SPATIO-TEMPORAL EXTREME EVENT MODELING OF TERROR INSURGENCIES

A PREPRINT

Lekha Patel*

Lyndsay Shand

J. Derek Tucker

Gabriel Huerta

Statistical Sciences, Sandia National Laboratories, Albuquerque, NM

October 29, 2021

ABSTRACT

Extreme events with potential deadly outcomes, such as those organized by terror groups, are highly unpredictable in nature and an imminent threat to society. In particular, quantifying the likelihood of a terror attack occurring in an arbitrary space-time region and its relative societal risk, would facilitate informed measures that would strengthen national security. This paper introduces a novel self-exciting marked spatio-temporal model for attacks whose inhomogeneous baseline intensity is written as a function of covariates. Its triggering intensity is succinctly modeled with a Gaussian Process prior distribution to flexibly capture intricate spatio-temporal dependencies between an arbitrary attack and previous terror events. By inferring the parameters of this model, we highlight specific space-time areas in which attacks are likely to occur. Furthermore, by measuring the outcome of an attack in terms of the number of casualties it produces, we introduce a novel mixture distribution for the number of casualties. This distribution flexibly handles low and high number of casualties and the discrete nature of the data through a Generalized ZipF distribution. We rely on a customized Markov chain Monte Carlo (MCMC) method to estimate the model parameters. We illustrate the methodology with data from the open source Global Terrorism Database (GTD) that correspond to attacks in Afghanistan from 2013-2018. We show that our model is able to predict the intensity of future attacks for 2019-2021 while considering various covariates of interest such as population density, number of regional languages spoken, and the density of population supporting the opposing government.

Keywords Self exciting point processes \cdot Discrete extreme value distribution \cdot Bayesian methods \cdot Global terrorism data base \cdot Afghanistan

1 Introduction

Recent global developments have seen a surge in high-consequence events with life-threatening consequences. The emergence of such extreme situations has heightened the urgency of their detection, prevention, and deterrence. While many quantitative methods endeavor to address this topic, few are able to fully consider the spatio-temporal impact of historic records on the risk of subsequent lethal events. In this paper, we propose a novel approach to detecting and predicting terrorism attacks using concepts from extreme value analysis and a spatio-temporal self-exciting point process framework. The focus of this work is in modeling *terror insurgencies*, violent uprisings organized by groups such as the Taliban, Boko Haram, and the Islamic State (IS), using a repertoire of terror-based attacks against civilians in an attempt to overthrow a government of interest. Application of our model will be highlighted by studying the insurgency of the Taliban in Afghanistan between 2013-2021, a highly relevant socio-political topic that is currently generating significant global coverage.

^{*}Corresponding author: lpatel@sandia.gov

In the literature, traditional statistical models seek to describe the overall distribution of such rare events of modeling interest, while extreme value analysis prioritizes the characterization of events that lie in the tails of these distributions. Extreme value methods have commonly been used to predict and quantify uncertainty around environmental or climatological events associated with high costs such as a high impact on human casualties (e.g., earthquakes, hurricanes, flooding, wildfires), as described in for example, Coles (2001). From a statistical point of view, extreme value data can be modeled based on two types of approaches. The first relies on calculating a sequence of maximum (or minimum) values over blocks of data, e.g. monthly or yearly maxima (minima) and fitting these values to their large sample distribution, the Generalized extreme value (GEV) distribution. The second finds observations that exceed (or fall below) a given threshold and fits the Generalized Pareto distribution (GPD) to these exceedance values. The latter method is also known as a Peaks over Threshold (PoT) approach and can make a more effective use of the data rather than only considering the block maxima. Additionally, a PoT approach can be more flexible than analyses of block maxima as it can allow for the simultaneous prediction of the time and characteristics of an extreme event. The utilization of both methods are application specific but have not been considered extensively to assess extreme high consequence terrorism events. However, the specialized literature includes extensive developments of both block-maxima and PoT approaches for time series and point-referenced spatial data. The interested reader is directed to the works of Castro-Camilo et al. (2019); Bortot and Gaetan (2014); Bortot and Tawn (1998); Bracken et al. (2016); Cooley and Sain (2010); Cooley et al. (2007); Gaetan and Grigoletto (2007); Huerta and Sanso (2007); Sharkey and Winter (2019) for more information on temporal or space-time extreme value analysis and some of its applications.

In this work, we propose novel methodology to analyze spatio-temporal patterns of extreme terror attacks with potential life-threatening implications. Our model is intended to quantify and predict the time, location, and likelihood of a future extreme event, with the extremity of an event defined by the number of casualties it produces. A specific aspect of our model for this application considers point process modeling and therefore follows the works of Patel et al. (2020); Tucker et al. (2019); Zammit-Mangion et al. (2012). Point process models are useful to describe, in a probabilistic manner, phenomena that occur in space-time, and are therefore highly suitable to model complex terror insurgencies. In such situations, *homogeneous* point processes, which would assume a constant rate of attacks, insufficiently capture event clustering in space-time. Instead, we rely on self-exciting processes to capture the spatio-temporal dependencies that naturally arise in terror formations and resurgences, as illustrated previously by Mohler et al. (2011) which focused on crime modeling and dispersion in Chicago and Tucker et al. (2019) which looked at issues of incomplete data to characterize terrorism events in Colombia.

Our methodology is sufficiently flexible to characterize the spatio-temporal behavior of region-specific attacks and their associated number of casualties, which represent the marks of a marked spatio-temporal point process underlying terror behavior. To account for the extreme behavior of these marks, we develop a novel mixture model to estimate a discrete *mark* distribution that simultaneously models situations with a low number of casualties and cases where casualties can be considered extreme or *above a threshold*. In particular, we show that the marks above a threshold are well estimated via a *Generalized-ZIPF* distribution which gives an analog representation to the well known GPD or PoT approach as described previously. Therefore, we can appropriately model the number of casualties for extreme terrorism events separately from events which fall within the expected casualty range. In the context of financial time series, Grothe et al. (2014) proposed a multivariate time series model within the PoT framework where clustering of events is captured via a self exciting point process. The model parameters are estimated with MLE and illustrated with the analysis of daily log-returns obtained from financial markets. To the best of our knowledge, the combination of self-exciting mark point processes with a discrete extreme value distribution has not been considered before for the analysis and prediction of terror insurgencies.

This paper is organized as follows. Section 2.1 introduces our motivating dataset, the Global Terrorism Database previously considered by Tucker et al. (2019) and the different aspect of focus for this work. Section 3 presents our novel and flexible marked point process framework and details the mark distribution choices made via a mixture representation. It also discusses an inference approach based upon a novel application of Markov chain Monte Carlo (MCMC) methods. Section 4 applies our model to data from the GTD and in difference to Tucker et al. (2019) we focus on data from 2013-2018 corresponding to Afghanistan and illustrate the capability of our model in predicting the intensity of future events. Section 5 gives our conclusions and discusses potential extensions of our work.

2 Data

2.1 Global Terrorism Database

The methodology proposed in this paper focuses on the Global Terrorism Database, GTD (2017). The GTD is an open-source database which can be downloaded at https://www.start.umd.edu/gtd/, and includes information on more than 190,000 terrorist attacks from 199 different countries occurring from 1970 to 2018. Along with the attack

time (day, month and year) and location (latitude and longitude) of each recorded terrorism event, detail is also given on the attack type, weapon(s) used, nature of the target, casualties, injuries, and the group responsible (when available). While extreme value theory has been shown to be a valid approach for predicting the likelihood of catastrophic terrorism events in Mohtadi and Murshid (2009), using a more limited data set than the GTD, these analyses had not been extended to appropriately incorporate spatial or other critical descriptive information of the terrorism events.

Although the GTD contains numerous accounts of attacks, events are not consistently distributed evenly across the globe or across time. Data limitations for some countries e.g. the USA and UK from which there is insufficient data, additionally prove challenging for extreme value analysis (EVA). This therefore limits our focus to specific regions in the world (e.g. countries) where more events occur and thus the resulting analyses can provide meaningful interpretations. The thresholds which define extreme events will be region dependent as seen in Figure 1a which shows the number of casualties per country for cases which reported the highest casualties due to terrorism events since 1970.

Data quality may also affect subsequent analyses of extreme events. For example, the GTD is incomplete and contains missing entries on one or more of the *explanatory covariates* and *time stamps* of events. The *spatial* locations of events, given by latitude and longitude, are estimated at the nearest city and thus can be inaccurate. As highlighted by Tucker et al. (2019), these aspects of the data are important to account for when developing a statistical model. The data for the year 1993 is completely missing and although there have been efforts to recollect this data, these have not been successful. Additionally, it is possible that many countries, especially those with attacks recorded before 2000, have *underreported* the severity of the attack. Including this data, therefore, would likely lead to misleading results. Although the GTD (2017) team has made improvements of the methodology used to compile the database, by balancing the strengths of "artificial and human intelligence" based on more diverse sets of news media from around the world, we choose to omit this data from our analysis and focus on data from 2013 to 2018.

In order to develop an appropriate statistical framework for extreme global terror insurgency attacks, we first require an extreme event to be defined and assess the viability of these distributions to fit the data. Depending on the question of interest, an extreme attack could either be characterized as occurring in an unusual or unexpected space-time location, or as an extremely rare type such as an event perpetrated by a new terrorist group or an event reporting a high magnitude of fatalities or causalities. The former of the two requires a further definition of what constitutes as an "unexpected" space-time location and could therefore be difficult to formally formulate in comparison to the latter. We therefore turn to rare type modeling and begin, in a similar manner to Mohtadi and Murshid (2009), by defining an extreme attack as one that results in a high reporting of casualties. The number of casualties here is calculated as the number of deaths and injuries resulting from an attack.

On a country specific level, the countries which have sufficient data and quoted the most deadly attacks in the GTD are: Iraq, Pakistan, Afghanistan, India and Colombia as shown in Figure 1a. Log scaled empirical cumulative distribution functions (CDFs) and quantile plots shown in Figures 1b and 1c respectively, suggest that the distributions of casualties reported from the most deadly countries (Iraq/Afghanistan, India/Pakistan/Colombia) are similar and heavy tailed. These exploratory plots indicate that extreme valued type of distributions would be applicable for this data. Figure 2 shows the location of terrorist attacks across the region where the per-country number of casualties is above the 97% empirical quantile. The patterns observed in this figure give an indication that there is a significant spatial contribution on the underlying distribution of extreme events in India, Afghanistan, Iraq and Pakistan, also rendering that high-casualty events lead to clustered behavior.

3 Materials and methods

The temporal nature of terrorism events can be identified as being *self-exciting* in nature as described by (Mohler et al., 2011; Mohler, 2013; Clark and Dixon, 2018; Tucker et al., 2019), indicating that events occurring in time are likely to be dependent on any previous event in the history of the process. It is also understood that spatially, violent crime activity is likely to occur in or around locations that have had previously successful attacks (Mohler et al., 2011). Such behavior therefore induces clustering patterns and Figure 2 shows similar clustering for high causality events. On the other hand, new terrorist activity often arises in regions with denser populations (Zammit-Mangion et al., 2012). Other variables such as gross domestic product (GDP), ethnicity distribution and geographical terrain additionally indicate regions that are prone to terror events and could be considered as explanatory variables.

To address the nuanced spatio-temporal behavior of terror events and the casualties that are reported from them, we utilize a marked spatio-temporal modeling approach. This model considers a terror event as an attack taking place at an arbitrary space-time location such that the *mark* of each event defined the number of casualties that occur from it.

Let $\{X(\mathbf{s},t,m) : \mathbf{s} \in W, t \in [0,T], m \in \mathbb{N}_0\}$ denote the marked point process of terror attacks indexed by the spatial location \mathbf{s} observed in a bounded region W, the continuous time $t \leq T$, and the number of casualties m. The counting

Number of casualties per country



(a) Histogram of all casualties from countries with the highest casualties across all attacks since 1970.

Quantile plots per country on log scale



(c) Quantile function.

Figure 1: Plots of all global casualties during 1970-2018.

measure of events associated with this process is denoted $N_X(A \times M) = |X \cap (A \times M)|$ for $A \subseteq W \times [0, T], M \subseteq \mathbb{N}_0$. For any $A \times M$, its *intensity measure* conditional on the history of the process \mathcal{H}_t is defined as the expected number of events of X in $A \times M$ given \mathcal{H}_t , i.e. as $\mathbb{E}(N_X(A \times M)|\mathcal{H}_t)$. Its *intensity function* denoted $\lambda(\mathbf{s}, t, m|\mathcal{H}_t)$ is related to this measure via $\mathbb{E}(N_X(A \times M)|\mathcal{H}_t) = \sum_{m \in M} \int_A \lambda(\mathbf{s}, t, m|\mathcal{H}_t)$ ds dt (Møller and Waagepetersen, 2004).

3.1 Intensity characterization

For a given time t and location s in our domain of interest $(A = W \times [0, T])$, the likelihood of an event with mark m can be expressed through its intensity function, whose general form (Reinhart, 2017) is

$$\lambda(\mathbf{s}, t, m; \mathcal{H}_t) = f_M(m; \mathcal{Z}_m(\mathbf{s}, t))\lambda^*(\mathbf{s}, t; \mathcal{H}_t, \mathcal{Z}_{\mathbf{s}, t}).$$
(1)

Here, $f_M(m; \mathcal{Z}_m(\mathbf{s}, t))$ is a discrete probability density function for the marks given the mark covariates $\mathcal{Z}_m(\mathbf{s}, t)$ observed at location \mathbf{s} and time t. Further, $\lambda^*(\mathbf{s}, t; \mathcal{H}_t, \mathcal{Z}_{\mathbf{s},t})$ is the conditional intensity function of the spatio-temporal process at time $t \in [0, T]$ and spatial location $\mathbf{s} \in W$ given the history of events \mathcal{H}_t and spatio-temporal covariates $\mathcal{Z}_{\mathbf{s},t}$.





Figure 2: PoT casualty data between 2002-2018. The locations shown corresponding to cases where the per-country casualty level is above the empirical 97% quantile

3.2 Mark distribution

Prior data analysis studying continuous extreme valued distributions (GPD and GEV) to country level datasets provided by the GTD, have shown good fits for the number of casualties (Patel et al., 2020). In particular, Patel et al. (2020) detail that the GPD or POT approach provides the best overall fit to the tails of casualty data, caveatted by a requirement for an exceedance threshold to be determined and the distribution's prime suitability for continuous data. To deal with this, we follow the work of MacDonald et al. (2011) and utilize a discrete *mixture* distribution that is able to model an extreme number of (discrete) casualties produced from each attack and also its zero-inflated behavior, as highlighted by the quantile behavior shown in Figures 1b-1c.

After fixing an exceedance threshold $u \in \mathbb{N}$, we utilize a zero-inflated discrete distribution for all marks $m \in [0, u]$ below a threshold u that supports observations of attacks with a low number of casualties. Above u, we consider a discrete extreme value distribution addressing attacks resulting in an extreme number of casualties. We model $f_M(m; \mathcal{Z}_m(\mathbf{s}, t))$ via the mixture distribution

$$f_M(m; \pi_M, u, \boldsymbol{\theta}_{ZI}, \boldsymbol{\theta}_{EV}) = \pi_M \frac{f_{ZI}(m; u, \boldsymbol{\theta}_{ZI})}{\sum_{m' \le u} f_{ZI}(m'; u, \boldsymbol{\theta}_{ZI})} \mathbb{1}_{m \le u} + (1 - \pi_M) \frac{f_{EV}(m; u, \boldsymbol{\theta}_{EV})}{\sum_{m' > u} f_{EV}(m'; u, \boldsymbol{\theta}_{EV})} \mathbb{1}_{m > u},$$
(2)

where $f_{ZI}(m; u, \theta_{ZI}, \mathcal{Z}_m(\mathbf{s}, t))$ denotes a general zero-inflated distribution to model the number of casualties up to threshold u given parameters θ_{ZI} and $f_{EV}(m; u, \theta_{EV}, \mathcal{Z}_m(\mathbf{s}, t))$ denotes a general extreme valued distribution to model the number of casualties *above* u given parameters θ_{EV} . The parameter π_M denotes the mixing proportion, or weight, of the distribution corresponding to each case. For simplicity in notation, we drop the explicit dependence on $\mathcal{Z}_m(\mathbf{s}, t)$ for both f_{ZI} and f_{EV} .

3.2.1 Casualty model selection

Suitable probability mass functions for f_{ZI} include the zero-inflated Poisson (ZIP) and Negative-Binomial (ZINB) distributions, both of whose parameters θ_{ZI} can be written as functions of the exogenic covariates $\mathcal{Z}_m(\mathbf{s}, t)$. For all $m \leq u$, their probability distributions are given by

$$f_{ZIP}(m; u, \boldsymbol{\theta}_{ZI}) = \alpha \mathbb{1}_{m=0} + (1 - \alpha) \frac{\lambda^m \mathrm{e}^{-\beta}}{m!} \qquad \qquad \alpha \in [0, 1], \ \beta > 0 \qquad (3)$$

$$f_{ZINB}(m; u, \boldsymbol{\theta}_{ZI}) = \alpha \mathbb{1}_{m=0} + (1 - \alpha) \binom{m + r - 1}{m} (1 - p)^r p^m \qquad \alpha, p \in [0, 1], \ r > 0.$$
(4)

Determining a form for f_{EV} requires considering concepts from *discrete extremal functions* which are suitable to this problem. Specifically, for $m \ge u$, we consider the *Generalized-ZIPF* distribution (GZD) defined by Hitz et al. (2017) as

$$f_{GZD}(m; \boldsymbol{\theta}_{EV}) \propto \begin{cases} \left(1 + \frac{\xi(m-u)}{\sigma}\right)^{-\frac{1}{\xi}-1} & \xi > 0\\ \exp\left(-\frac{(m-u)}{\sigma}\right) & \xi = 0, \end{cases}$$
(5)

a flexible class of extreme valued discrete distributions decaying via the Zipf-Mandelbrot power-law (Mandelbrot, 1953). Zipf's law created by Zipf (1949) originally studied the distribution of word frequencies in common language and has resulted in a wide range of Zipf-type distributions within different frameworks (Li, 1992). An advent of its usage in the social sciences, for example, to analyze voting histories at different levels of social systems (Lyra, M. L. et al., 2003), population densities in cities (Gabaix, 1999) and fake social media trends (Rastogi, 2016) complementing its extreme valued properties, renders the GZD a suitable candidate in modeling extreme numbers of casualties resulting from terror insurgencies. While the power law distribution utilized in Clauset and Woodard (2013) to quantify the probability of high casualty-producing terror events belongs to a single parameter family, the GZD belongs to a two-parameter family of distributions and can be seen as the discrete counterpart of the Generalized Pareto Distribution (GPD), giving extra flexibility in model fitting and easily incorporating a threshold u, beyond which extreme events are modeled.

We performed preliminary studies considering casualty data from Iraq, Pakistan, Afghanistan, and India with combinations of the zero-inflated Poisson/Negative Binomial distributions for f_{ZI} and the GZD/GPD distributions for f_{EV} in Equation 2. The available data for each country was fitted to our proposed distribution with maximum likelihood estimation for the parameters θ_{ZI} , θ_{EV} , π and u. Figure 3 shows fits of f_M with either a zero inflated Negative Binomial or a zero-inflated Poisson distributions for f_{ZI} for Iraq, Pakistan, Afghanistan, and India respectively. The threshold u was considered at a few values indicated by vertical dashed lines and the combination of possible distributions/thresholds assessed via the Akaike information criterion (AIC) as reviewed by (Cavanaugh and Neath, 2019). Distributional fits measured by AICs, are shown on top of the casualty data histograms of each country. In all cases, AIC favors a ZIP distribution for f_{ZI} and a GZD form for f_{EV} , with a threshold of value u = 2. This is the mark distribution that will be adopted in the remainder of the paper.

3.3 Conditional intensity

The conditional intensity $\lambda^*(\mathbf{s}, t; \mathcal{H}_t)$ appearing in Equation (1) describes the instantaneous measure of a terror event occurring at a space-time location (\mathbf{s}, t) given historic events \mathcal{H}_t occurring before time t. The self-exciting nature of terror events renders a Hawkes type intensity Hawkes (1971), suitable for this conditional measure. The Hawkes process formulation specifically addresses attributes of self-excitation or contagion, accounting for an arbitrary event's contribution to the likelihood of future event occurrence. This flexible model aspect has provided an important basis for its broad utility and has seen application in modeling financial stock crashes and surges (Da Fonseca and Zaatour, 2014), predicting origin times and magnitudes of earthquakes (Ogata, 1988) and propagation of social media events (Rizoiu et al., 2017). To do so, the Hawkes conditional intensity takes the general form (Reinhart, 2017)

$$\lambda^*(\mathbf{s}, t; \mathcal{H}_t) = \mu^*(\mathbf{s}, t; \boldsymbol{\theta}_{\mu^*}) + \sum_{t_i < t} \phi(t - t_i, |\mathbf{s} - \mathbf{s}_i|, m_i; \boldsymbol{\theta}_{\phi}),$$
(6)

where \mathcal{H}_t denotes the history of the events and $\mu^*(\mathbf{s}, t; \boldsymbol{\theta}_{\mu^*})$ is the baseline spatio-temporal intensity of terror attacks and $\phi(\cdot)$ denotes the *triggering* function, describing the self-exciting intensity of a future event given the space-time locations and marks associated with previous attacks. Depending on the kernel $\phi(\cdot)$, the excitation may be local or have longer term effects in space and/or time (Hawkes, 1971). The terms $\boldsymbol{\theta}_{\mu^*}$ and $\boldsymbol{\theta}_{\phi}$ denote hyper-parameters associated with the baseline and triggering intensities, respectively.

In the following, we denote an individual event by its measured time, spatial location and casualty mark as $\mathbf{x}_t := (\mathbf{s}_{\mathbf{x}_t}, t, m_{\mathbf{x}_t})$, and the point pattern of events as $\mathcal{X} = \bigcup_t \mathbf{x}_t$: a realization of the *conditional* point process $X(\mathbf{s}, t; \mathcal{H}_t)$. The conditional intensity in Equation (6) describes the instantaneous probability of an arbitrary event \mathbf{x}_t belonging *either* to the background via the intensity contribution of μ^* or being triggered by a background point.

From this observation, we can consider a Hawkes process as a *branching process* (Hawkes and Oakes, 1974) of marked spatial events ordered by their time stamps. Here, events are seen to arrive either through immigration or birth. An immigrant corresponds to an event that results from the baseline intensity μ^* , whereas events that are caused by self-excitation are birthed descendants from previous events, that are excited via the triggering kernel ϕ . We note that when $\phi \equiv 0$ almost everywhere, the Hawkes process as defined in Equation (6) reduces to an inhomogeneous Poisson process with rate $\mu^*(\mathbf{s}, t; \boldsymbol{\theta}_{\mu^*})$. In this manner, background points form the centers of *inhomogeneous Poisson clusters* of triggered events drawn from the self-exciting kernel ϕ with respect to the immigrant event. The interested reader is directed to Hawkes and Oakes (1974) for more information about this Cluster-based Hawkes Process.



Figure 3: Comparisons of mark distributions fit with the GZD/GPD and ZIP/ZINB for Iraq (left, top), Pakistan (right, top), Afghanistan (left, bottom) and India (right, bottom)

Histogram of casualties for Iraq

Histogram of casualties for Pakistan

The point pattern \mathcal{X} can be therefore be succinctly written as the set union of the background \mathcal{X}_{μ^*} and triggered pattern \mathcal{X}_{ϕ} , written as

$$\mathcal{X} = \mathcal{X}_{\mu^*} \cup \mathcal{X}_{\phi}$$
$$\mathcal{X}_{\phi} = \cup_{\mathbf{x}'_t \in \mathcal{X}_{\mu^*}} \mathcal{X}_{\mathbf{x}'_t}$$

where $\mathcal{X}_{\mathbf{x}'_t}$ denotes the triggered pattern of inhomogeneous clusters from background point $\mathbf{x}'_t \in \mathcal{X}_{\mu^*}$ at arbitrary time $t \in [0, T]$.

3.3.1 Baseline intensity

The baseline intensity $\mu^*(\mathbf{s}, t; \boldsymbol{\theta}_{\mu^*})$ is intended to characterize untriggered terror events, for example, attacks that are initiated by new insurgency groups. In the simplest case, a (constant baseline) homogeneous Poisson intensity i.e. where $\mu^*(\mathbf{s}, t; \boldsymbol{\theta}_{\mu^*}) \equiv \mu^*$ describes a process in which the expected number of baseline attacks per unit area is μ^* . In general, baseline attacks appear inhomogeneous across areas of interest, and are seen to be correlated with population density, distance to nearest cities and other socio-economic factors (Zammit-Mangion et al., 2012). To deal with this inhomogeneous Poisson process

$$\mu^*(\mathbf{s},t;\boldsymbol{\theta}_{\mu^*}) = \exp(\mathcal{Z}_{\mathbf{s},t}^{\top}\boldsymbol{\theta}_{\mu^*}),$$

where $\theta_{\mu^*} \in \mathbb{R}^n$ is a vector of coefficients and $\mathcal{Z}_{s,t}$ is a vector of recorded covariates. Covariates of interest include population density, government voting statistics and number of different languages spoken at the district level with GDP, geographical distances to nearest cities and terrains used at the country level. These covariates are available from Statistical Year Books of countries of interest. In the case of Afghanistan, year books are available at https://nsia.gov.af/library. Given a background point pattern \mathcal{X}_{μ^*} , the density of the inhomogeneous Poisson intensity (with respect to the unit rate Poisson on $W \times T$) takes the form (Baddeley and Turner, 2000)

$$f_{\mu^*}(\mathcal{X}_{\mu^*};\boldsymbol{\theta}_{\mu^*}) = \prod_{\mathbf{x}'_t \in \mathcal{X}_{\mu^*}} \mu^*(\mathbf{s}_{\mathbf{x}'_t}, t; \boldsymbol{\theta}_{\mu^*}) \exp\left(-\int_{W \times T} (\mu^*(t', \mathbf{s}'; \boldsymbol{\theta}_{\mu^*}) - 1) \mathrm{d}\mathbf{s}' \mathrm{d}t'\right),\tag{7}$$

where the product is over all events with arbitrary time $t \in [0, T]$. Given \mathcal{X}_{μ^*} , the point process density in Equation (7) can be numerically maximized to determine θ_{μ^*} . Here, the integral term can be computed using the quadrature scheme described in Baddeley and Turner (2000) and permits a computationally efficient maximization algorithm.

3.3.2 Triggering intensity

In order to derive the log-likelihood associated with the triggered point pattern \mathcal{X}_{ϕ} , it is necessary to understand the baseline origin of such points within the cluster-based Hawkes formulation previously mentioned. In particular, the marginal distribution $p(\mathcal{X}_{\phi}; \boldsymbol{\theta}_{\phi})$ requires conditioning on all background events in \mathcal{X}_{μ^*} and data sequences \mathcal{D}_{ϕ} from \mathcal{X}_{ϕ} that give time-ordered clusters of triggered events originating from a background point. Following Walder and Bishop (2018), the log density of $p(\mathcal{X}_{\phi}; \boldsymbol{\theta}_{\phi})$ is shown to follow

$$\log p(\mathcal{X}_{\phi}; \boldsymbol{\theta}_{\phi}) = \sum_{\substack{d \in \mathcal{D}_{\phi} \\ u \leq v \\ \{d \ni \mathbf{x}_{s}: s \in (u, v)\} = \emptyset}} \sum_{\substack{\mathbf{x}_{t}' \in \mathcal{X}_{\mu^{*}} \\ d \in \mathbf{x}_{s}: s \in (u, v)\} = \emptyset}} \log \phi(\delta(\mathbf{x}_{t}', \mathbf{x}_{u}, \mathbf{x}_{v})) - \int_{W \times [0, T-t]} \phi(t', \mathbf{s}', m_{\mathbf{x}_{t}'}) \mathrm{d}\mathbf{s}' \mathrm{d}t',$$
(8)

where the second summand is over *subsequent* event pairs $\mathbf{x}_u, \mathbf{x}_v$ with v > u and

$$\delta(\mathbf{x}'_t, \mathbf{x}_u, \mathbf{x}_v) = (v - u, |\mathbf{s}_{\mathbf{x}_v} - \mathbf{s}_{\mathbf{x}_u}|, m_{\mathbf{x}'_t}).$$

This density utilizes the branching structure of the Hawkes' immigration-birth formation through the clusters in \mathcal{D}_{ϕ} as additional variables to consider. Since $\mathcal{D}_{\mu^*} = \mathcal{D} \setminus \mathcal{D}_{\phi}$ can be readily obtained from \mathcal{D}_{ϕ} , the self-exciting and background intensity parameters can be evaluated from their derived conditional distributions if this structure is known.

The triggering function $\phi(\cdot)$ is typically written with a known parametric form. In the one dimensional case, Hawkes (1971) proposed the triggering function to be an exponential decay function $\phi(s) = \alpha e^{-\beta s}$ with $\alpha, \beta > 0$ or power law functions of the form $\phi(s) = \frac{k}{(c+s)^p}$ for k, c, p > 0.

For this problem, however, the nuances in terror insurgence excitation and uncertainty surrounding the quantitative effect of historic events render non-parametric models to be more appealing. To show this explicitly, a simplified non-parametric Hawkes process utilized in Mohler et al. (2011), with intensity

$$\lambda^*(\mathbf{s}, t; \mathcal{H}_t) = \nu_1(t)\nu_2(\mathbf{s}) + \sum_{t_i < t} g_1(t - t_i)g_2(|\mathbf{s} - \mathbf{s}_i|),$$

was considered for exploratory data analysis. In relation to the generalized intensity function given in Equation (1), we have in this case that $\mu^*(\mathbf{s}, t; \boldsymbol{\theta}_{\mu^*}) = \nu_1(t)\nu_1(\mathbf{s})$ and $\phi(t - t_i, |\mathbf{s} - \mathbf{s}_i|, m_i) = g_1(t - t_i)g_2(|\mathbf{s} - \mathbf{s}_i|)$, where the functions ν_1, ν_2, g_1, g_2 are to be determined. Empirically based kernel density estimates of the functions ν_1, ν_2 and g_1 fitted to Afghanistan's terror events in the GTD are shown in Figure 4. It is seen here that the estimated functions for the temporal background and triggering functions are highly nonlinear and difficult to parametrize. Additionally, this approach cannot easily incorporate additional covariates $\mathcal{Z}_{\mathbf{s},t}$ into the procedure, thereby owing to crude spatial intensity estimates.

In order to better understand and characterize terror excitation, the demonstration in Figure 4 motivates our novel method, presented next, for modeling the triggering function.



Figure 4: Nonparametric Hawkes intensity fits for Afghanistan (2013-2018): Estimated background temporal $\nu_1(t)$ (left), background spatial $\nu_2(s)$ (center) and triggering temporal $g_1(\cdot)$ (right) intensities.

Specifically, we place a Gaussian Process prior over the class of all triggering functions $\phi(\cdot)$. In a similar manner to the works of Flaxman et al. (2017), the triggering kernel $\phi(t - t_i, |\mathbf{s} - \mathbf{s}_i|, m_i)$ is written non-parametrically as $\phi(\cdot) = af^2(\cdot)$, where f is a Gaussian Process and a > 0, to ensure positivity of the intensity function. The *covariance* function of f can be decomposed into its eigenvalues $\{\tilde{\lambda}_1, \tilde{\lambda}_2, \ldots, \}$ and eigenfunctions $\{e_1(\cdot), e_2(\cdot), \ldots\}$ via *Mercer's theorem* to give

$$k(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^{\infty} \tilde{\lambda}_i e_i(\mathbf{x}) e_i(\mathbf{y}).$$

The above representation implies that $f(\cdot) = \boldsymbol{\omega}^\top \mathbf{e}(\cdot)$, where $\boldsymbol{\omega} \sim \mathcal{N}(0, \tilde{\Lambda} = \text{diag}(\tilde{\lambda}_1, \tilde{\lambda}_2, \ldots))$ and $\mathbf{e}(\cdot) = [e_1(\cdot), e_2(\cdot), \ldots]$ (Zhang et al., 2018). In practice, a low-rank approximation of the above infinite sum is used to ensure computational stability.

Typical examples of the covariance function $k(\mathbf{x}, \mathbf{y})$ include the Squared exponential $k(\mathbf{x}, \mathbf{y}) = \exp\left(-\|\mathbf{x} - \mathbf{y}\|^2/2l^2\right)$ with known $\tilde{\lambda}_i, e_i(\mathbf{x})$ with respect to the Gaussian measure, the Rational quadratic, $k(\mathbf{x}, \mathbf{y}) = \left(1 + \|\mathbf{x} - \mathbf{y}\|^2/2\alpha_s l^2\right)^{-\alpha}$ and the Periodic $k(\mathbf{x}, \mathbf{y}) = 1 + \sum_{m=1}^{\infty} 2\cos(2\pi m(\mathbf{x} - \mathbf{y}))/(2\pi m)^{2s}$, on [0, 1] with known $\tilde{\lambda}_i, e_i(\mathbf{x})$ with respect to Lebesgue measure (Flaxman et al., 2017). Due to this characterization, not all covariance functions have known Mercer expansions.

To deal with this, and to allow full flexibility in the choice of a covariance function for f, we utilize the work of Flaxman et al. (2017). Specifically, for a positive definite kernel function $k : S \times S \to \mathbb{R}$ over a non-empty domain S, a unique reproducing kernel Hilbert space (RKHS) \mathcal{H}_k exists and defines a space of functions $f : S \to \mathbb{R}$ from which f can be represented by an inner product $f(\mathbf{x}) = \langle f, k(\mathbf{x}, \cdot) \rangle_{\mathcal{H}_k}$. Using this set-up, for *triggering* data $\mathbf{x}_i \in \mathcal{X}_{\phi}$ arising from the event $\mathbf{x}_i \in \mathcal{X}_{\mu^*}$, the function f satisfies

$$f = \arg\min_{f' \in \mathcal{H}_k} \left\{ -\sum_i \log(af'^2(t_j - t_i, |\mathbf{s}_i - \mathbf{s}_j|, m_i)) + a \int_{W \times T} f'^2(t_j - t, |\mathbf{s} - \mathbf{s}_j|, m_i) \, \mathrm{d}\mathbf{s} \, \mathrm{d}t + \gamma \|f'\|^2 \right\},$$
(9)

including a regularization term γ that corresponds to the squared Hilbert space norm of f. The integral term of 9 does not however guarantee a solution with covariance $k(\mathbf{x}, \mathbf{y})$, stemming from the possibility of non-explicit Mercer expansions belonging to the most flexible class of covariances. To deal with this, Flaxman et al. (2017) constructs an alternative kernel $\tilde{k}(\mathbf{x}, \mathbf{y})$, which can be computed directly if the covariance function has an explicit Mercer expansion

with respect to the Lebesgue measure. In general, using the covariance

$$\tilde{k}(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^{K} \frac{\eta_i}{a\eta_i + \gamma} \hat{e}_i(\mathbf{x}) \hat{e}_i(\mathbf{y}),$$

leads to the existence of a solution f.

 $f(\cdot) = \boldsymbol{\omega}^{\top} \mathbf{\hat{e}}(\cdot),$

In the case when Mercer's expansion is unknown, the adjusted eigenvalues $\eta_i/a\eta_i + \gamma$ and eigenvectors \hat{e}_i of kernel \tilde{k} are sampled via the uniform sampling approach described in detail in Flaxman et al. (2017) and briefly described here. By scaling $W = (0, 1)^2$, T = 1, M = (0, 1), a uniform form grid u over all possible values of s, t, m can be constructed and the covariance function K_{uu} computed on the matrix. Implementing *Lanczos iterations* invokes the p highest eigenvalues $\tilde{\lambda}_i$ and eigenvectors e_i^u to be found of the truncated eigen-decomposition of K_{uu} . Finally, after determining the covariance K_{xu} between data x and u, $\tilde{k}(x, y)$ can be evaluated explicitly by computing

$$\eta_i = \frac{\tilde{\lambda}_i}{p}, \qquad \hat{e}_i(\mathbf{x}) = \frac{\sqrt{p}}{\tilde{\lambda}_i} K_{\mathbf{x}\mathbf{u}} e_i^{\mathbf{u}}$$

The above representation thereby allows computation of $\phi(\cdot) = af^2(\cdot)$, where

$$\hat{\mathbf{e}}(\cdot) = [\hat{e}_1(\cdot) \dots \hat{e}_p(\cdot)], \qquad (10)$$

$$\boldsymbol{\omega} \sim \mathcal{N}\left(0, \tilde{\Lambda} = \operatorname{diag}\left(\frac{\eta_1}{a\eta_1 + \gamma}, \dots, \frac{\eta_p}{a\eta_p + \gamma}\right)\right), \qquad \boldsymbol{\omega} \in \mathbb{R}^p.$$
(11)

Here, the unknown parameter set $\theta_{\phi} = \{\theta_K, \omega, a, \gamma\}$, with θ_K containing all parameters of the intended covariance structure, is to be determined.

3.4 MCMC estimation

Using $\lambda(\mathbf{s}, t, m; \mathcal{H}_t)$ as given by Equation (1), the log-likelihood function of data \mathcal{X} given the unknown parameter sets $\boldsymbol{\theta}_1 = \{\pi, \boldsymbol{\theta}_{ZI}, \boldsymbol{\theta}_{EV}\}$ and $\boldsymbol{\theta}_2 = \{\boldsymbol{\theta}_{\mu^*}, \boldsymbol{\theta}_{\phi}\}$ is

$$\ell(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2; \mathcal{X}) = \sum_{i=1}^n \left[\log \left(\mathcal{Z}_{\mathbf{s}, t}^\top \boldsymbol{\theta}_{\mu^*} + \sum_{t_j < t_i} \phi(t_i - t_j, |\mathbf{s}_i - \mathbf{s}_j|, ;m_i) \right) + \log(f_M(m_i; \boldsymbol{\theta}_1, \mathcal{Z}_m(\mathbf{s}, t))) \right] - \int_{W \times T} \lambda^*(\mathbf{s}, t | \mathcal{H}_t, \mathcal{Z}_{\mathbf{s}, t}, \boldsymbol{\theta}_2) \, \mathrm{d}\mathbf{s} \, \mathrm{d}t.$$
(12)

Since the intensity function comes in a separable form between the mark and spatio-temporal parameters, we can estimate θ_1 and θ_2 in parallel. The parameter set θ_2 can either be estimated through maximum likelihood methods or via a Metropolis Hastings MCMC algorithm (MCMC). We elect here to use MCMC which samples θ_1 from the mark vector $\mathbf{m} = (m_1, m_2, \dots, m_n)$ via

$$p(\boldsymbol{\theta}_1; \mathbf{m}, \mathcal{Z}_{\mathbf{m}}(\mathbf{s}, t)) \propto \pi(\boldsymbol{\theta}_1) \prod_{i=1}^n f_M(m_i; \boldsymbol{\theta}_1, \mathcal{Z}_{m_i}(\mathbf{s}, t)),$$

with the prior specifications

$$\pi(\boldsymbol{\theta}_{\beta}), \pi(\boldsymbol{\theta}_{\xi}), \pi(\boldsymbol{\theta}_{\sigma}) \sim \mathcal{N}(0, I)$$
$$\log\left(\frac{1-\alpha}{\alpha}\right), \log\left(\frac{1-\pi_{M}}{\pi_{M}}\right) \sim \mathcal{N}(0, 1).$$

We perform estimation of θ_2 within a Bayesian Hierarchical framework using a hybrid Metropolis-within-Gibbs particle MCMC algorithm. This utilizes the aforementioned branching structure of the Hawkes' immigration-birth formation as an additional variable, from which both self-exciting and background intensity parameters can be updated using the conditional distributions we will henceforth derive.

Specifically, the branching variable, denoted $\mathcal{B} = \{C_i\}_i$ consists of inhomogeneous Poisson clusters containing information of all *independent* events $\mathcal{X}_{\mu^*} \in \mathcal{X}$ immigrating from the background intensity function and the triggered events $\mathcal{X}_{\phi} \in \mathcal{X} \setminus \mathcal{X}_{\mu^*}$ that are birthed from them. For example, having cluster $C_1 = (\mathbf{x}_{t_0}, \mathbf{x}_{t_1}, \mathbf{x}_{t_2})$ where $t_0 < t_1 < t_2$, would imply that background event \mathbf{x}_{t_0} triggers subsequent events \mathbf{x}_{t_1} and \mathbf{x}_{t_2} and that this structure belongs to \mathcal{B} . From herein, we denote the background event triggering arbitrary cluster C as $C_0 \in \mathcal{X}_{\mu^*}$ and the births as $C_{>0}$.

In order to derive Gibbs conditional distributions parameters in θ_2 , the branching structure \mathcal{B} must be sampled from. Sampling \mathcal{B} from its posterior $p(\mathcal{B}; \mathcal{X}, \theta_2)$ consists of computing probabilities that an arbitrary event \mathbf{x}_{t_i} is either an immigrant or birth. To do so, we define the probabilities p_{ij} for all $t_j < t_i$ as the probability event \mathbf{x}_{t_i} is triggered by event \mathbf{x}_{t_j} and p_{i0} as the probability event \mathbf{x}_{t_i} is birthed from the background. Using the intensity in Equation (1), we have that

$$p_{i0} = \mu^*(\mathbf{s}_i, t_i; \boldsymbol{\theta}_{\mu^*}) \left(\mu^*(\mathbf{s}_i, t_i; \boldsymbol{\theta}_{\mu^*}) + \sum_{t_k < t_i} \phi(t_i - t_k, |\mathbf{s}_i - \mathbf{s}_k|, m_i) \right)^{-1}$$
(13)

$$p_{ij} = \phi(t_i - t_j, |\mathbf{s}_i - \mathbf{s}_j|, m_i) \left(\mu^*(\mathbf{s}_i, t_i; \boldsymbol{\theta}_{\mu^*}) + \sum_{t_k < t_i} \phi(t_i - t_k, |\mathbf{s}_i - \mathbf{s}_k|, m_i) \right)^{-1} \quad t_j < t_i.$$
(14)

These probabilities can be determined in each MCMC iteration given the current value of θ_2 thereby enabling the branching structure \mathcal{B} to be sampled by sampling the clusters from these probability vectors.

Once \mathcal{B} is updated, the algorithm then proceeds to sample from the background posterior $p(\theta_{\mu^*}; \mathcal{X}, \mathcal{B}, \theta_2)$ and self-exciting posterior $p(\theta_{\phi}; \mathcal{X}, \mathcal{B}, \theta_2)$.

For the background density, we note that using the likelihood function given by Equation (7), we may update θ_{μ^*} using a Metropolis-Hastings step which samples from

$$p(\boldsymbol{\theta}_{\mu^*}; \mathcal{X}, \mathcal{B}, \mathcal{Z}_{\mathbf{s},t}) \propto f_{\mu^*}(\mathcal{X}_{\mu^*} | \boldsymbol{\theta}_{\mu^*}) \pi(\boldsymbol{\theta}_{\mu^*}) \\ \pi(\boldsymbol{\theta}_{\mu^*}) \sim \mathcal{N}(0, I),$$

where \mathcal{X}_{μ^*} can be determined from \mathcal{X} using \mathcal{B} . We note that Gaussian priors are used where applicable for controlling the efficacy of the subsequent MCMC algorithm(s) in accepting a theoretically optimal rate of samples.

We note here that in the case of a *constant* intensity θ_{μ^*} resulting from a single constant background covariate, that a Poisson-Gamma conjugacy structure may be used. Specifically, since $|\mathcal{X}_{\mu^*}| \sim \text{Poisson}(\theta_{\mu^*})$ can be determined from \mathcal{B} , we place

$$\begin{split} p(\theta_{\mu^*}; \mathcal{X}, \mathcal{B}, \boldsymbol{\theta}_{\phi}) &\propto p(\mathcal{X}_{\mu^*}; \theta_{\mu^*}, \boldsymbol{\theta}_{\phi}) \pi(\theta_{\mu^*}) \\ &\pi(\theta_{\mu^*}) \sim \operatorname{Gamma}(1, 1) \\ &\pi(\theta_{\mu^*}; |\mathcal{X}_{\mu^*}|) \sim \operatorname{Gamma}(|\mathcal{X}_{\mu^*}| + 1, 2|W \times T|). \end{split}$$

For the triggering parameters, we note that given \mathcal{B} , the log density given in Equation (8) can be reduced to

$$\log p(\mathcal{X}_{\phi}; \boldsymbol{\theta}_{\phi}, \mathcal{X}, \mathcal{B}) = \sum_{\substack{C \in \mathcal{B} \\ u < v \\ \{C \ni \mathbf{x}_{s}: s \in (u, v)\} = \emptyset}} \sum_{\substack{u < v \\ \{c \in \mathbf{x}_{s}: s \in (u, v)\} = \emptyset}} \log a(\boldsymbol{\omega}^{\top} \hat{\mathbf{e}}(\delta(C_{0}, \mathbf{x}_{u}, \mathbf{x}_{v})))^{2} - a \int_{\substack{W \times \\ [0, T - t_{C_{0}}]}} (\boldsymbol{\omega}^{\top} \hat{\mathbf{e}}(t', \mathbf{s}', m_{C_{0}}))^{2} \mathrm{d}\mathbf{s}' \mathrm{d}t'.$$
(15)

Given this, updates for the fixed parameters in θ_{ϕ} can be determined via

$$p(\boldsymbol{\theta}_{\phi}; \mathcal{X}, \mathcal{B}) \propto p(\mathcal{X}_{\phi}; \boldsymbol{\theta}_{\phi}, \mathcal{X}, \mathcal{B}) \pi(\boldsymbol{\theta}_{\phi})$$
$$\boldsymbol{\omega}|a, \gamma \sim \mathcal{N}\left(0, \tilde{\Lambda} = \operatorname{diag}\left(\frac{\eta_{1}}{a\eta_{1} + \gamma}, \dots, \frac{\eta_{K}}{a\eta_{K} + \gamma}\right)\right)$$
$$\log(a), \log(\gamma) \sim \mathcal{N}(0, 1).$$

In particular, the log posterior $p(\boldsymbol{\omega}; \mathcal{X}, \mathcal{B}, \boldsymbol{\theta}_{\phi})$ is given by

$$\log p(\boldsymbol{\omega}; \mathcal{X}, \mathcal{B}, \boldsymbol{\theta}_{\phi} \setminus \boldsymbol{\omega}) = -\frac{K}{2} \log 2\pi - \frac{1}{2} \left(\log |\tilde{\Lambda}| + \boldsymbol{\omega}^{\top} \tilde{\Lambda}^{-1} \boldsymbol{\omega} \right) + \sum_{\substack{C \in \mathcal{B} \\ \mathbf{x}_{u}, \mathbf{x}_{v} \in C_{>0} \\ u < v \\ \{C \ni \mathbf{x}_{s}: s \in (u, v)\} = \emptyset}} \sum_{\substack{C \in \mathcal{B} \\ u < v \\ \{C \ni \mathbf{x}_{s}: s \in (u, v)\} = \emptyset}} \log a(\boldsymbol{\omega}^{\top} \hat{\mathbf{e}}(\delta(C_{0}, \mathbf{x}_{u}, \mathbf{x}_{v})))^{2} - a\boldsymbol{\omega}^{\top} \left(\int_{\substack{W \times \\ [0, T - t_{C_{0}}]}} E(t', \mathbf{s}', m_{C_{0}}) \mathrm{d}\mathbf{s}' \mathrm{d}t' \right) \boldsymbol{\omega},$$
(16)

where for any $\mathbf{s}, t, m \in W \times [0, T] \times \mathbb{N}_0$, the outer product of $\hat{\mathbf{e}}(t, \mathbf{s}, m)$ is defined via

$$E(t, \mathbf{s}, m) = \hat{\mathbf{e}}(t, \mathbf{s}, m) \hat{\mathbf{e}}^{\top}(t, \mathbf{s}, m).$$

We note here that computing the integral term in Equation (16) within the algorithm corresponds to

$$\int_{[0,T-t_{C_{0}}]} E(t',\mathbf{s}',m_{C_{0}}) d\mathbf{s}' dt' = \int_{[0,T-t_{C_{0}}]} \hat{\mathbf{e}}(t',\mathbf{s}',m_{C_{0}}) \hat{\mathbf{e}}(t',\mathbf{s}',m_{C_{0}})^{\top} d\mathbf{s}' dt'$$

$$= \int_{[0,T-t_{C_{0}}]} \left[\frac{\sqrt{p}}{\tilde{\lambda}_{1}} K_{\mathbf{x}\mathbf{u}} e_{1}^{\mathbf{u}} \cdots \frac{\sqrt{p}}{\lambda_{p}} K_{\mathbf{x}\mathbf{u}} e_{p}^{\mathbf{u}} \right] \left[\frac{\sqrt{p}}{\tilde{\lambda}_{1}} K_{\mathbf{x}\mathbf{u}} e_{1}^{\mathbf{u}} \cdots \frac{\sqrt{p}}{\lambda_{p}} K_{\mathbf{x}\mathbf{u}} e_{p}^{\mathbf{u}} \right]^{\top} d\mathbf{s}' dt'$$

$$\approx \frac{1}{M} \sum_{i=1}^{M} \left[\frac{\sqrt{p}}{\tilde{\lambda}_{1}} K_{\mathbf{x}'_{i}\mathbf{u}} e_{1}^{\mathbf{u}} \cdots \frac{\sqrt{p}}{\lambda_{p}} K_{\mathbf{x}'_{i}\mathbf{u}} e_{p}^{\mathbf{u}} \right] \left[\frac{\sqrt{p}}{\tilde{\lambda}_{1}} K_{\mathbf{x}'_{i}\mathbf{u}} e_{1}^{\mathbf{u}} \cdots \frac{\sqrt{p}}{\lambda_{p}} K_{\mathbf{x}'_{i}\mathbf{u}} e_{p}^{\mathbf{u}} \right]^{\top}, \quad (17)$$

where the M importance samples or particles satisfy $\mathbf{x}'_i = (t_i(1-t_{C_0}), |\mathbf{s}_{C_0} - \mathbf{s}_i|, m_{C_0})$, \mathbf{s}_i, t_i are uniformly distributed on $W \times [0, T]$ and \mathbf{u} is a fixed uniform grid in each iteration. Although Equation (17) denotes an approximation to the true value of the integral, we note that using a renewed particle set $\{\mathbf{x}'_i\}_{i=1}^M$ and even a renewed uniform grid \mathbf{u} at each iteration will yield an *unbiased* estimate of the true integral, for any M. Within an MCMC framework, it is sufficient to compute an unbiased estimate of the log posterior(s) of interest (Christophe et al., 2010) when computing Metropolis-Hastings acceptance ratios, resulting in matching MCMC convergence properties while limiting the potentially large particle sample size usually required to achieve a favorable approximation.

Given the varying covariance structures that can be placed on the Gaussian process prior and the unbiased update form for $\phi(\cdot)$, our approach updates the parameter set θ_{ϕ} with a Metropolis-Hastings kernel. However, up to an arbitrary rank p, studying the posterior distribution of Gaussian process parameter ω within this framework is challenging due to potentially high dimensions and variability resulting from the modeling of such nuanced datasets. On the other hand, the ability to differentiate (16) motivates an approach for updating ω using a Hamiltonian Monte Carlo (HMC) step, used extensively in the literature to sample from challenging target distributions (Neal, 2011; Betancourt, 2018). Specifically, by defining $U(\omega) = -\log p(\omega; \mathcal{X}, \mathcal{B}, \theta_{\phi} \setminus \omega)$ termed as the *potential energy* and the momentum vector $\rho \sim \mathcal{N}(0, \Sigma)$ providing the *kinetic energy* $-\log p(\rho; \Sigma)$, HMC is used to jointly sample ω and ρ using the *Hamiltonian* or the joint negative log posterior given by

$$H(\boldsymbol{\rho}, \boldsymbol{\omega}) = U(\boldsymbol{\omega}) - \frac{1}{2} \boldsymbol{\rho}^{\top} \Sigma^{-1} \boldsymbol{\rho}.$$

The joint system (ω, ρ) consisting of current parameter vector ω and momentum ρ , initially sampled from its Gaussian prior at each iteration, is evolved via numerical integration of the resulting Hamilton system of equations (Betancourt, 2018). This requires computation of $\frac{\partial U(\omega)}{\partial \omega}$, which can be given by direct differentiation of Equation (16) as

$$\frac{\partial U(\boldsymbol{\omega})}{\partial \boldsymbol{\omega}} = \tilde{\Lambda}^{-1} \boldsymbol{\omega} + 2 \sum_{\substack{C \in \mathcal{B} \\ u < v \\ \{C_{>0} \ni \mathbf{x}_s: s \in (u,v)\} = \emptyset}} \sum_{\substack{W \times \\ [0,T-t_{C_0}]}} E(t', \mathbf{s}', m_{C_0}) \mathrm{d}\mathbf{s}' \mathrm{d}t' \right) \boldsymbol{\omega} - \frac{\hat{\mathbf{e}}(\delta(C_0, \mathbf{x}_u, \mathbf{x}_v))}{\boldsymbol{\omega}^{\top} \hat{\mathbf{e}}(\delta(C_0, \mathbf{x}_u, \mathbf{x}_v))} \right].$$
(18)

Once solved to find $(\boldsymbol{\omega}^*, \boldsymbol{\rho}^*)$, the solution is accepted with (Metropolis-Hastings') probability $\min(1, \exp(H(\boldsymbol{\rho}, \boldsymbol{\theta}) - H(\boldsymbol{\rho}^*, \boldsymbol{\theta}^*)))$ to account for numerical integration errors that are typically encountered in practice.

4 **Results**

We first validate our inference method on simulated data to demonstrate its accuracy in estimating all unknown parameters of the model. We then apply the method to study and predict terror insurgencies in Afghanistan based upon event data detailed in the Global Terrorism Database between 2013 and 2018.

4.1 Simulated study

To study the accuracy of the developed MCMC inferential algorithm on the conditional intensity estimation, we provide parameter estimates from a simulated example. In this study, 768 events were simulated on the unit hypercube from the self-exciting point process with constant background rate $\mu^* = 50$ and with triggering prior covariance

$$K(\mathbf{x}, \mathbf{y}) = \exp\left(-\frac{|t_{\mathbf{x}} - t_{\mathbf{y}}|^2}{2l_t^2}\right) \left(1 + \frac{||(\mathbf{s}_x, m_x) - (\mathbf{s}_y, m_y)||^2}{2\alpha_s l_s}\right)^{-\alpha_s},$$
(19)



Figure 5: Posterior plots for all conditional intensity parameters estimated from our algorithm from simulated data. Here, true values (blue) are shown alongside estimated values (red) and 95% credible intervals (black, dashed).

where $\lambda_t = 0.3, \lambda_s = 1, \alpha = 1$ and $a = 1, \gamma = 0.1$. The background points were simulated by first generating $\mu' \sim \text{Poi}(\mu^*)$ and then uniformly sampling μ' points over $W \times T \times M$. Second, for each background point $\mathbf{x}_c := (t_c, \mathbf{s}_c, m_c)$, inhomogeneous Poisson clusters with intensity $\phi(t - t_c, |\mathbf{s} - \mathbf{s}_c|, m_c)$ for $t > t_c$ were generated through rejection sampling that approximated the integrated intensity over the domain via importance weights. Last, the marks were sampled uniformly on (0, 1).

Using this methodology, the particle MCMC algorithm was run to obtain 25000 samples after an initial burn in of 2000. Posterior parameter estimates for all unknown parameters in the spatio-temporal conditional intensity given in Equation (1) were inferred by thinning of the Markov Chain at every tenth sample. Posterior distributions along with posterior modes and 95% credible intervals for all parameters of interest are given in Figure 5. The true values (blue) and the posterior mode estimates (red) are also shown in the figure. We notice that the most informative posterior distributions correspond to the parameters λ_t , α and a also having posterior modes close to the true values. On the other hand, while the posterior mode of γ is close to the true value, its distribution shows high variability with heavy tails across its interval. The credible interval for μ^* is much wider and contains its true value, caveatted by the fact that the simulated background rate was estimated from a single realization of a Poisson(μ^*) variate, and is therefore subject to high posterior variance. It is worth noting that for a relatively small point pattern, all simulated parameters are shown to lie within their 95% posterior credible intervals.

4.2 Modeling Afghan terror (2013-2018)

We now detail model specificities and results used to portray terror insurgency patterns in Afghanistan between 2013 and 2018.

4.2.1 Data preprocessing

The raw event data from the GTD for this study contains some missing values in the form of unknown event times (as days), locations (as latitude/longitude values) and covariates. To deal with this issue, the data used was preprocessed and missing values imputed through the Multivariate Imputation via Chained Equations (MICE) approach (Schafer, 1997) to generate a full dataset. Further, times were jittered uniformly across the day to become a continuous variable needed in the model. Latitude and longitude values were also jittered using a Gaussian distribution centered around their given coordinate with a small standard deviation, thereby adding some randomness to event data determined to the nearest city of true event locations. To be able to use the aforementioned methodology in the computation of the conditional intensity, we standardize the times $t_i \in [0, T]$, locations $s_i \in W$ and marks $m_i \in M$ to the unit hypercube and post-scale the resulting spatio-temporal intensity by the area $T \times W$.

4.2.2 Model fitting

In this study, over 10^4 terror events were collected, data preprocessed and the analysis run with our model. To be able to incorporate terror patterns that show dependence between spatial locations and casualties, an initial MCMC run was used to represent the covariance for $f(\cdot)$ as the space-time separable function given in Equation (19), additionally positing a shorter relative time scale for terror excitation. When estimated, however, prohibitively large values for α_s deemed the covariance function

$$K(\mathbf{x}, \mathbf{y}) = \exp\left(-\frac{|t_{\mathbf{x}} - t_{\mathbf{y}}|^2}{2l_t^2}\right) \exp\left(\frac{||(\mathbf{s}_x, m_x) - (\mathbf{s}_y, m_y)||^2}{2l_s^2}\right),$$

more suitable to be used for this dataset.

The MCMC algorithm enabling estimation of the conditional intensity and mark parameters was run under these specifications. For both MCMC algorithms, 50000 samples were computed after an initial burn in of 2000 and parameters inferred by thinning of the Markov Chain at every tenth sample.

To account for the dependency of parameters θ_{ZI} and θ_{EV} on potential mark covariates $\mathcal{Z}_m(s,t)$, we write the parameters of f_M as functions of $\mathcal{Z}_m(\mathbf{s},t)$ in the following way

$$\beta(\mathbf{s},t) = \exp(\mathcal{Z}_m(\mathbf{s},t)^\top \boldsymbol{\theta}_\beta) \qquad \xi(\mathbf{s},t) = \exp(\mathcal{Z}_m(\mathbf{s},t)^\top \boldsymbol{\theta}_\xi) \qquad \sigma(\mathbf{s},t) = \exp(\mathcal{Z}_m(\mathbf{s},t)^\top \boldsymbol{\theta}_\sigma), \tag{20}$$

where $\boldsymbol{\theta}_{ZI} = \{\alpha, \boldsymbol{\theta}_{\beta}\}$ and $\boldsymbol{\theta}_{EV} = \{\boldsymbol{\theta}_{\xi}, \boldsymbol{\theta}_{\sigma}\}.$

For the mark distribution given by Equation (2), the hyper-parameters θ_{ZI} and θ_{EV} were estimated using the forms of the parameters given in Equation (20). Here, the general vector of covariates is chosen as

$$\mathcal{Z}_m(\mathbf{s},t)^{\top} = \begin{pmatrix} 1 & t & \operatorname{pop}_{\mathbf{c}_{\mathbf{s}},\tilde{y}} \exp(-a\|\mathbf{s} - \mathbf{c}_{\mathbf{s}}\|_2) \end{pmatrix},$$

where $\exp(-a \|\mathbf{s} - \mathbf{c_s}\|_2)$ denotes an inverse distance measure from spatial point s to its nearest city $\mathbf{c_s}$, $a \in \mathbb{R}$ and $\operatorname{pop}_{\mathbf{c_s}, \tilde{y}}$ is the raw population estimate at $\mathbf{c_s}$ in year \tilde{y} , utilized to inform the number of casualties occurring within proximity to denser populated areas.

While all three parameter sets θ_{β} , θ_{ξ} and θ_{σ} were tested using this covariate form, anticipated over-parameterizations invoked the Deviance information criterion (DIC) to be computed for parameter combinations that excluded certain parameters conveying posterior modes close to zero. By defining the *deviance* function as (Spiegelhalter et al., 2002)

$$D(\boldsymbol{\theta}_1) = -2\log p(\boldsymbol{\theta}_1; \mathbf{m}, \mathcal{Z}_{\mathbf{m}}(\mathbf{s}, t)),$$

the DIC is given by

$$DIC = 2\overline{D(\boldsymbol{\theta}_1)} - D(\hat{\boldsymbol{\theta}}_1),$$

where $\hat{\theta}_1$ denotes the modal value of the parameter vector θ_1 and $\overline{D(\cdot)}$ denotes the mean deviance. This criterion is particularly useful since the first expression can be easily calculated as the average deviance over the obtained MCMC samples, and the second as the deviance evaluated at their corresponding MAP estimates. Similarly to the well known Bayesian Information Criterion (BIC), the model with the smallest DIC value is favored. Using this criterion, we found the most suitable model regulating the mark distribution, as shown in Equations (2) and (20), is parametrized via

$$\beta(\mathbf{s},t) = \exp(\boldsymbol{\theta}_{\beta,1} + \boldsymbol{\theta}_{\beta,t}t + \boldsymbol{\theta}_{\beta,c}\mathrm{pop}_{\mathbf{c},\tilde{y}}\exp(-a_{\beta}\|\mathbf{s} - \mathbf{c}_{\mathbf{s}}\|_{2}))$$
(21)

$$\xi(\mathbf{s},t) = \exp(\boldsymbol{\theta}_{\xi,1} + \boldsymbol{\theta}_{\xi,t}t)$$
(22)

$$\sigma(\mathbf{s},t) = \exp(\boldsymbol{\theta}_{\sigma,1} + \boldsymbol{\theta}_{\sigma,t}t + \boldsymbol{\theta}_{\sigma,c}\mathrm{pop}_{\mathbf{c},\tilde{y}}\exp(-a_{\sigma}\|\mathbf{s} - \mathbf{c}_{\mathbf{s}}\|_{2})).$$
(23)

For casualties over the space-time unit hypercube, this parameterization resulted in the posterior estimates and 95% credible intervals presented in Table 1.

Parameter	Estimate	95% CI	Parameter	Estimate	95% CI
π_M	0.415	(0.406, 0.425)	$\theta_{\mathcal{E},1}$	-0.744	(-0.904, -0.570)
α	0.131	(0.099, 0.149)	$\boldsymbol{\theta}_{\mathcal{E},t}^{r}$	0.311	(-0.136, 0.712)
$oldsymbol{ heta}_{eta,1}$	0.648	(0.542, 0.774)	$\theta_{\sigma,1}$	1.233	(1.150, 1.319)
$oldsymbol{ heta}_{eta,t}$	-0.539	(-0.815, -0.256)	$oldsymbol{ heta}_{\sigma,t}$	0.953	(0.763, 1.180)
$\boldsymbol{ heta}_{eta,c}$	-0.507	(-0.687, -0.324)	$oldsymbol{ heta}_{\sigma,c}$	0.202	(0.066, 0.368)
a_{eta}	-0.126	(-1.741, 1.971)	a_{σ}	0.393	(-1.728, 2.033)

Table 1: Posterior estimates and 95% credible intervals for parameters in $\beta(\mathbf{s}, t), \xi(\mathbf{s}, t)$ and $\sigma(\mathbf{s}, t)$ (see Equations (21-23)) of the mixed mark distribution given in Equation (2).

For the self-exciting model, the baseline intensity is assumed to have the form $\mu^*(\mathbf{s}, t; \boldsymbol{\theta}_{\mu^*}) = \exp(\mathcal{Z}_{\mathbf{s},t}^{\top} \boldsymbol{\theta}_{\mu^*})$, where the vector of covariates

 $\mathcal{Z}_{\mathbf{s},t}^{\top} = \left(\begin{array}{cccc} 1 & x_{\mathbf{s}} & y_{\mathbf{s}} & x_{\mathbf{s}}^2 & y_{\mathbf{s}}^2 & t \end{array} \text{Pop.dens}_{\mathbf{s},t} \quad \text{Alt}_{\mathbf{s}} \quad \text{Lang}_{\mathbf{s}} \quad \text{Elect.opp}_{\mathbf{s},t} \end{array}\right),$

includes a log-quadratic contribution of the locations to add extra spatial curvature to the resulting intensity. Further, population density Pop.dens_{s,t}, altitude Alt_s, number of languages spoken Lang_s and the density of population supporting the opposing government Elect.opp_{s,t}, at location s and time t (where applicable) are included as measurable variables affecting the rate of insurgency. Population and altitude covariates have been used successfully to study events detailed in the Afghan war dairy (Zammit-Mangion et al., 2012), and here extended to account for continuous spatial variables by exploiting the absolute distance to the nearest city within a province where the covariable has been measured. Further, language distribution and density of population supporting the opposing government were added to account for the effects of ethnic diversity and Taliban repression on civilians to control the country's governing (Tuerk, 2019). Other covariates including density of civil workers, yearly district revenue and population density in support of the governing party that were initially tested, invoked poor fits due to missing values and were therefore subsequently removed from the model. Unknown covariates, including those measured *yearly* such as population densities, were linearly interpolated to create a full dataset of values over continuous time. Covariates were found from https://nsia.gov.af/library and https://www.iec.org.af/results/en/home.

While the baseline kernel parameters θ_{μ^*} were updated using a standard Metropolis-Hastings with a Gaussian proposal adjusted to conceive an approximate 23% acceptance rate, sampling the triggering parameters $\theta_{\phi} \setminus \omega$ in a similar fashion resulted in poor mixing due to high correlations induced by the Gaussian process formulation. Further, formulating a similar HMC algorithm used to update ω anticipated analytical challenges due to heavy computations of the log likelihood derivative. To deal with this, we implemented the Adaptive Metropolis (AM) algorithm which is able to exploit the inherent correlation structure between the parameters of $\theta_{\phi} \setminus \omega$. This algorithm constructs a random walk proposal centered at the previous sampled value and using the empirical covariance of previously accepted samples. A further description of this sampler, including a proof detailing its ergodicity, can be found in Haario et al. (2001).

For marked spatio-temporal events over the unit space-time hypercube, inference of θ_{ϕ} utilizing the AM algorithm resulted in the posterior densities presented in Figure 6. The correlation between l_t and l_s appearing in the Gaussian process covariance $K(\mathbf{x}, \mathbf{y})$ utilized for this dataset is shown by their similarly shaped heavy-tailed posteriors, also highlighting nuanced spatio-temporal insurgency patterns in Afghanistan between 2013-2018. On the other hand, the lower posterior mode of $\log(a)$ in comparison to θ_{μ^*} emphasizes the importance of the baseline intensity and its regulating covariables in measuring insurgent events, primarily arising in densely populated areas and districts with low civilian support for insurgent repression. Nonetheless, the relatively high magnitude of triggering events reflected by $\log(a)$ confirms succinct self-excitation patterns across varying scales in space-time, reflecting organized terror events by such groups invoking violence for insurgency success. For the conditional intensity, posterior estimates and 95% credible intervals of θ_{μ^*} and θ_{ϕ} are presented in Table 2.

Results using our methodology and posterior parameter estimates are shown in Figure 7, where the space-time intensity is computed using importance samples of the time variable across the year, therefore presenting average yearly intensities. This intensity is then transformed back to its original units of the observation windows $T \times W \times M$ and shown on the log-scale. Here, the spatio-temporal intensities of attacks leading to more than 20 casualties across the region between 2013 and 2018 are shown, along with recorded events. In line with (Zammit-Mangion et al., 2012), we see higher intensities along areas targeted by terrorists along the country's "ring-road" which visits Kabul, Kandahar, Herat and Balkh, as shown in Figure 8. An increase in extreme attacks resulting from such activity suggests the gradual insurgence in districts that follow this road.



Figure 6: Posterior density plots of the triggering intensity parameters θ_{ϕ} , with the posterior modes (red) shown with 95% credible intervals (black).

Parameter	Estimate	95% CI	Parameter	Estimate	95% CI
$oldsymbol{ heta}_{\mu^*,1}$	8.413	(7.284, 8.523)	$\log(l_t)$	-1.339	(-3.760, 7.729)
$oldsymbol{ heta}_{\mu^*,x_s}$	2.830	(2.343156, 3.812476)	$\log(l_s)$	0.1221	(-2.042, 27.586)
θ_{μ^*,y_s}	8.140	(7.705, 10.239)	$\log(a)$	4.937	(3.55, 7.811)
$\boldsymbol{ heta}_{\mu^*,x_2^2}$	-1.862	(-2.712, -1.078)	$\log(\gamma)$	-5.867	(-6.628, 6.716)
$\boldsymbol{\theta}_{\mu^*,y_a^2}$	-8.315	(-10.536, -7.821)	$oldsymbol{\omega}_1$	-0.0061	(-0.170, 0.064)
$\boldsymbol{\theta}_{\mu^*,t}$	-0.190	(-0.392, -0.037)	$oldsymbol{\omega}_2$	-0.0161	(-0.225, 0.059)
$\theta_{\mu^*, \text{Pop.dens}_{e_t}}$	1.235	(1.152, 1.843)	$oldsymbol{\omega}_3$	-0.0068	(-0.094, 0.097)
$\hat{\theta}_{\mu^*, \text{Alt}_s}$	-1.196	(-1.402, -1.094)	$oldsymbol{\omega}_4$	-0.0015	(-0.164, 0.021)
$\boldsymbol{ heta}_{\mu^*,\mathrm{Lang}_2}$	0.047	(0.036, 0.103)	$oldsymbol{\omega}_5$	-0.0004	(-0.208, 0.028)
$\theta_{\mu^*, \text{Elect.opp}_{s,t}}$	-0.879	(-1.253, -0.802)	$oldsymbol{\omega}_6$	$8.71923 imes 10^{-5}$	(-0.0193, 0.0314)

Table 2: Posterior estimates and 95% credible intervals for parameters in θ_{μ^*} and θ_{ϕ} , appearing in the conditional intensity function log likelihood (see Equation (12)), as determined using a hybrid MCMC approach.

4.3 Extreme Afghan terror prediction 2019-2021

To be able to predict future events in 2019, 2020 and 2021 for which there is no data in the GTD, the time variable T was extended to December 31st 2021. Figure 9 shows predicted yearly spatio-temporal average intensities of attacks (log-scaled) leading to 20 or more casualties across the region between these years and confirms the fitted intensities for





Log intensity of attacks 2014

Log intensity of attacks 2015





Figure 7: Afghanistan spatial-temporal intensity (log-scale) fits for attacks producing at least 20 casualties, between 2013-2018 with events shown as black diamonds.

the previous years in a forecasted surge of attacks, leading to the recent insurgency success. This is especially heightened around the "ring-road" previously mentioned, with the extremity of insurgent terror around the district of Kabul shown to be increasing, more than doubling the yearly average extreme casualty (≥ 20) intensity of 0.001 in 2013 to a predicted of 0.112 in 2021, over pixelated grids of 0.06×0.1 longitude/latitude units. This increase in attacks also focuses on regions that have not previously seen much terror activity. These areas include the North-eastern Badakhshan Province, a mountainous region which recently has seen increased presence of the Uyghur armed insurgency group, the East Turkestan Islamic Movement (Devonshire-Ellis, 2021). In particular, in this region lies the Wakhan Corridor, connecting Afghanistan to neighboring countries and has been subject to recent road plannings by the Chinese government, as part of the Belt and Road construction (Goulard, 2021), for greater market access. Figure 8 shows a disrict map of Afghanistan. Amiri (2017) shows a planned map of new roads to have been developed since 2017, of which districts in the north and center of Afghanistan had been planned for greater infrastructure, aligning closely with an increased risk of terror insurgence predicted by our model. Our analysis therefore emphasizes the increase in extreme terror related activities from new insurgent groups and infrastructure developments.



Figure 8: Afghanistan districts by name. Planned new road developments in Afghanistan in 2017 (Amiri, 2017) highlight districts in which increased extreme terror activity has been predicted by our model (see Figure 9).



Figure 9: Afghanistan spatial-temporal predicted intensity (log-scale) for attacks producing at least 20 casualties, between 2019-2021.

5 Conclusion

In this paper, we have formulated a flexible spatio-temporal model to study extreme terror insurgencies and study Afghanistan between 2013 and 2021 to show the effectiveness of our model. Our framework importantly focuses on quantifying the risk of *extreme* attacks using a mixed distribution comprising a discrete extreme-valued and a zero inflated model of casualties produced per attack. By incorporating this measure on casualties defined as marks within a spatio-temporal point process framework, the conditional modeling of spatio-temporal trends of extreme insurgencies has been developed. Here, a self-exciting model for terror attacks and resurgences has been constructed through an inhomogeneous Poisson baseline and a Gaussian Process formulated triggering kernel. The former accounts for important covariates related to standard terror behavior, which in the study of Afghanistan includes population density at the country's district level and voting density for the opposing government. The latter develops a flexible prior over

all Gaussian Process covariance structures to determine triggering functions that can succinctly incorporate nuanced terror patterns observed in both space and time. All unknown parameters of the model are quantified via a hybrid particle MCMC algorithm which utilizes Mercer's approximation to Gaussian Process covariances. Spatio-temporal prediction of extreme terror insurgencies in Afghanistan using our model, corroborate strongly with recent insurgent events in the country and infrastructure developments that had planned to take place in areas previously unperceived to terror.

The methodology we have developed provides a basis for several technical extensions. For example, the handling of missing data as studied in Tucker et al. (2019) for the GTD could be fully accounted for in the constructed MCMC algorithm, however, is likely to suffer from severe computational drawbacks when analyzing much larger datasets than those studied in this paper. A potential scope for future research would therefore be to investigate avenues providing faster inferential schemes or approximate inference that are compatible with the Gaussian Process methodology we have presented and the missing data problem inherent to the data. Further, differing covariance structures could be studied by our broad methodology to analyze the effect of the unknown triggering function on the overall organization of planned events.

While this paper focuses on developing a self-exciting point process framework for characterizing spatio-temporal terror patterns of insurgency and their societal risk, the flexibility inherent to our model is likely to be applicable to a broad range of events that are observed in space-time and produce a measurable outcome of interest. For example, stochastic processes underlying stock market return data, extreme weather conditions and cyber-network traffic events could all benefit from our model formulation and inference mechanism to model and predict the risk of extreme events. The novel inclusion of the nonparametric Gaussian Process prior placed upon the triggering kernel particularly highlights the applicability of our work to understudied self-exciting processes that are beyond the scope of this paper.

Acknowledgement

This work was supported by a Sandia National Laboratories Laboratory Directed Research and Development (LDRD) grant. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525 and SAND2021-12962 O. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

References

Amiri, A. (2017). Kabul's plan to realize Afghanistan's geographic dividend. The Diplomat.

- Baddeley, A. and R. Turner (2000). Practical maximum pseudolikelihood for spatial point patterns. *Australian & New Zealand Journal of Statistics* 42(3), 283–322.
- Betancourt, M. (2018). A conceptual introduction to Hamiltonian Monte carlo. arXiv preprint arXiv:1701.02434.
- Bortot, P. and C. Gaetan (2014). A latent process model for temporal extremes. *Scandinavian Journal of Statistics* 41(3), 606–621.
- Bortot, P. and J. A. Tawn (1998). Models for the extremes of Markov chains. *Biometrika* 85(4), 851–867.
- Bracken, C., B. Rajagopalan, L. Cheng, W. Kleiber, and S. Gangopadhyay (2016). Spatial Bayesian hierarchical modeling of precipitation extremes over a large domain. *Water Resources Research* 52(8), 6643–6655.
- Castro-Camilo, D., R. Huser, and H. Rue (2019). A spliced Gamma-generalized Pareto model for short-term extreme wind speed probabilistic forecasting. *Journal of Agricultural, Biological and Environmental Statistics* 24(3), 517–534.
- Cavanaugh, J. E. and A. A. Neath (2019). The Akaike information criterion: Background, derivation, properties, application, interpretation, and refinements. *WIREs Computational Statistics 11*(3), e1460.
- Christophe, A., A. Doucet, and R. Holenstein (2010). Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society B* 72(3), 269–342.
- Clark, N. J. and P. M. Dixon (2018). Modeling and estimation for self-exciting spatio-temporal models of terrorist activity. *The Annals of Applied Statistics* 12(1), 633–653.
- Clauset, A. and R. Woodard (2013). Estimating the historical and future probabilities of large terrorist events. *The Annals of Applied Statistics* 7(4), 1838–1865.
- Coles, S. (2001). An introduction to statistical modeling of extreme values. Springer Series in Statistics. Springer-Verlag.

- Cooley, D., D. Nychka, and P. Naveau (2007). Bayesian spatial modeling of extreme precipitation return levels. *Journal* of the American Statistical Association 102(479), 824–840.
- Cooley, D. and S. R. Sain (2010). Spatial hierarchical modeling of precipitation extremes from a regional climate model. *Journal of Agricultural, Biological, and Environmental Statistics* 15(3), 381–402.
- Da Fonseca, J. and R. Zaatour (2014). Hawkes process: Fast calibration, application to trade clustering, and diffusive limit. *Journal of Futures Markets* 34(6), 548–579.
- Devonshire-Ellis, C. (2021). The Afghanistan-China belt & road initiative. Silk Road Briefing.
- Flaxman, S., Y. W. Teh, and D. Sejdinovic (2017). Poisson intensity estimation with reproducing kernels. *Electron. J. Statist.* 11(2), 5081–5104.
- Gabaix, X. (1999). Zipf's law for cities: An explanation. The Quarterly Journal of Economics 114(3), 739-767.
- Gaetan, C. and M. Grigoletto (2007). A hierarchical model for the analysis of spatial rainfall extremes. *Journal of Agricultural, Biological, and Environmental Statistics* 12(4), 434.
- Goulard, S. (2021). Does the belt and road have a future in Taliban-ruled Afghanistan? The Diplomat.
- Grothe, O., V. Korniichuk, and H. Manner (2014). Modeling multivariate extreme events using self-exciting point processes. *Journal of Econometrics* 182(2), 269–289.
- GTD (2017). Global Terrorism Database: National consortium for the study of terrorism and responses to terrorism. https://www.start.umd.edu/gtd/. Accessed March 2020.
- Haario, H., E. Saksman, and J. Tamminen (2001). An adaptive Metropolis algorithm. Bernoulli 7(2), 223-242.
- Hawkes, A. G. (1971, 04). Spectra of some self-exciting and mutually exciting point processes. *Biometrika* 58(1), 83–90.
- Hawkes, A. G. and D. Oakes (1974). A cluster process representation of a self-exciting process. *Journal of Applied Probability 11*(3), 493–503.
- Hitz, A., R. Davis, and G. Samorodnitsky (2017). Discrete extremes. arXiv preprint arXiv:1707.05033.
- Huerta, G. and B. Sanso (2007). Time-varying models for extreme values. *Environmental and Ecological Statistics* 14(3), 285–299.
- Li, W. (1992). Random texts exhibit Zipf's-law-like word frequency distribution. *IEEE Transactions on Information Theory* 38(6), 1842–1845.
- Lyra, M. L., Costa, U. M. S., Costa Filho, R. N., and J. S. Andrade jr. (2003). Generalized Zipf's law in proportional voting processes. *Europhys. Lett.* 62(1), 131–137.
- MacDonald, A., C. J. Scarrott, D. Lee, B. Darlow, M. Reale, and G. Russell (2011). A flexible extreme value mixture model. *Computational Statistics & Data Analysis* 55(6), 2137–2157.
- Mandelbrot, B. (1953). Contribution à la théorie mathématique des jeux de communication. Publications de l'Institut de statistique de l'Université de Paris. Paris: Laboratoires d'électronique et de physique appliquées.
- Mohler, G. (2013). Modeling and estimation of multi-source clustering in crime and security data. *The Annals of Applied Statistics* 7(3), 1525–1539.
- Mohler, G. O., M. B. Short, P. J. Brantingham, F. P. Schoenberg, and G. E. Tita (2011). Self-exciting point process modeling of crime. *Journal of the American Statistical Association 106*(493), 100–108.
- Mohtadi, H. and A. P. Murshid (2009). Risk of catastrophic terrorism: an extreme value approach. *Journal of Applied Econometrics* 24(4), 537–559.
- Møller, J. and R. Waagepetersen (2004). *Statistical Inference and Simulation for Spatial Point Processes*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability.
- Neal, R. M. (2011). MCMC using Hamiltonian dynamics. Chapman and Hall/CRC.
- Ogata, Y. (1988). Statistical models for earthquake occurrences and residual analysis for point processes. *Journal of the American Statistical Association 83*(401), 9–27.
- Patel, L., L. Shand, J. D. Tucker, and G. Huerta (2020). Assessing extreme value analysis to predict rare events from the Global Terrorism Database. *ASA Joint Statistical Meetings Proceedings*, 1–26.
- Rastogi, T. (2016). A power law approach to estimating fake social network accounts. arXiv preprint arXiv:1605.07984.
- Reinhart, A. (2017). A review of self-exciting spatio-temporal point processes and their applications. *Statistical Science* 33, 299–318.

- Rizoiu, M.-A., Y. Lee, S. Mishra, and L. Xie (2017). A tutorial on Hawkes processes for events in social media. *Frontiers of Multimedia Research*, 191–218.
- Schafer, J. (1997). Analysis of multivariate incomplete data. London, UK: Chapman & Hall.
- Sharkey, P. and H. C. Winter (2019). A Bayesian spatial hierarchical model for extreme precipitation in Great Britain. *Environmetrics 30*(1), 25–29.
- Spiegelhalter, D. J., N. G. Best, B. P. Carlin, and A. Van Der Linde (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 64(4), 583–639.
- Tucker, J. D., L. Shand, and J. R. Lewis (2019). Handling missing data in self-exciting point process models. *Spatial Statistics 29*, 160–176.
- Tuerk, A. (2019). Afghan polls close as voters face Taliban threats. NPR.
- Walder, C. J. and A. N. Bishop (2018). Fast Bayesian intensity estimation for the permanental process. *arXiv preprint* arXiv:1701.03535.
- Zammit-Mangion, A., M. Dewar, V. Kadirkamanathan, and G. Sanguinetti (2012). Point process modelling of the Afghan war diary. *Proceedings of the National Academy of Sciences 109*(31), 12414–12419.
- Zhang, R., C. Walder, M. A. Rizoiu, and L. Xie (2018). Efficient Non-parametric Bayesian Hawkes Processes. arXiv preprint arXiv:1810.03730.
- Zipf, G. K. (1949). Human behaviour and the principle of least effort. Reading, MA: Addison-Wesley.