International Journal of Forecasting 34 (2018) 431–439

Contents lists available at ScienceDirect

International Journal of Forecasting

journal homepage: www.elsevier.com/locate/ijforecast

Improving social harm indices with a modulated Hawkes process

George Mohler [a,*], Jeremy Carter [b], Rajeev Raje [a]

[a] Computer and Information Science, Indiana University - Purdue University Indianapolis, United States
[b] School of Public and Environmental Affairs, Indiana University - Purdue University Indianapolis, United States

A B S T R A C T

Communities are affected adversely by a range of social harm events, such as crime, traffic crashes, medical emergencies, and drug use. The police, fire, health and social service departments are tasked with mitigating such social harm through various types of interventions. While various different social harm indices have been proposed for allocating resources to spatially fixed hotspots, the risk of social harm events is dynamic, and new algorithms and software systems that are capable of quickly identifying risks and triggering appropriate public safety responses are needed. We propose a novel modulated Hawkes process for this purpose that offers flexible approaches to both (i) the incorporation of spatial covariates and leading indicators for variance reduction in the case of rarer event categories, and (ii) the capture of dynamic hotspot formation through self-excitation. We present an efficient l1-penalized EM algorithm for estimating the model that performs feature selection for the spatial covariates of each incident type simultaneously. We provide simulation results using data from the Indianapolis Metropolitan Police Department in order to illustrate the advantages of the modulated Hawkes process model of social harm over various recently introduced social harm indices and property crime Hawkes processes.

1. Introduction

Crime is highly concentrated in urban communities, and hotspot or "predictive" policing efforts aim to apply limited resources to high intensity geographic areas and time intervals in order to disrupt crime opportunities, leading to aggregate crime rate reductions (Braga & Bond, 2008; Mohler et al., 2015; Ratcliffe, Taniguchi, Groff, & Wood, 2011; Weisburd et al., 2006). A number of different algorithmic methods have been proposed for estimating crime hotspot risk, including multivariate models (Kennedy, Caplan, & Piza, 2011; Liu & Brown, 2003; Wang, Brown, & Gerber, 2012), kernel density estimation (Bowers, Johnson, & Pease, 2004; Chainey, Tompson, & Uhlig, 2008; Fielding & Jones, 2012; Johnson, 2007; Johnson, Bowers, Birks, & Pease, 2009) and spatio-temporal point processes (Mohler, 2014; Mohler, Short, Brantingham, Schoenberg, & Tita, 2011). Point processes and density estimation have the advantage of capturing near-repeat effects and only require event data as inputs, whereas multivariate models achieve variance reduction through the introduction of spatial covariates (though the variance can be increased if irrelevant covariates are included). Mohler et al. (2015) conducted field trials of predictive policing using a property crime Hawkes process, where patrols directed through the Hawkes process led to statistically significant crime rate reductions compared to analyst-directed patrols.

However, the police also fill other roles in the community beyond crime response and prevention, including traffic enforcement, emergency medical services (EMS) response, and more generally dealing with events related to social harm (Ratcliffe, 2015). Despite these multiple and disparate daily challenges, the existing hotspot and predictive policing algorithms and intervention strategies focus on individual or groups of related sub-categories of social harm events. Scholars have recently called for the next evolution of hotspot policing to move beyond crime counts in space and time to the more expansive and hierarchical approach of policing "social harm" (Ratcliffe, 2015; Sherman, Neyroud, & Neyroud, 2016; Weinborn, Ariel, Sherman, & O'Dwyer, 2017). A recent approach to quantifying the impact of crime on society has been the development of crime harm indices (Curtis-Ham & Walton, 2017; Sherman

* Corresponding author.
E-mail address: gmohler@iupui.edu (G. Mohler).

https://doi.org/10.1016/j.ijforecast.2018.01.006
0169-2070/

et al., 2016) that attempt to weight crime offenses based on sentencing guidelines, as opposed to the simpler count measures of crime occurrences. In this context, harm is operationalized as the impact upon society, depending on the qualitatively different levels of severity across incidents of crime. This approach to quantifying crime has implications for the development of more effective police interventions for reducing harm, as opposed to reducing crime counts. With this end-goal in mind, Ratcliffe (2015) extends the idea of a crime harm index to a social harm index that includes incidents to which police must respond that fall outside the traditional definition of crime but still inflict harm on society, such as vehicle crashes. The current study takes the notion of social harm a step further to include additional incidents within the police purview that affect society. Here, harm is operationalized beyond the impact of crime on society to be more inclusive of the nature of police work. Put simply, a focus on social harm builds on hotspots policing by applying similar methodological approaches, but broadens the list of harm incidents to reflect day-to-day policing demands (e.g., crime, medical emergencies, vehicle crashes, etc.) more accurately, while weighting these various incidents to reflect the degree of severity of the harm each may inflict upon society.

The preliminary findings in social harm research suggest that the inclusion, and weighting, of various harm incidents holds substantive promise for police practice and intervention. To date, the most common approach to weighting social harms has been to map sentencing data to specific crime offenses. This method has taken the form of both actual sentencing outcomes (Babyak, Alavi, Collins, Halladay, & Tapper, 2009; Bangs, 2016; Curtis-Ham & Walton, 2017) and prescriptive sentencing guidelines, often referred to as gravity or severity scales (Ratcliffe, 2015; Sherman et al., 2016; Weinborn et al., 2017), and makes use of suggested sentencing lengths to rank the "harm" of a given offense. For example, criminal homicide may have a sentencing guideline of 24 years in prison, while armed robbery may elicit a 12-year sentence, and residential burglary a 6-year sentence. In such a weighting scenario, criminal homicide would be twice as severe or harmful as armed robbery and four times as harmful as residential burglary. Weighting by sentencing guidelines can take many forms, and the discussion presented here is limited to the importance of weighting crimes and other incidents by their severity. Indeed, "neither criminology nor the adjacent social sciences have made a serious effort to systematically identify, evaluate or compare the harms associated with different crimes" (Greenfield & Paoli, 2013) and "focusing merely on counts, rather than on the severity or harm of crime, is somewhat crude and imprecise" (Weinborn et al., 2017). Sherman et al. (2016) provide a robust discussion of varying weighting procedures using sentencing guidelines.

Studies employing this approach have concluded that social harm varies across police patrol districts (Ratcliffe, 2015) and that a small proportion of crime victims are exposed to greater levels of social harm (Dudfield, Angel, Sherman, & Torrence, 2017). Most closely related to the current study, Weinborn et al. (2017) employed the Cambridge Harm Index (CHI) of Sherman et al. (2016) wherein crimes are weighted by the number of days in prison for a given offense, as outlined in the Home Office Sentencing Guidelines, to examine the spatio-temporal concentration of crime counts versus CHI social harm. Their results indicated social harm to be three times as concentrated as crime counts alone across 15 councils in England and Wales over a 12-month period. Interestingly, and critically in pointing out the need for scholars to consider a variety of social harms beyond traditional hotspot policing strategies, the authors observed that only 25% of their crime count hotspots overlapped with their social harm locations, or "harmspots". Thus, while it can be insightful to conduct spatiotemporal analyses of crime counts alone in order to focus police strategies, it seems prudent to account for the severity of the harm that crime may cause, as crimes are not all created equal, and the spatiotemporal variation of more harmful incidents may differ from that of less harmful events. Moreover, given that harmspots exhibit different spatiotemporal patterns from hotspots, they may also have different corollary relationships with community structure from hotspots; hence one focus of the present study.

The present study contributes further to the social harm policing literature through the inclusion of multiple harm types that are yet to be examined in a single study. The present study includes a range of Part 1 (the most serious crimes that occur regularly across all jurisdictions and are likely to be reported to the police) and Part 2 (other crimes) criminal offenses, as well as vehicle crashes and drug overdoses — the latter of which is currently regarded as one of the most concerning social harms, as drug overdose deaths across the United States have more than quadrupled since 1999 (Rudd, Aleshire, Zibbell, & Gladden, 2016). Part 1 and Part 2 criminal offenses are defined by the Federal Bureau of Investigation (2016) as a tiered classification system for the Uniform Crime Reporting and National Incident Based Reporting System. Furthermore, unlike static social harm indices that are estimated over a fixed historical window of observation, our methodology produces a dynamic harm index that incorporates new event data daily to account for spatiotemporal fluctuations in the social harm risk.

This study introduces a modulated Hawkes process for modeling dynamic social harm hotspots. The model combines several advantageous aspects of existing multivariate regression and point process models. In particular, the model consists of a background modulated Poisson process that links spatial covariates (census variables, average crime rate, etc.) to the risk of each social harm event category. Our estimation procedure also includes automatic variable selection in order to prevent over-fitting and determine important covariates for model explanation. Secondly, the point process approach allows for the incorporation of the self-excitation that is present in some event categories. Because the output of the modulated Hawkes process is a conditional intensity for each event type, a dynamic social harm index can be defined easily by calculating the expected cost of a given spatial region and time interval.

The outline of this paper is as follows. Section 2 gives an overview of the data set used in our study and the methods that are used to estimate the average societal cost of each event type. Section 3 provides the mathematical details of the modulated Hawkes process and a l1-penalized