (2020) for further details. Most of the data, including data on accidents and congestion, is publicly available via the National Traffic Information Service (NTIS).

In this paper, we use data from NTIS to model the distribution of motorway incidents as a spatio-temporal process comprised of a background component and a self-excitation component, see Section 4 for details. We focus on the M25 London Orbital, one of the busiest motorways in the United Kingdom. The objectives of the study are twofold. The first is to quantify how accident risk on the M25 varies in space and time relative to the baseline. The second is to use the self-excitation component of the process to quantify the likely contribution of secondary incidents to the observed totals. We would like this model to be helpful in addressing the question of how best to target interventions when the baseline accident rate is low in absolute terms. We therefore perform extensive in-sample and out-of-sample validation to verify the models performance on seen and unseen data.

2 | LITERATURE REVIEW

There is much work focusing on spatial and temporal analysis of traffic incidents, with a discussion of the evidence for spatial autocorrelation in incident data given in Aguero-Valverde and Jovanis (2008). Here, the authors discussed previously used descriptive analysis methods, including K-function analysis and comparison to complete spatially random patterns, and detail how they showed evidence of spatial correlation among event locations. They furthered this by incorporating spatial components into an autoregressive model, finding it improved upon models disregarding the spatial correlation in the data. Additional analysis is provided in Fan et al. (2018), where the authors analysed the evolution of event hot spots through time on an urban network in China. They extended and applied kernel density estimation on networks, Moran's I and Local Indicators of Mobility Associations (LIMA) to offer exploratory analysis of the data set from multiple perspectives. Their work was focused on analysis of the data, rather than formulating a model incorporating the observed features, as we do here.

Further spatial-temporal analysis was completed in Song et al. (2018), where the authors used Kulldorff's space-time scan statistics to determine statistically significant clusters of traffic accidents across the entire United Kingdom in 2016. They found two significant clusters, both in the north of the country, but conceded that they do not explicitly account for the network structure in their analysis. We also consider data from the United Kingdom, but focus on a single motorway rather than the entire country. Our approach is fundamentally different as we aim to model the dynamics as a point process, not just discover locations of statistically significant clusters. Readers interested in the discussed spatial analysis methodologies can find details of K-function analysis in Dixon (2014), Kulldorff's space-time scan statistic in Kulldorff (2001), and LIMA in Rey (2016).

A point of particular interest to our work arises from the modelling of traffic incidents in Acker and Yuan (2019). Here, logistic regression and random forest models were used to predict the likelihood of event occurrences. In particular, the models incorporated a significant 'cascading effect' variable, in which the presence of an event showed significant influence on the likelihood of another nearby in space and time. Although a different type of model and formulation to the one we use, there is clearly a sense that cascading effects are a real component in some traffic data that one may want to incorporate into a model. It is unclear what time and length scales are associated with this effect however, and if smart motorway features may remove this effect from the data.

Based on this analysis, it is natural to consider applying spatiotemporal point-process models to traffic events. Recent work considering these models on linear networks is given in Moradi and Mateu (2019), with an application to traffic event modelling. Considering road networks in Huston, Medellin and Eastbourne, the authors investigated what features the network-constrained data showed. They found statically significant evidence that events on the network did not follow a uniform spatial-temporal Poisson process and that tests indicated favouring clustering of data in space and time. While their focus was on urban road networks, ours is on a smart motorway. As we consider one continuous ring road, not an urban network, we can also use a standard distance measure, rather than a graph distance based one as in the cited work. Our choice of data set is not to simplify the fitting procedure onto continuous space rather than a network, but instead it offers some insight into a much discussed topic of smart motorway safety.

From our review of the literature, it is clear that one could ask if point-process models incorporating a self-excitation component capture this previously discussed cascading effect. Such models have been applied to a wide range of real-world problems, with a recent review of these models given in Reinhart (2018). One application discussed in this review is earthquake modelling, as in Zhuang et al. (2004) and Zhuang (2011). This is a natural application for such models, as there is strong evidence that initial large earthquakes lead to aftershocks, and hence there is a clearly interpretable 'self-excitation' component to the application. Another application discussed is crime forecasting, as in Mohler et al. (2011) and Mohler (2014), where self-excitation can be seen in physical terms as retaliation crimes, among other things. Alternative applications discussed include epidemic forecasting in Schoenberg et al. (2019), and modelling events on social networks as in Fox et al. (2016). Very recently, similar models have been applied to modelling the spread of COVID-19 Lesage (2020), although this is still in its infancy.

Clearly, there is significant work on applying self-exciting point-process models to problems in crime, earthquakes and epidemics, but there is little on applying them to traffic events. One paper that does look at this is Li et al. (2018), where a self-exciting point-process model is proposed that could theoretically be fit to real data. However, only simulated data are used, generated from the proposed model, to then show how the fitting and evaluation would work. Additionally, a somewhat similar idea was considered in Lim et al. (2016); however importantly, here the authors considered traffic flow data to be 'events' and tried to use self-excitation to model the idea that often traffic flow occurs in

clusters. They then applied the methods to model traffic flow in Sydney; however, there was no clear conclusion as to if the model was statistically defensible and captured all features of the data, or if alternative traffic forecasting methods were preferable. There is still an enormous amount of work to be done applying this methodology to real data, and understanding what components of it are important, and the amount of self-excitation present in traffic data, along with appropriate time and length scales it occurs on. Throughout our modelling, we are able to offer some discussions of these questions.

In the discussed works using self-exciting point processes, there is some sense of a ‘background’ component that models the typical behaviour and a ‘triggering’ component that allows for self-excitation. There is much discussion as to what functional form the components should take, in particular for the triggering components. Typically, one supposes some reasonable parametric forms, then determines which is most appropriate through inspection of the log-likelihood value or information criteria. Standard choices include some form of exponential, Gaussian or power-law decay of triggering in time and space. However, recent work in Zhuang and Mateu (2019) showed how one could determine both the background and triggering components in a non-parametric way, through kernel-smoothing of data. In particular, the authors proposed to model crime data using a background comprised of periodic daily, periodic weekly, long-term trend and spatial components, as well as triggering in space and time. However, every component is determined without assuming any functional form, instead the authors show how when basing their methods on work in Zhuang (2006), one can determine which events in the data appear to be background and which appear to be triggered, and then smooth data based on this to reconstruct the desired components. It is on this we base most of our work in this paper, and in Section 4, we give an overview of the methodology and how we adapt it for our use.

3 | DATA COLLECTION AND PREPROCESSING

3.1 | Collection & selection

The data for this study are taken from the National Traffic Information Service (NTIS)[†]. NTIS provides both historic and real-time traffic data for all roads in the United Kingdom that lie on the ‘Strategic Road Network’. This network includes all motorways and major A-roads in the United Kingdom but we choose to focus our analysis on one the countries busiest motorways, the M25 London Orbital, pictured in Figure 1. Inside NTIS, roads are represented by a directed graph, with edges (referred to as links from now on) being segments of road that have a constant number of lanes and no slip roads joining or leaving. We extract all links that lie on the M25 in the clockwise direction, yielding a subset of the road network to collect data for. Placed along links are physical sensors called ‘loops’, which record passing vehicles and report averaged quantities each minute.

The most relevant components of NTIS to our work are event flags that are manually entered by traffic operators. These flags specify an event type, for example accident or obstruction, the start and end time and the link the event occurred on. We extract all accidents and obstructions that occurred on the chosen links between 1 September 2017 and 30 September 2018. However, we require more fine-grained locations for events than just the link it occurred on, so we perform further localization. This preprocessing step is highly specific to our data set, so we only offer an overview here, with more details given in the supplementary material. To understand why this is needed, one should note that links vary in size from around 500 to 10,000 m. However, the average distance between successive

[†]Technical details of the NTIS data feeds are available at <http://www.trafficengland.com/services-info>



FIGURE 1 The M25 London Orbital, roughly 180 km in length. The Dartford crossing, located in the east, is a short segment of road that we do not have data for [Colour figure can be viewed at wileyonlinelibrary.com]

loop sensors is roughly 500 m. As a result, the data are not initially appropriate to model with point processes; however, one can use the time series provided by individual loop sensors to determine higher resolution location data for events. Given an event flag and link, we localize the event in between the two loop sensors on that link that show the largest drop in speed and rise in occupancy when going from the sensor downstream to upstream. This is motivated by existing work in incident detection, with an example being Payne et al. (1976).

4 | METHODOLOGY

4.1 | Model formulation

Our objective is to model the number of events (incidents) observed between any two positions x_1 and x_2 in any time interval between t_1 and t_2 . We denote this quantity by $N_{[x_1, x_2], [t_1, t_2]}$. Since we know that accidents cluster in both space and time, the simplest model is a non-homogeneous Poisson point process. This model is specified by an underlying intensity function, $\lambda(x, t)$, which is the local event probability per unit length per unit time. It is then assumed that for any intervals, $[x_1, x_2]$ and $[t_1, t_2]$, $N_{[x_1, x_2], [t_1, t_2]}$ has a Poisson distribution:

$$\mathbb{P}(N_{[x_1, x_2], [t_1, t_2]} = n) = \frac{\Lambda_{[x_1, x_2], [t_1, t_2]}^n}{n!} \exp(-\Lambda_{[x_1, x_2], [t_1, t_2]}), \quad (1)$$

where

$$\Lambda_{[x_1, x_2], [t_1, t_2]} = \int_{t_1}^{t_2} \int_{x_1}^{x_2} \lambda(x, t) dx dt. \quad (2)$$

It is natural to assume that the intensity is multiplicatively decomposable:

$$\lambda(t,x) = \mu_0 \mu_d(t) \mu_w(t) \mu_t(t) \mu_s(x). \quad (3)$$

Here μ_0 is a uniform background intensity which has units of events per unit time per unit length. This uniform background is then modulated by the functions $\mu_d(t)$, $\mu_w(t)$, $\mu_t(t)$ and $\mu_s(x)$ to capture spatiotemporal variation. The spatial modulation, $\mu_s(x)$, accounts for the fact that different locations will naturally have differing rates of accidents. For example, junctions with low visibility having higher rates than straight, simple sections of road. The temporal modulation, $\mu_d(t) \mu_w(t) \mu_t(t)$ consists of three components: $\mu_d(t)$ represents daily variation, $\mu_w(t)$ represents weekly variation and $\mu_t(t)$ represents any long-term trend that may be present. Daily and weekly seasonality is a ubiquitous feature of traffic data reflecting daily ‘rush-hour’ commuting patterns, and weekly differences between the 5-day working week and the 2-day weekend. Annual seasonality may also be present but since our data spans 1 year, any such variation is captured by the trend.

Inhomogeneous point processes describe clustering solely in terms of variations in the intensity function. A more sophisticated model is a self-exciting point process, known as a Hawkes process, to capture the distinction between primary and secondary incidents. Self-excitation means that when an event occurs, the probability of observing a subsequent event nearby increases. A second term depending on the previous events is added to Equation (3) to give what is called the *conditional* intensity function:

$$\lambda(t,x) = \mu_0 \mu_d(t) \mu_w(t) \mu_t(t) \mu_s(x) + A \sum_{\substack{t_i < t \\ x_i > x}} g(t-t_i)h(x-x_i), \quad (4)$$

where A is the triggering rate, g and h are triggering functions that describe how the triggering mechanism decays in time and space, respectively, and (t_i, x_i) are the times and locations of the observed incidents. The word ‘conditional’ here reflects the dependence of the intensity on the realization of the process.

Our methodology is based on the work of Zhuang and Mateu (2019), where it is shown how to construct the conditional intensity of a Hawkes process in a non-parametric way by applying kernel smoothing to the observed data on crime. We modify this model to make it applicable to traffic data. The supplementary material contains the model derivation, estimators of model components and details of the fitting algorithm, following closely Zhuang and Mateu (2019). There are three main changes. First, the spatial triggering mechanism is one-dimensional and unidirectional: secondary incidents cannot occur downstream of the primary incident. Second, we enforce monotonicity of the triggering functions, g and h , in Equation (4) as a constraint to help with identifiability. Third, we apply boundary correction to the kernel density estimates to reduce bias. In the remainder of this section, further detail is provided on these changes.

4.2 | Reconstructing one-dimensional and unidirectional spatial triggering

First, we adapt the work in Zhuang and Mateu (2019) to reconstruct a one-dimensional and unidirectional spatial triggering function. To begin, we state that for a spatiotemporal point process N_p with conditional intensity function $\lambda(t, x)$ and a predictable process $f(t, x) \geq 0$, over a time domain $[T_1, T_2]$ and space domain S , we can write:

$$\mathbb{E} \left[\int_{[T_1, T_2] \times S} f(t, x) dN_p(dt \times dx) \right] = \mathbb{E} \left[\int_{T_1}^{T_2} \int_S f(t, x) \lambda(t, x) dx dt \right]. \quad (5)$$

To determine our triggering functions in an entirely data-driven way, we first consider two data points $(\tau^{(1)}, \chi^{(1)})$ and $(\tau^{(2)}, \chi^{(2)})$ where $\tau^{(1)} < \tau^{(2)}$ and $\chi^{(1)} > \chi^{(2)}$. We then define $\rho(\tau^{(1)}, \chi^{(1)}, \tau^{(2)}, \chi^{(2)})$ as:

$$\rho(\tau^{(1)}, \chi^{(1)}, \tau^{(2)}, \chi^{(2)}) = \begin{cases} \frac{Ag(\tau^{(2)} - \tau^{(1)})h(\chi^{(2)} - \chi^{(1)})}{\lambda(\tau^{(2)}, \chi^{(2)})}, & \text{if } \tau^{(1)} < \tau^{(2)} \text{ and } \chi^{(1)} > \chi^{(2)}. \\ 0, & \text{otherwise} \end{cases} \quad (6)$$

If we apply this in Equation (5), letting $f(\tau, \chi) = \rho(t_i, x_i, \tau, \chi) \mathbb{1}_{\chi - x_i \in [x - \Delta x, x + \Delta x]}$ for some small positive Δx , on the temporal domain $[0, T]$ and spatial domain $[0, X]$, we attain:

$$\begin{aligned} & \sum_j \rho(t_i, x_i, \tau_j, \chi_j) \mathbb{1}_{\chi - x_i \in [x - \Delta x, x + \Delta x]} \\ & \approx \int_0^T \int_0^X \rho(t_i, x_i, \tau, \chi) \mathbb{1}_{\chi - x_i \in [x - \Delta x, x + \Delta x]} \lambda(\tau, \chi) d\chi d\tau \\ & = A \int_{t_i}^T \int_0^{x_i} \frac{g(\tau - t_i)h(\chi - x_i)}{\lambda(\tau, \chi)} \mathbb{1}_{\chi - x_i \in [x - \Delta x, x + \Delta x]} \lambda(\tau, \chi) d\chi d\tau \\ & = A \left(\int_{t_i}^T g(\tau - t_i) d\tau \right) \left(\int_0^{x_i} h(\chi - x_i) \mathbb{1}_{\chi - x_i \in [x - \Delta x, x + \Delta x]} d\chi \right). \end{aligned} \quad (7)$$

If we then let $s = \chi - x_i$, we have:

$$\begin{aligned} & \sum_j \rho(t_i, x_i, \tau_j, \chi_j) \mathbb{1}_{\chi_j - x_i \in [x - \Delta x, x + \Delta x]} \\ & \approx A \left(\int_{t_i}^T g(\tau - t_i) d\tau \right) \left(\int_{-x_i}^0 h(s) \mathbb{1}_{s \in [x - \Delta x, x + \Delta x]} ds \right) \\ & = A \left(\int_{t_i}^T g(\tau - t_i) d\tau \right) \left(\int_{x - \Delta x}^{x + \Delta x} h(s) ds \right) \\ & \propto h(x). \end{aligned} \quad (8)$$

and hence:

$$\hat{h}(x) \propto \sum_{(i,j)} \rho_{i,j} \mathbb{1}_{x_j - x_i \in [x - \Delta x, x + \Delta x]} \quad (9)$$

with:

$$\rho_{i,j} = \frac{Ag(t_j - t_i)h(x_j - x_i)}{\lambda(t_j, x_j)}, \text{ for all } (i,j) \text{ s.t. } t_i < t_j \text{ and } x_i > x_j. \quad (10)$$

This derivation follows closely that in Zhuang and Mateu (2019), with an additional restriction on $\rho_{i,j}$, and h being a univariate function. One then smooths this initial estimate using kernel functions, with more details on this and all other model components in the supplementary material.

4.3 | Constraining triggering functions

We note that this model aims to explain any residuals from the background process using the triggering component. As such, the model has significant freedom to adapt to the data. We can limit this

freedom to ensure that the triggering component truly reflects increased rates of events on a short time scale (compared to the periodic components) in the wake of a particular event. A natural way to do this and extension of the original methodology is to enforce the triggering functions to be monotonic. This is a reasonable constraint that ensures interpretability of the triggering functions, while still allowing them to be constructed by the data. Triggering decaying in space and time suggests we expect reduced influence of an event as we move further from its location and as time passes. To accomplish this, we follow the method in Hall and Huang (2001). In standard kernel smoothing, one writes a smoothed estimate of some data set (X, Y) as:

$$\hat{v}(x) = \frac{1}{N} \sum_{i=1}^N \frac{1}{h} K\left(\frac{x-X_i}{h}\right) Y_i \quad (11)$$

with some kernel K , some bandwidth h , and attain a smoothed estimate $\hat{v}(x)$. Specifically, in our application, X represents differences in event locations or event times, and all Y values equal 1; however, the technique works for more general cases. Hall and Huang (2001) generalized Equation (11) to incorporate a weight p_i into each data point used in the smoothing to ‘adjust’ the initial smoothed fit to be monotonic.

One writes this adjusted fit as:

$$\hat{v}_{mono}(x|p_1, \dots, p_N) = \frac{1}{N} \sum_{i=1}^N \frac{1}{h} K\left(\frac{x-X_i}{h}\right) Y_i p_i. \quad (12)$$

Our goal is now to choose a weight p_i for each data point i used in the construction of the function, while altering the original estimate as little as possible. There are an infinite number of sets of $\{p_1, p_2, \dots, p_N\}$ one could choose to enforce a monotonic function, and to identify a unique solution, we choose the set that is as close to the uniform distribution $\left\{\frac{1}{N}, \dots, \frac{1}{N}\right\}$ as possible. One distance measure used in the work to compare the p_i 's to a uniform distribution is:

$$D_0(p_1, \dots, p_N) = - \sum_{i=1}^N \log(Np_i). \quad (13)$$

Using this, we can then introduce a step in our model fitting where we solve a further optimization problem, determining each p_i value to produce a monotonic triggering function. The optimization problem is specified as:

$$\begin{aligned} & \min_{p_1, \dots, p_N} D_0(p_1, \dots, p_N) \\ & \text{s.t. } \frac{d\hat{v}_{mono}(x|p_1, \dots, p_N)}{dx} \leq \epsilon \\ & p_i \geq 0 \text{ for all } i \in [1, 2, \dots, N] \\ & p_i \leq 1 \text{ for all } i \in [1, 2, \dots, N] \\ & \sum_{i=1}^N p_i = 1. \end{aligned} \quad (14)$$

Since we are enforcing the triggering to decay with increasing time and distance, ϵ constrains the gradient of the triggering functions to be below some value. It should be noted that in practice, constraining the triggering function only alters the resulting functions in Section 5.5, where we fit the model to small subsets of data. When fit to the entire year, the resulting functions are already monotonic.

4.4 | Boundary correction

It is well known that, on a truncated domain, kernel density estimates are biased near the boundary, with discussion of this in Chiu (2000) and references within. There are many ways to account for this throughout the literature, with a simple method being to ‘mirror’ the data, supposing we have extra data points. If we have our ‘real’ data points X_1, X_2, \dots, X_N and wish to truncate our domain at 0, we introduce extra points $-X_1, -X_2, \dots, -X_N$. The smoothed estimate $\hat{v}(x)$, using a bandwidth h , is then:

$$\hat{v}(x) = \frac{1}{Nh} \sum_{i=1}^N \left[K\left(\frac{x-X_i}{h}\right) + K\left(\frac{x+X_i}{h}\right) \right]. \quad (15)$$

We apply this correction to both triggering functions and the background trend.

Other components in our model are not truncated at a specific point, but rather periodic over some domain. This is true for the daily and weekly background components, but we further assume that the spatial background is periodic, as the M25 is an almost continuous ring around London. Recall Fig. 1, a small section of the M25 is not a motorway and reports no data, but it is negligible in comparison to the wider motorway, so assuming a spatial background on a ring is reasonable. These components do not strictly need boundary correction, however care must be taken in their evaluation. Suppose we take the daily background as an example. We observe events times t_i on a domain $[0, T]$ and map these onto some periodic domain $[0, m_d]$, where m_d represents the number of minutes in a day. The mapped times would then be given by $t_i - m_d \left\lfloor \frac{t_i}{m_d} \right\rfloor$. However, an event may have a non-negligible contribution from its location mapped to the previous and following day that is we observe a contribution from times $t_i - m_d \left\lfloor \frac{t_i}{m_d} \right\rfloor - m_d$ and $t_i - m_d \left\lfloor \frac{t_i}{m_d} \right\rfloor + m_d$. As a result, the periodic estimates are the sum over all data points of the kernel evaluated both inside the domain, and at one preceding and following period. This could actually be an infinite sum, over the entire past and future days; however, in practice, the decay of a Gaussian kernel leads to it being sensible to truncate this sum to just these terms. More details of this are provided in the supplementary material.

5 | RESULTS

5.1 | Bandwidth selection, model selection and prevalence of triggering

Selection of smoothing bandwidths is an open problem for models of this form; however, the application discussed offers natural choices with time scales inherent to the system. We choose the daily, weekly and trend bandwidths to be 60, 10×60 and $60 \times 24 \times 14$ min, respectively. The choice of daily bandwidth is selected due to the ‘rush-hour’ behaviour in the United Kingdom typically varying on a time scale of around an hour, while the weekly and trend components capture variation across larger time scales. The spatial bandwidth is chosen as 5500 m, which appears small enough to capture differing features across the M25, while not introducing superfluous oscillations. This is also larger than the uncertainty we would expect in event locations, therefore accounting for potential uncertainty in the data. Finally, the temporal and spatial bandwidths for the triggering functions are chosen to be 30 min and 500 m, respectively. We considered variations on all of these values, finding those listed provided a reasonable compromise of model interpretability and identify known components of traffic flow.

We first consider some measure of goodness of fit for models containing variations of the discussed components. The first is a homogeneous Poisson process, used as the simplest reference model

TABLE 1 Log-likelihood values for models with various components

Model	A	Log-likelihood
Fixed rate poisson process	—	-28861.55
Daily + weekly background	—	-28028.05
Daily + weekly + trend background	—	-27929.60
Daily + weekly + triggering	0.068495	-27864.48
Daily + weekly + trend + triggering	0.065462	-27781.38

one could construct. We then compare models with: daily and weekly background components; daily, weekly and trend background components; daily, weekly background components and triggering; and daily, weekly, trend background components and triggering. To compare these, we consider the log-likelihood, given by:

$$\log(L) = \sum_{i=1}^N \log(\lambda(t_i, x_i)) - \int_0^T \int_0^X \lambda(t, x) dx dt, \quad (16)$$

with a larger value suggesting a better model. Results for this are given in Table 1. Note that in using log-likelihood to judge the goodness of fit, we ignore model complexity, the idea that we could continuously add components to any model and see increasingly marginal improvements as we attain a more complex model. However, our results throughout this section show the estimated functions present reasonable and interpretable behaviour that seems to catch marginal but still present features of the phenomenon of interest. To check the specification of the models we inspect the residuals and validate the results in Section 5.4, while a successful attempt of out-of-sample validation is attained in Section 5.5.

Inspecting Table 1, we see by far the worst model is a constant rate Poisson process, and adding periodic daily and weekly components to this shows the largest improvement in log-likelihood. Including a trend and triggering component further improve the log-likelihood. Using this methodology, the parameter A can be interpreted as the proportion of the impact of the triggering function on the total intensity. For our optimal model incorporating all components, this means about 6.55% of events appear to be the result of triggering, in practical terms about 100 events. One may want to consider this as an upper bound, as discussed in Section 5.3, along with an appropriate time scale.

5.2 | Background analysis

The background component of the model is strong, as seen when inspecting the changes in log-likelihood from Table 1. In our model, this constitutes a daily, weekly, spatial and trend component. We visualize the temporal background components in Figure 2.

Inspecting Figure 2(a), we see that the daily background increases to an initial peak during the morning rush hour, then remains roughly constant, before rising again to a peak at around 4 PM and decaying after this. From Figure 2(b), there appears to be much less variation in the intensity across the week compared to all other identified components, but a slightly higher intensity on Thursdays and Fridays, and the lowest on Tuesdays and Saturdays. Finally, we see a small increase in the trend during the first 7 weeks of the data, then it remains reasonably flat until week 28, where it begins to rise again. Around week 40, it stabilizes again. This could be due to an increase in actual event

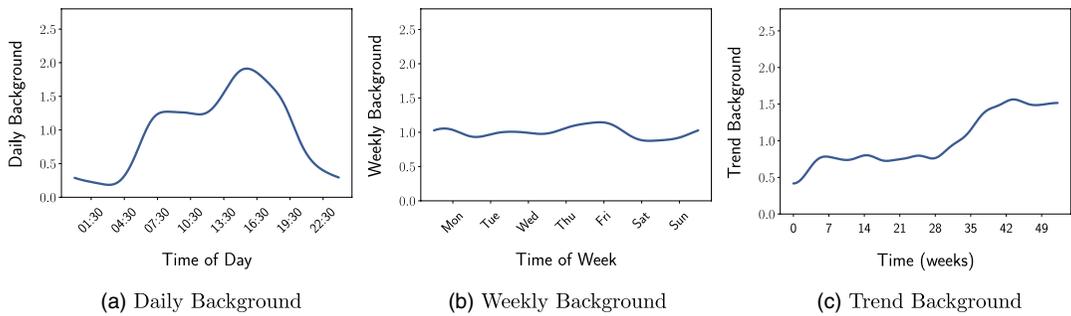


FIGURE 2 Temporal background components fit to 1 year of data [Colour figure can be viewed at wileyonlinelibrary.com]

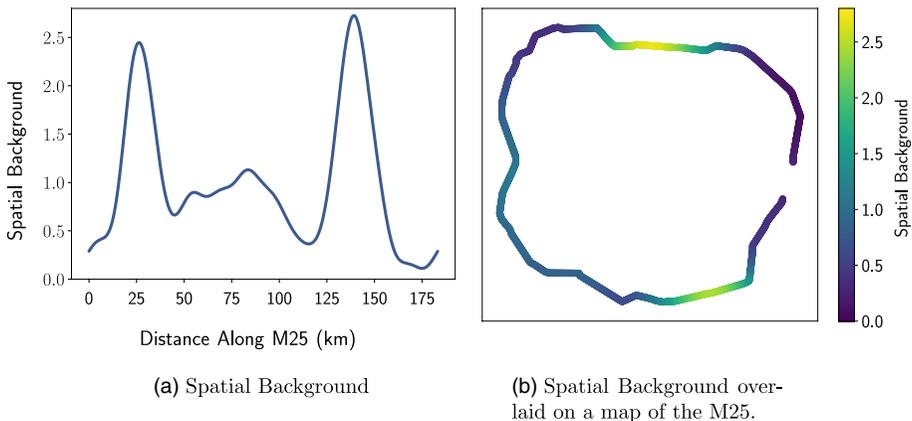


FIGURE 3 Spatial background components fit to 1 year of data [Colour figure can be viewed at wileyonlinelibrary.com]

intensity, or more comprehensive reporting after a certain point, more operators and so forth, however no changes in reporting are known to us.

As well as the temporal background, where events are most common around the M25 is of interest. We show the spatial background in Figure 3, from which a clear spatial structure is visible. We see two distinct peaks in Figure 3(a), and a smaller spread out peak in between these two. The largest peak, around 140 km along the motorway, is located near the ‘Potters Bar’ junction. The second largest, around 25 km into the motorway, is located between where the M25 meets the M26, and where the M25 meets the M23. This background intensity, imposed onto a map schematic of the real M25 is plotted in Figure 3(b). We investigated if the temporal background in the vicinity of these spatial peaks differed to that across the entire M25, but found only minor changes. Interested readers can find this analysis in the supplementary material.

5.3 | Triggering analysis

Triggering does appear to improve the log-likelihood of our model, and as seen in Table 1, it explains around 6.55% of events in the data. We visualize the resulting functions in Figure 4. From Figure 4(b), it is clear that spatial triggering is limited to around 2 km, after which we do not see any non-zero values. However, temporal triggering, pictured in Figure 4(a), appears to have quite a ‘long tail’ in the

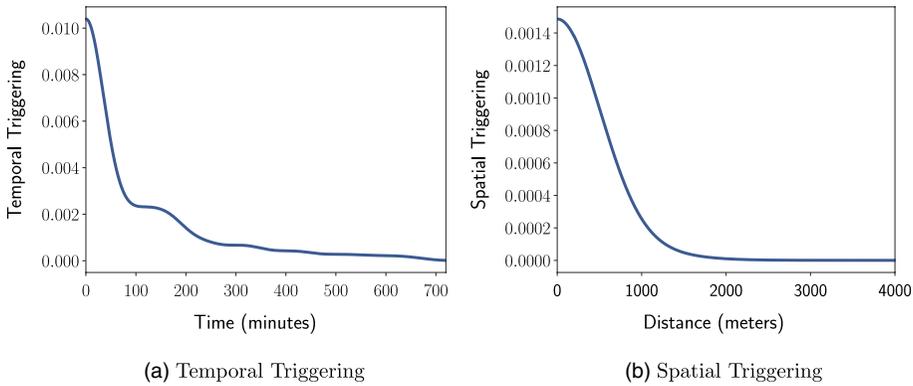


FIGURE 4 Triggering functions, fit to 1 year of data [Colour figure can be viewed at wileyonlinelibrary.com]

sense that it decays over a very long range. However, this is a clear time scale in this result of around 100 min. As a result of this, one should take 6.55% as an upper bounds of sorts, understanding this, combined with the identified time and length scales is an informative conclusion to draw.

5.4 | Model validation

To validate if our model captures the relevant features of a process, one often follows the work of Ogata (1988), Brown et al. (2002) and Zhuang (2006). Specifically, we use the transformed time sequence given by:

$$t_i \rightarrow \Lambda_i = \int_0^{t_i} \int_0^X \lambda(u,x) dx du, \tag{17}$$

then the resulting sequence of values Λ_i for all $i \in \{1, 2, \dots, N\}$ will follow a unit rate Poisson process if the model is correctly specified. To prove this, one can derive that the sequence of values $\Lambda_i - \Lambda_{i-1}$ are i.i.d. random variables that follow an exponential distribution with parameter 1, which is done in Brown et al. (2002). Given that we know the expected distribution of $\Lambda_i - \Lambda_{i-1}$, we can then transform this to follow a standard uniform distribution by computing:

$$z_i = 1 - e^{-(\Lambda_i - \Lambda_{i-1})}. \tag{18}$$

Now, the computed z_i values should follow the simplest known distribution, which we can evaluate by comparing the measured and theoretical cumulative distribution functions (CDF), as well as the measured and expected quantiles. One can generate confidence bounds for the comparison of two CDFs by inversion of the Kolmogorov–Smirnov statistic, and for quantile–quantile (QQ) plots by using the fact that the order statistics of a uniform distribution follow a beta distribution. The distribution of the k -th order statistic has parameters k and $n + 1 - k$, with n being the number of sample points. Using this, we then generate CDF and QQ plots for our different modelling scenarios, given in Figure 5(a) and (b). It is clear from Figure 5 that the model is statistically defensible when inspecting both the CDF and QQ plots of the results. Some of the quantiles in Figure 5(b) are just on the edge of acceptable, but do not deviate outside of the confidence bands. These results show that our model is well specified, however extra components could continue to be added to improve the fit.

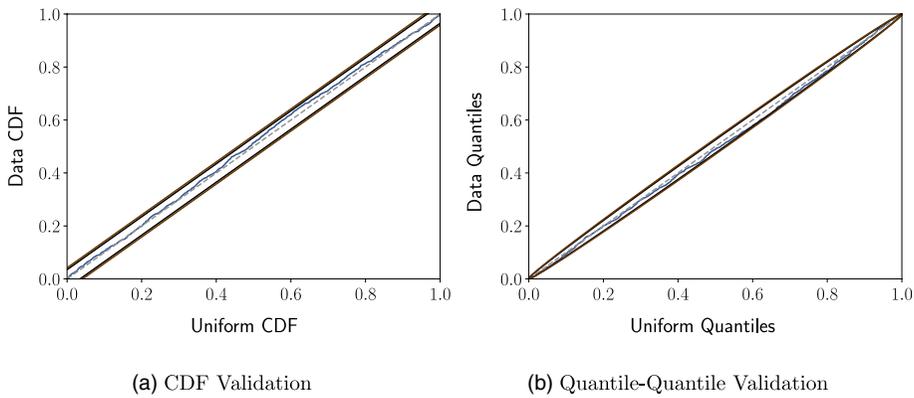


FIGURE 5 Validation results for 1 year of data. Model results are shown with —, 95% limits: —, 99% limits: — and a reference line: - - - (colour online) [Colour figure can be viewed at wileyonlinelibrary.com]

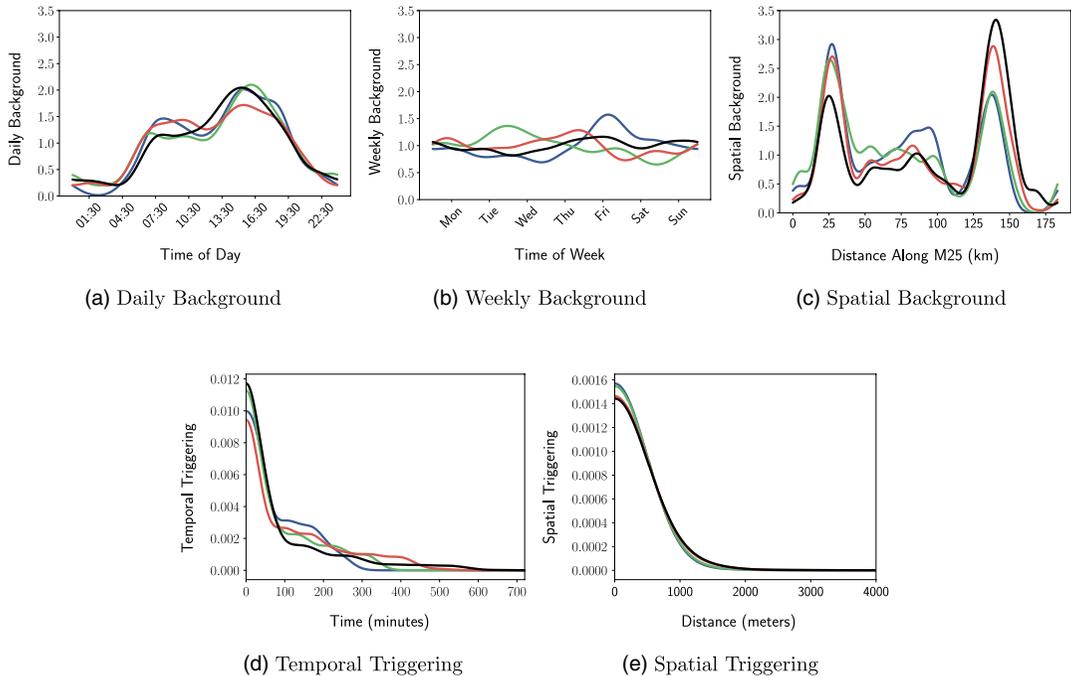
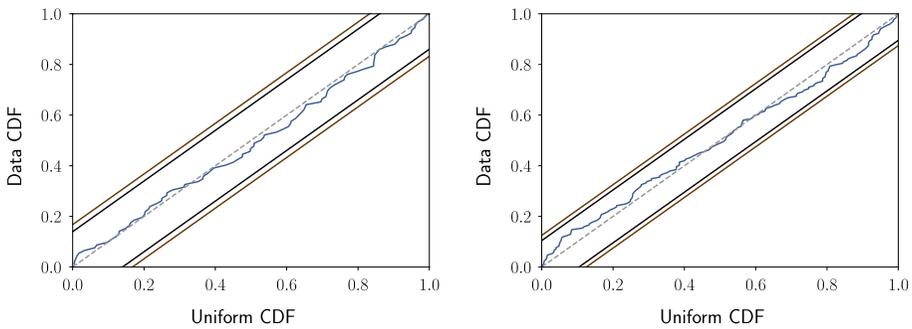


FIGURE 6 Background and triggering components compared across different 3-month data sets. Data sets are: 09/2017–11/2017: —, 12/2017–02/2018: —, 03/2018–05/2018: — and 06/2018–08/2018: — (colour online) [Colour figure can be viewed at wileyonlinelibrary.com]

5.5 | Do components change with season?

Given our results for the year of data, we can question how resilient the background and triggering components are by inspecting subsets of the data. To test this, we partition our data into 3-month seasonal periods, and fit the model to each subset. We then overlay the components in Figure 6. It is clear from our results in Figure 6 that there is varying amounts of consistency in components of the model. Generally, we see in Figure 6(a) that the daily background component is constant throughout the year,



(a) Out of sample CDF Validation, fit to data 12/2017-02/2018, validated on data for 03/2018.

(b) Out of sample CDF Validation, fit to data 06/2018-08/2018, validated on data for 09/2018.

FIGURE 7 Out-of-sample model validation for two data sets. Model results are shown with —, 95% limits: ———, 99% limits: ——— and a reference line: - - - (colour online) [Colour figure can be viewed at wileyonlinelibrary.com]

the only variation being in the 3-month data starting on 03/2018, where the morning peak is a little more pronounced than any of the other data sets, and the evening is a little less. It is difficult to make conclusions about the weekly component, when we have only 12 instances of each day of the week for a given 3-month subset. Of particular interest is the spatial background through time, shown in Figure 6(c), where we see that the peak around 140 km along the M25 actually becomes more pronounced as time progresses. Both the peak at 25 and 140 km are present throughout all subsets of data, but later periods appear to show that the peak around Potters Bar is more significant later in the data set compared to earlier. It is unclear if a physical change occurred leading to this, but would be of interest to investigate with more data. Finally, the temporal and spatial triggering functions are generally quite consistent across all data sets. We still see somewhat of a long decay in the temporal triggering, but time scales of around 100 min remain. Considering the triggering in each of these subsets, we attain A values of 0.034, 0.063, 0.068 and 0.075 for the periods starting 2017/09, 2017/12, 2018/03 and 2018/06, respectively. Without more data, we cannot say if these values vary throughout the year, or if they are increasing over time.

We note that, for small temporal subsets of data, the trend component becomes less impactful in the model. Recalling Figure 2(c), the trend is almost flat for long periods, suggesting we can omit it for some temporal subsets. As a result, we can perform out-of-sample model validation, something we are not aware of being done previously in the literature for this type of model. Since the trend is omitted, and we see the triggering components are consistent through time, we can train the model on some 3-month subset of data, then perform our validation on unseen data. We show two examples of this in Figure 7(a) and (b). It is clear from Figure 7 that on short time scales, one can use the model for acceptable performance out-of-sample. This is not possible over longer time scales, however, due to the clear trend and varying spatial background evident in the data.

5.6 | Do components change for significant events?

While all events flagged in NTIS should correspond to actual traffic incidents, many of them may not have had a significant impact on the traffic state. Consider the case where two vehicles have a minor collision, create no debris, and the drivers pull into the closed hard-shoulder to exchange insurance information. In such a case, there should be little impact on flow, travel time and speed. Additionally,

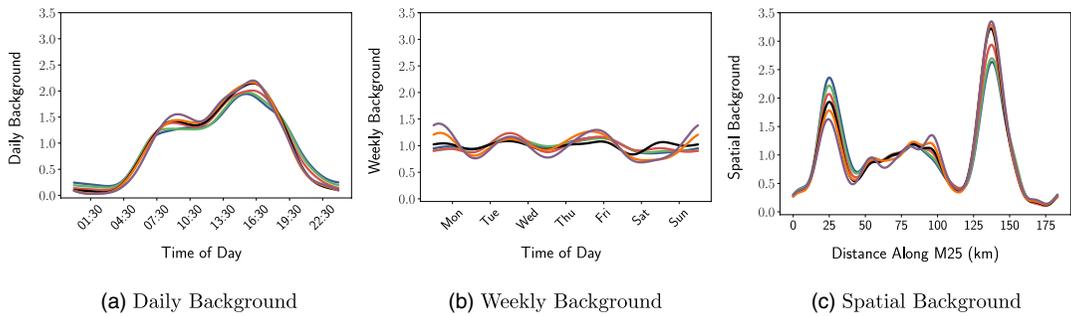


FIGURE 8 Background components compared across significant events, varying what speed decrease is required to define an event as significant. Thresholds are: 0%: —, 10%: —, 20%: —, 30%: —, 40%: — and 50%: — (colour online) [Colour figure can be viewed at wileyonlinelibrary.com]

if a vehicle breaks down and pulls into the hard-shoulder, and the road is far from capacity, we should again see little drop in average speed. To consider only the behaviour of events that have some significant impact, we inspect the link-level data, which contains significantly less noise than the loop level. For a given event window, we consider the largest percentage drop in speed between a simple historical median segmentation profile and measured values across the entire window. If this percentage drop is above some threshold, then we say the event caused a significant impact on the traffic state. Further discussion of this seasonal model is given in the supplementary material. As we raise this threshold, we isolate more extreme events but discard so much data that it is no longer reasonable to fit a model. To retain enough data for fitting, we consider only thresholds between 0 and 50%. We split our data set into subsets containing only events that lead to a speed decrease of at-least 0%, 10%, ..., 50%, then re-run our model fitting on these subsets. The resulting background components are visualized in Figure 8.

Figure 8(a) shows the daily background is reasonably stable across different thresholds of significance. As we raise the threshold, the morning and evening peak structure becomes clearer, and the periods very early and late in the day are lowered. Figure 8(b) shows that weekends have lower intensity than weekdays as we raise the threshold for significance. This is likely due to demand being significantly lower on weekends compared to weekdays, and hence when an event does occur, there is less chance of a queue forming as the road is further from capacity than on a weekday. Finally, the spatial background in Figure 8(c) clearly shows that the second peak, identified around 140 km along the M25, appears to experience more significantly impactful events than the first peak, observed around 25 km along the M25. This suggests that not only is Potters Bar a ‘hot-spot’ for events, but also that the events here are some of the more extreme on the network. Analysis of triggering during different significance thresholds shows that the time scales and A values remain consistent throughout. The only visible difference in the triggering functions is that for larger thresholds, the functions decay more quickly.

6 | CONCLUSION

We have analysed the spatiotemporal variation in the incident rate on London's M25 motorway over a period of 1 year using a model that distinguishes between primary and secondary events. This variation is found to be strongly inhomogeneous. The temporal variation shows a strong daily double peak structure reflecting commuting patterns superimposed in a weaker weekly variation with a peak on

Fridays and a trough on Saturdays. This pattern of temporal variation remains stable over the data period. The spatial variation shows two primary peaks in intensity. The first and largest is in the vicinity of the Potters Bar Interchange. The other is in the vicinity of Junctions 5 and 6 where the M26 and M23 join the M25. The peak at Potters Bar appears to increase in intensity during the data period, and is more pronounced when we condition on the most significant events in terms of impact on traffic speed. We find that 6–7% of the observed incidents are most probably secondary incidents under the assumptions of our model. Plausible time and length scales emerge for the range of the triggering effects: 100 min in the temporal triggering, and 1 km for spatial triggering. From these figures, we conclude that the effects of secondary incidents are a small but detectable feature of the M25 incident data set. We suggest that, on the M25, the scope to further reduce incident rates by reducing secondary incidents is limited compared to what could be achieved by reducing the peaks at specific times or ‘hot-spot’ locations.

We have extended existing work to limit the freedom of the self-excitation components for non-parametric Hawkes processes, along with considering unidirectional spatial triggering. We further proposed an out-of-sample validation method which is viable when the trend component is roughly constant over time. Applying this to our data set showed we have avoided overfitting the model, despite its significant freedom. These advancements are not specific to the application discussed here, and can be used when models of this form are applied to different domains with similar practical considerations. The modelling framework is novel in the context of traffic incident analysis, and our methodological refinements ensure the model captures relevant properties of the application.

In terms of further work, it would be informative to repeat our analysis for a major road without the MIDAS system or smart motorway features since we expect this infrastructure to reduce the risk of secondary incidents. Additionally, it would be interesting to apply this analysis to urban road network, however directly extending our methodology to account for this structure appears difficult. While smoothing of point-processes on a network has been explored recently, for example in Moradi and Mateu (2019), further complication remains regarding smoothing of the triggering functions on a network. When we take data points in \mathbb{R}^N and compute distances between them, we attain a value in \mathbb{R}^+ , and differences between these distances lie in \mathbb{R} . However, the same is not necessarily true for differences between distances computed on a network, which still attain real values, but belong to the sub-space stemming from their network-constrained spatial structure. Taking this into account makes smoothing the triggering functions complex and deserves further attention in future work. General comments about clustering of events on urban roads, made in Acker and Yuan (2019), suggest that secondary incidents remain a feature of urban roads. One could also extend our out-of-sample validation to include a trend component. To do so, a parametric form of the trend is required, as we cannot extrapolate our current non-parametric estimate. Finally, one could incorporate more variables that may influence traffic incidents, for example weather conditions.

ACKNOWLEDGEMENTS

We thank Dr. Steve Hilditch, Thales UK for sharing expertise on NTIS and UK transportation systems. This work was supported by the EPSRC (grant number EP/L015374/1).

REFERENCES

- Acker, B. & Yuan, M. (2019) Network-based likelihood modeling of event occurrences in space and time: A case study of traffic accidents in Dallas, Texas, USA. *Cartography and Geographic Information Science*, 46, 21–38.
- Aguero-Valverde, J. & Jovanis, P.P. (2008) Analysis of road crash frequency with spatial models. *Transportation Research Record*, 2061, 55–63.

- Brown, E.N., Barbieri, R., Ventura, V., Kass, R.E. & Frank, L.M. (2002) The time-rescaling theorem and its application to neural spike train data analysis. *Neural Computation*, 14, 325–346.
- Chiu, S.-T. (2000) Boundary adjusted density estimation and bandwidth selection. *Statistica Sinica*, 10, 1345–1367.
- Department for Transport. (2019) Reported road casualties in Great Britain: main results 2018. *Tech. rep.* Available from: <https://www.gov.uk/government/statistics/reported-road-casualties-great-britain-main-results-2018>.
- Dixon, P.M. (2014) *Ripley's K function*. Atlanta: American Cancer Society.
- Fan, Y., Zhu, X., She, B., Guo, W. & Guo, T. (2018) Network-constrained spatio-temporal clustering analysis of traffic collisions in Jiangnan District of Wuhan, China. *PLoS ONE*, 13, 1–23.
- Fox, E.W., Short, M.B., Schoenberg, F.P., Coronges, K.D. & Bertozzi, A.L. (2016) Modeling e-mail networks and inferring leadership using self-exciting point processes. *Journal of the American Statistical Association*, 111, 564–584.
- Hall, P. & Huang, L.-S. (2001) Nonparametric kernel regression subject to monotonicity constraints. *The Annals of Statistics*, 29, 624–647.
- Highways England. (2020) How to drive on a smart motorway. Available from: <https://www.gov.uk/guidance/how-to-drive-on-a-smart-motorway>.
- Kulldorff, M. (2001) Prospective time periodic geographical disease surveillance using a scan statistic. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 164, 61–72.
- Lesage, L. (2020) A Hawkes process to make aware people of the severity of COVID-19 outbreak: application to cases in France. *Research report*, Université de Lorraine; University of Luxembourg.
- Li, Z., Cui, L. & Chen, J. (2018) Traffic accident modelling via self-exciting point processes. *Reliability Engineering & System Safety*, 180, 312–320.
- Lim, K.W., Wang, W., Nguyen, H., Lee, Y., Cai, C. & Chen, F. (2016) Traffic flow modelling with point processes. In *Proceedings of the 23rd World Congress on Intelligent Transport Systems*, 1–12. ITS.
- Mohler, G. (2014) Marked point process hotspot maps for homicide and gun crime prediction in Chicago. *International Journal of Forecasting*, 30, 491–497.
- Mohler, G.O., Short, M.B., Brantingham, P.J., Schoenberg, F.P. & Tita, G.E. (2011) Self-exciting point process modeling of crime. *Journal of the American Statistical Association*, 106, 100–108.
- Moradi, M.M. & Mateu, J. (2019) First- and second-order characteristics of spatio-temporal point processes on linear networks. *Journal of Computational and Graphical Statistics*, 29, 1–21.
- Ogata, Y. (1988) Statistical models for earthquake occurrences and residual analysis for point processes. *Journal of the American Statistical Association*, 83, 9–27.
- Payne, H.J., Helfenbein, E.D. & Knobel, H.C. (1976) Development and testing of incident detection algorithms. vol. 2, research methodology and detailed results. *Tech. Rep. FHWA-RD-76-20*, Federal Highway Administration, Washington D.C.
- Reinhart, A. (2018) A review of self-exciting spatio-temporal point processes and their applications. *Statistical Science*, 33, 299–318.
- Rey, S.J. (2016) Space-time patterns of rank concordance: Local indicators of mobility association with application to spatial income inequality dynamics. *American Association of Geographers Annals*, 106, 788–803.
- Schoenberg, F.P., Hoffmann, M. & Harrigan, R.J. (2019) A recursive point process model for infectious diseases. *Annals of the Institute of Statistical Mathematics*, 71, 1271–1287.
- Song, J., Wen, R. & Yan, W. (2018) Identification of traffic accident clusters using Kulldorff's space-time scan statistics. In *2018 IEEE International Conference on Big Data (Big Data)*, 3162–3167.
- World Health Organization. (2018) *Global status report on road safety 2018*. Geneva: World Health Organization.
- Zhuang, J. (2006) Second-order residual analysis of spatiotemporal point processes and applications in model evaluation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68, 635–653.
- Zhuang, J. (2011) Next-day earthquake forecasts for the Japan region generated by the ETAS model. *Earth, Planets and Space*, 63, 5.
- Zhuang, J. & Mateu, J. (2019) A semiparametric spatiotemporal hawkes-type point process model with periodic background for crime data. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 182, 919–942.
- Zhuang, J., Ogata, Y. & Vere-Jones, D. (2004) Analyzing earthquake clustering features by using stochastic reconstruction. *Journal of Geophysical Research*, 109, 17 p. Available from: <https://agupubs.onlinelibrary.wiley.com/doi/full/10.1029/2003JB002879>.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section.
Supplementary Material

How to cite this article: Kalair K, Connaughton C, Alaimo Di Loro P. A non-parametric Hawkes process model of primary and secondary accidents on a UK smart motorway. *J R Stat Soc Series C*. 2021;70:80–97. <https://doi.org/10.1111/rssc.12450>