

Space-Time Modelling of Terrorism and Counter-Terrorism

Stephen Ashley Tench

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
of the
University College London

UCL Department of Security and Crime Science
UCL Centre for Advanced Spatial Analysis
University College London

May 2018

Declaration

I, Stephen Tench, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

In this thesis multiple approaches are presented which demonstrate the effectiveness of mathematical modelling to the study of terrorism and counter-terrorism strategies. In particular, theories of crime science are quantified to obtain objective outcomes. The layout of the research findings is in four parts.

The first model studied is a Hawkes point process. This model describes events where past occurrence can lead to an increase in future events. In the context of this thesis a point process is used to capture dependence among terrorist attacks committed by the Provisional Irish Republican Army (PIRA) during “The Troubles” in Northern Ireland. The Hawkes process is adapted to produce a method capable of determining quantitatively temporally distinct phases within the PIRA movement. Expanding on the Hawkes model the next area of research introduces a time-varying background rate. In particular, using the Fast Fourier Transform a sinusoidal background rate is derived. This model then enables a study of seasonal trends in the attack profile of the Al Shabaab (AS) group. To study the spatial dynamics of terrorist activity a Dirichlet Process Mixture (DPM) model is examined. The DPM is used in a novel setting by considering the influence of improvised explosive device (IED) factory closures on PIRA attacks. The final research area studied in this thesis is data collection methods. An information retrieval (IR) tool is designed which can automatically obtain terrorist event details. Machine learning techniques are used to compare this IR data to a manually collected dataset. Future research ideas are introduced for each of the topics covered in this dissertation.

Impact Statement

The aim of the research conducted in this dissertation is to develop methods which can quantitatively describe phenomena observed in terrorist activity. To achieve this aim the fields of mathematical modelling and crime science have been utilised to generate numerical insights into a range of terrorism and counter-terrorism datasets. In particular, spatial and temporal patterns in these databases have been derived with real-world interpretations which can be of benefit to both academics and practitioners.

From an academic perspective all the research conducted in this thesis acts as a good foundation for future studies. For example, researchers studying terrorist groups will likely find the mathematical models which have been developed here useful in a wide variety of contexts. All of the modelling techniques developed throughout this dissertation are likely adaptable to other terrorist groups and settings. In addition, the insights found by utilising the frameworks derived in this thesis could prove highly beneficial to complementing qualitative and quantitative studies. An example of complementing qualitative research in this dissertation is the use of a Hawkes process model to quantify temporal shifts in terrorist group organisations. These shifts were found to demonstrate gaps in current social science findings. By forming a bridge between qualitative and quantitative research this thesis proves useful in a wide variety of disciplines.

The research in this dissertation has also proven its academic impact via journal publication. In 2016 research from this thesis was published in the European Journal

of Applied Mathematics with the title “Spatio-temporal patterns of IED usage by the Provisional Irish Republican Army”. Moreover, this research has been distributed to a wider audience through other media. In particular, The Washington Post and the Society for Industrial and Applied Mathematics have featured articles covering some of the research findings in this thesis. Further papers are also currently under preparation for future publication.

As well as having a wider academic impact there is potential for the models studied in this thesis to be used in a variety of ways by practitioners. Tackling terrorism and constructing counter-terrorism strategies are highly important security issues facing societies across the globe. Therefore, methods which can assist in understanding the dynamics underpinning terrorist activity are essential. Mathematical modelling can also provide an objective framework to analyse terrorist activities leading to more precise counter-terrorism measures. The research conducted as part of this thesis provides practitioners many different ways to collect and analyse terrorism data. Moreover, it has been shown in this dissertation that the models presented have real-world interpretations. This linking between numerical findings and real-world phenomena would likely prove useful to practitioners unfamiliar with the mathematical concepts being presented.

Overall, the research covered in this thesis will likely prove useful in many different ways. Within academia the research covered in this dissertation can provide useful techniques to assist in future terrorism studies. On the other hand, the models, and their practical interpretations, will likely be of benefit to practitioners aiding both their decision making and design of counter-terrorism approaches.

Acknowledgments

I would like to express my sincere gratitude to my two supervisors Dr Hannah Fry and Dr Paul Gill. Their insights and guidance have helped immensely throughout the completion of this thesis. I also wish to acknowledge both the UCL Department of Security and Crime Science and the UCL Centre for Advanced Spatial Analysis (CASA) for hosting me during this research project. Finally, I would like to extend my personal thanks to my parents for their steadfast support during my academic career.

The research presented here was made possible through the grant provided via the EPSRC Security Science Doctoral Training Centre with reference number EP/G037264/1

CONTENTS

Chapter 1	Introduction and Literature Review	22
1.1	Introduction	23
1.2	Motivation	23
1.3	Crime Science and Terrorism	27
1.3.1	Rational Choice Perspective	28
1.3.2	Routine Activities Theory	28
1.3.3	Crime Pattern Theory	29
1.3.4	Crime Reduction	30
1.3.5	Patterns in Terrorism	31
1.4	Research Questions	33
1.5	Case Studies	34
1.5.1	Provisional Irish Republican Army	34
1.5.2	Al Shabaab	37
1.6	Research Methodology	38
1.6.1	Hawkes Process	39
1.6.2	Geographic Profiling	45
1.7	Data Issues	47

1.8	Dissertation Structure	50
Chapter 2 Hawkes Process Modelling of the PIRA		52
2.1	Introduction	53
2.2	Mathematical Formulation of Point Process Models	53
2.2.1	Poisson Process Model	55
2.2.2	Hawkes Process Model	55
2.3	Parameter Estimation	58
2.4	Model Goodness of Fit	62
2.5	Applying the Hawkes Process	65
2.5.1	IED Data	65
2.5.2	Change Point Detection	66
2.5.3	Change Point Analysis Based on Sociological Boundaries	67
2.5.4	Change Point Analysis Independent of Sociological Boundaries	73
2.6	Interpreting Phase Boundaries	83
2.7	Interpreting Model Parameters	87
2.7.1	KS Method Parameters	87
2.7.2	AIC Method Parameters	91
2.8	Discussion	92
Chapter 3 Seasonal Hawkes Process Modelling of Al Shabaab		96
3.1	Datasets	98
3.1.1	Al Shabaab Event Data	98
3.1.2	Weather Data	103

3.2	Modelling Framework	104
3.3	Identifying Seasonality	106
3.3.1	Detrending Data	107
3.3.2	Fourier Series	109
3.3.3	Discrete Fourier Transform	112
3.3.4	Fast Fourier Transform	115
3.3.5	Fourier Analysis Application	119
3.3.6	Results of Fourier Series Analysis	123
3.4	Hawkes Process Modelling	133
3.4.1	Hawkes Process with Time Independent Background Rate	133
3.5	Hawkes Process with Seasonal Background Rate	136
3.6	Results	139
3.7	Discussion	141

Chapter 4 Geographic Profiling using the Dirichlet Process Mixture Model **148**

4.1	Datasets	153
4.2	Geographic Profiling	155
4.2.1	Criminal Geographic Profiling	155
4.2.2	Bayesian Framework of CGT	157
4.3	Dirichlet Process Mixture Model	161
4.3.1	Model Theory	161
4.3.2	DPMM Formulae	164
4.3.3	Markov Chain Monte Carlo Method of Inference	167

4.3.4	Gibbs Sampling for DPMM	168
4.3.5	DPMM Parameter Fitting	171
4.4	DPMM Analysis of PIRA IED Factories	181
4.5	Analysis	182
4.5.1	Model Calibration	182
4.5.2	Results	183
4.6	Discussion	191
Chapter 5 Manual vs Automatic Data Collection Methods		199
5.1	Introduction	200
5.2	Datasets	205
5.2.1	Manually Collected Dataset	205
5.2.2	Automatically Collected Dataset	206
5.3	Comparison Techniques	208
5.3.1	Dynamic Time Warping	208
5.3.2	Term Frequency - Inverse Document Frequency	211
5.4	Dataset Comparisons	214
5.4.1	Event Date Comparisons	214
5.4.2	Location Comparisons	214
5.4.3	Event Note Comparisons	216
5.4.4	Comparisons Framework	216
5.5	Results	218
5.5.1	ACLED Non-Local Source Comparison	219
5.5.2	ACLED Only Local Source Comparison	220

5.6 Discussion	222
Chapter 6 Conclusion and Future Research	227
Appendix A Spatio-temporal patterns of IED usage by the Provi- sional Irish Republican Army (Published Paper)	234
Bibliography	264

LIST OF FIGURES

2.1 Figure (a) shows the values of $\Delta = D_\alpha - D_n$ for each point added from phase 1 to phase 2. Figures (b)-(d) show the corresponding changes in the MLE parameter values. The vertical line indicates the change point between the phases as determined by maximising $\Delta = D_\alpha - D_n$. The x-axis runs backwards in time representing each new pointed added from phase 1 to phase 2. 69

2.2 Figure (a) shows the values of $\Delta = D_\alpha - D_n$ for each point added from phase 2 to phase 3. Figures (b)-(d) show the corresponding changes in the MLE parameter values. The vertical line indicates the change point between the phases as determined by maximising $\Delta = D_\alpha - D_n$. The x-axis runs backwards in time representing each new pointed added from phase 2 to phase 3. 70

-
- 2.3 Figure (a) shows the values of $\Delta = D_\alpha - D_n$ for each point added from phase 3 to phase 4. Figures (b)-(d) show the corresponding changes in the MLE parameter values. The vertical line indicates the change point between the phases as determined by maximising $\Delta = D_\alpha - D_n$. The x-axis runs backwards in time representing each new pointed added from phase 3 to phase 4. 71
- 2.4 Figure (a) shows the values of $\Delta = D_\alpha - D_n$ for each point added from phase 4 to phase 5. Figures (b)-(d) show the corresponding changes in the MLE parameter values. The vertical line indicates the change point between the phases as determined by maximising $\Delta = D_\alpha - D_n$. The x-axis runs backwards in time representing each new pointed added from phase 4 to phase 5. 72
- 2.5 Graphs (a)-(c) show the distribution of the parameter values for the Hawkes process in phase 1 (blue) and phase 2 (red) after each attempt at minimising the total AIC. Graph (d) is the corresponding distribution of the AIC values for each of these Hawkes process models. 77
- 2.6 Graphs (a)-(c) show the distribution of the parameter values for the Hawkes process in phase 1 (blue) and phase 2 (red) after each attempt at minimising the total AIC and restricting the results to the case $\omega \leq 1$. Graph (d) is the corresponding distribution of the AIC values for each of these Hawkes process models. 78

2.7	Graphs (a)-(c) show the distribution of the simulated parameter values for the Hawkes process in phase 1 (blue) and phase 2 (red). The parameters have been simulated from the Hawkes process models in phases 1 and 2 as defined by the AIC method for determining boundaries. The vertical lines represent the original model parameters. . . .	82
2.8	Distribution of IED events between phases under the original phase boundaries (blue), KS method boundaries (orange) and AIC boundary (red).	84
3.1	Map of Somalia. The area indicated in red indicates the location of the weather station from which the rainfall data used in this chapter was obtained.	105
3.2	Battle territory phase 3 binned data histogram. Bin widths are set to 30 days to correspond to approximately one month of data.	108
3.3	Battle territory phase 3 binned data histogram with a linear detrend. Bin widths are set to 30 days to correspond to approximately one month of data.	110

-
- 3.4 (a) Battle territory phase 3 weather data histogram. (b) Battle territory phase 3 weather data mean detrended histogram. (c) Battle territory phase 3 monthly weather data FFT frequencies power spectrum. The higher the power the greater the significance of the frequency. The highlighted points (black circles) correspond to the two frequencies with highest power which are used to construct a Fourier series approximation for the data. (d) Battle territory phase 3 weather data mean detrended histogram with Fourier series constructed from two frequencies with highest power. 126
- 3.5 (a) Battle territory phase 4 weather data histogram. (b) Battle territory phase 4 weather data mean detrended histogram. (c) Battle territory phase 4 monthly weather data FFT frequencies power spectrum. The higher the power the greater the significance of the frequency. The highlighted points (black circles) correspond to the two frequencies with highest power which are used to construct a Fourier series approximation for the data. (d) Battle territory phase 4 weather data mean detrended histogram with Fourier series constructed from two frequencies with highest power. 127

-
- 3.6 (a) Battle territory phase 3 event data histogram. (b) Battle territory phase 3 event data linearly detrended histogram. (c) Battle territory phase 3 monthly event data FFT frequencies power spectrum. The higher the power the greater the significance of the frequency. The highlighted points (black circles) correspond to the two frequencies with highest power which are used to construct a Fourier series approximation for the data. (d) Battle territory phase 3 event data linearly detrended histogram with Fourier series constructed from two frequencies with highest power. 128
- 3.7 (a) Battle territory phase 4 event data histogram. (b) Battle territory phase 4 event data mean detrended histogram. (c) Battle territory phase 4 monthly event data FFT frequencies power spectrum. The higher the power the greater the significance of the frequency. The highlighted points (black circles) correspond to the two frequencies with highest power which are used to construct a Fourier series approximation for the data. (d) Battle territory phase 4 event data mean detrended histogram with Fourier series constructed from two frequencies with highest power. 129

-
- 3.8 (a) Violence against civilians phase 3 event data histogram. (b) Violence against civilians phase 3 event data mean detrended histogram. (c) Violence against civilians phase 3 monthly event data FFT frequencies power spectrum. The higher the power the greater the significance of the frequency. The highlighted points (black circles) correspond to the two frequencies with highest power which are used to construct a Fourier series approximation for the data. (d) Violence against civilians phase 3 event data mean detrended histogram with Fourier series constructed from two frequencies with highest power. 130
- 3.9 (a) Violence against civilians phase 4 event data histogram. (b) Violence against civilians phase 4 event data mean detrended histogram. (c) Violence against civilians phase 4 monthly event data FFT frequencies power spectrum. The higher the power the greater the significance of the frequency. The highlighted points (black circles) correspond to the two frequencies with highest power which are used to construct a Fourier series approximation for the data. (d) Violence against civilians phase 4 event data mean detrended histogram with Fourier series constructed from two frequencies with highest power. 131

-
- 4.1 (a)-(o) Jeopardy surface distribution of PIRA IED attacks (red circles) before the closure of IED factories. Old factories closures are marked via blue squares whilst the new closure in each figure is marked with a black triangle. On the jeopardy surfaces white areas indicate likely source locations whilst red areas have a lower probability of containing a source. 190
- 4.2 Jeopardy surface covering all PIRA IED attacks which occurred after the final recorded factory closure. Red circles highlight IED attacks and blue squares mark the locations of all previously identified factory closures. On the jeopardy surface white areas indicate likely source locations whilst red areas have a lower probability of containing a source. 191
- 5.1 (a) Barcode plot comparing the timeseries of IR and ACLED events in days since each datasets first event. (b) Time warping path between the IR and ACLED timeseries. This path is derived from the dynamic time warping (DTW) method. 219
- 5.2 Location and event description comparison of the ACLED and IR datasets. For each event in the IR dataset the closest datapoint in the ACLED database in time and location is shown in red. Similar results are shown in blue for the closest event in time and event textual description. 219

5.3	(a) Percentage of event locations matched with a tf-idf value greater than or equal to threshold values $\{0.1, 0.2, \dots, 1.0\}$. (b) Percentage of event descriptions matched with a tf-idf value greater than or equal to threshold values $\{0.1, 0.2, \dots, 1.0\}$	220
5.4	(a) Barcode plot comparing the timeseries of IR and ACLED events in days since each datasets first event. (b) Time warping path between the IR and ACLED timeseries. This path is derived from the dynamic time warping (DTW) method.	220
5.5	Location and event description comparison of the ACLED and IR datasets. For each event in the IR dataset the closest datapoint in the ACLED database in time and location is shown in red. Similar results are shown in blue for the closest event in time and event textual description.	221
5.6	(a) Percentage of event locations matched with a tf-idf value greater than or equal to threshold values $\{0.1, 0.2, \dots, 1.0\}$. (b) Percentage of event descriptions matched with a tf-idf value greater than or equal to threshold values $\{0.1, 0.2, \dots, 1.0\}$	221

LIST OF TABLES

2.1	PIRA IED Dataset Event Fields	65
2.2	Change Point Analysis Results with KS Test Method Boundaries . . .	73
2.3	Parameters with KS Test Method Boundaries	73
2.4	Change Point Analysis Results with AIC Test Method Boundaries . . .	79
2.5	Parameters with AIC Test Method Boundaries	79
2.6	Phase 1 Bootstrap Statistics	83
2.7	Phase 2 Bootstrap Statistics	83
3.1	ACLED Data Fields	99
3.2	Number of Entries in BT Datasets	102
3.3	Number of Entries in VAC Datasets	103
3.4	Number of Entries in BT Weather Datasets	104
3.5	Number of Entries in VAC Weather Datasets	104
3.6	Battle Territory Phase 3 Linear Detrend Function Coefficients	109
3.7	Battle Territory Phase 3 Periods	125
3.8	Battle Territory Phase 4 Periods	125
3.9	Violence Against Civilians Phase 3 Periods	125
3.10	Violence Against Civilians Phase 4 Periods	125

3.11	Battle Territory Poisson Process Results	135
3.12	Violence Against Civilians Poisson Process Results	135
3.13	Battle Territory Time Invariant Background Hawkes Process Results	135
3.14	Violence Against Civilians Time Invariant Background Hawkes Pro- cess Results	136
3.15	Battle Territory Phase 3 Hawkes Process Results	140
3.16	Battle Territory Phase 4 Hawkes Process Results	140
3.17	Violence Against Civilians Phase 3 Hawkes Process Results	140
3.18	Violence Against Civilians Phase 4 Hawkes Process Results	140
4.1	PIRA IED Dataset	155
4.2	Factory Closure Dataset	155
4.3	Datasets Analysed using the DPMM	186
5.1	ACLED Data Headings	206
5.2	Automated Article Retrieval Data Headings	207
5.3	Example of Matching Events	217

CHAPTER 1

INTRODUCTION AND LITERATURE REVIEW

1.1 Introduction

No one solution exists to tackle the problem of terrorism. However, many approaches have been found, and are being developed, which can help build a clearer picture of the many facets of terrorism and so yield valuable insights into how it can be reduced. Amidst the plethora of sociological approaches a new tool is being harnessed for the depth of understanding it can provide, namely, mathematical modelling. Applied successfully in many other real world problems, from understanding the nature of earthquake occurrence (Hawkes, 1971) to uncovering the spatio-temporal pathways by which contagious diseases spread (Verity et al., 2014), mathematical modelling offers a way to condense hugely complicated phenomena into numerical quantities with real-world interpretations. In this dissertation it is shown that through the use of carefully considered and context relevant mathematical models a powerful tool can be created with the ability to numerically quantify the structural and tactical approaches of terrorist organisations in near real time.

1.2 Motivation

As discussed by Matusitz (2013) finding a universally accepted definition of terrorism is difficult. One definition that Matusitz states as being the most widely accepted is that “terrorism is the use of violence to create fear for political, religious, or ideological reasons.”

Terrorism has been a major security concern for governments around the world

for decades (Hanhimäki and Blumenau, 2013). A critical turning point in global attitudes that brought terrorism to the forefront of modern security challenges was the attacks of September 11, 2001 commonly referred to as 9/11. After this event research into the many facets of terrorism garnered much more attention as governments around the world sought ways to stop terrorist attacks occurring and ensure public safety (Chen et al., 2008). Despite this increase in research Horgan and Braddock (2012) finds that whereas 80% of studies in forensic psychology and 50% of studies in criminology use some type of statistical analysis only 25% of terrorism studies following 9/11 utilise some kind of statistical argument. The situation appears even more strikingly in research looking at ways to tackle terrorism where Lum, Kennedy and Sherley (2006) found that only 7 out of 20,000 terrorism studies had a sufficient evaluation of counter-terrorism approaches. This lack of research however seems odd particularly at a time when the digital era has produced huge quantities of data that mathematics is best placed to analyse.

One way in which the deficit of statistics in the terrorism literature is being addressed is via the use of mathematical modelling. This field of applied mathematics aims to create a bridge between the well developed tools and techniques that have been developed over many years in pure mathematics and real world problems. The use of mathematical tools to model real-world phenomena is prevalent in multiple fields of academic research. One area in particular where modelling has been found to yield important insights is in the study of social sciences where research in economics (Kendall, 1968; Ostaszewski, 1993; Shubik, 1967), psychology (Bush, 1956; Kempf and Repp, 1977; Restle, 1971) and sociology (Abell, 1971; Lenhard, Küppers

and Shinn, 2006; Sorensen, 1978) has gained much from interpreting problems in mathematical terms. In these fields the main outcome of studying mathematical models is that they provide a platform where policies regarding a range of issues from financial crisis to war and terrorism can be tested (Ball, 2012).

An emerging area where modelling has been particularly successful is the fields of crime and security studies. As discussed by Memon et al. (2009) applied maths proved invaluable during the Second World War from cryptography to crack and secure communications to game theory to pre-empt opponents moves. Despite this Memon et al. (2009) point out that government uptake of mathematical modelling for decision making has been slow in recent times. But as data collection and availability garners more attention by law enforcement agencies and governments the world over new opportunities are arising to apply well developed mathematical techniques to explore data in ways not possible previously (Johnson, Restrepo and Johnson, 2015). This approach has been successful in revealing new information about several different event types including burglary (Bernasco, Johnson and Ruiter, 2015; Davies and Bishop, 2013; Short et al., 2009), gang related violence (Egesdal et al., 2010; Hegemann, Lewis and Bertozzi, 2013; Stomakhin, Short and Bertozzi, 2011), insurgency (Arney and Arney, 2012; Braithwaite and Johnson, 2012; Farley, 2007) and civilian deaths in conflicts (Lewis et al., 2012). The aim of these sort of studies is to use available data, such as, times of a gang or terrorist attack, and build a mathematical representation which reveals the underlying forces leading to the events. It is this ability of mathematical models to provide deep insights with relative ease that provides the strongest case for their use. In particular, studying terrorism from the

perspective of any academic field is difficult due to the clandestine nature of terrorist groups who rely on secrecy in order to operate (Silke, 2001). But mathematical modelling can overcome some of these barriers by making full use of open source data, such as, times of an attack reported in a news article, to shed light on the different components of how groups are operating.

Where mathematical modelling differs from other forms of investigation is in its flexibility. When modelling a specific topic it is possible to incorporate the ideas of many different theories explaining various components of the topic. Moreover, mathematical modelling offers an iterative method of research whereby as models reveal new insights about existing notions these can be fed back into the model to improve its ability to capture true underlying dynamics. In other words modelling provides a way to both explore existing ideas and generate new ones (Epstein, 2008). This concept is especially widespread in the use of simulation studies. Here mathematical interpretations of the theoretical ideas surrounding terrorism can be studied using computer programs which simulate the results of different theories. An example of this type of approach is agent based modelling of terrorist attacks. Research conducted by Park et al. (2012) examines the use of agent based modelling in the context of simulating terrorist attacks at public venues to study how crowds might respond. This type of model leads to the development of effective practises for emergency response teams responding to terrorist incidents.

Another strength of using models is their predictive abilities. In particular, mathematical models can be used to capture the patterns of previous event data and try to project these patterns to simulate possible future occurrences. This was the ap-

proach taken by Clauset and Woodard (2013) when studying the probability of large scale terrorist events. By analysing historical database of terrorist events Clauset and Woodard were able to find a power-law distribution for the probability of a large scale event. With this distribution it was discovered that the worldwide probability of an attack on the scale of 9/11 since 1968 was 11-35%. Moreover, the power-law model also indicates a 19-46% chance of a large scale attack in the next decade. From this research Clauset and Woodard conclude that it is possible that global political and social drivers can be used to study terrorism without the need to study localised dynamics.

What this last example in particular shows is that there is a possibility to study terrorism on a macro level by using mathematical models. In this thesis this idea will be the foundation upon which all analysis is conducted. The necessary theory behind this shift from a micro to macro level will now be explained.

1.3 Crime Science and Terrorism

Before attempting to use mathematical modelling in an applied setting it is first necessary to understand the theoretical foundations of the problem to be studied. In the field of crime and security studies the underlying principle which has made modelling so successful is a shift from micro level analyses to ones on a macro level. Attention is moved away from individual dispositions or social drivers towards an exploration of environmental conditions which facilitate criminal or terrorist activity. This departure from a more traditional criminological person-centred understanding

is known as crime science. The theoretical foundations of this environmental approach are rooted in the rational choice perspective and the routine activity theory of crime.

1.3.1 Rational Choice Perspective

This approach advocates the idea that great potential in the field of crime reduction exists through a better understanding of the cost-benefit analysis undertaken by offenders (Clarke and Cornish, 1985). In particular, it is argued that those who commit crimes process information in a rational way in order to make decisions about a particular offence, such as, what target is appropriate and at what time. This perspective looks in detail at specific crimes and the decisions surrounding the involvement and commission of that crime. The driving force behind each stage of crime commission is therefore a cost-benefit analysis where maximum rewards are aimed for at minimum cost. However, as noted by Simon (1990) human rationality is bounded so that decision making is not assumed to be perfect. A similar theory is also found to hold in the case of terrorism with Caplan (2006) finding that a rational choice approach can be used to model terrorist acts.

1.3.2 Routine Activities Theory

Routine activity theory places a rationally motivated offender within an environmental context and explains how these two link up and lead to crime occurrence (Cohen and Felson, 1979). This approach to understanding crime was the result of an observation by Cohen and Felson that a crime rise following World War II may in

fact have been resulting from a societal transformation whereby the patterns of daily life were being fundamentally altered triggering new opportunities for crimes to be committed. Cohen and Felson hypothesised that crime was in fact the result of convergence in space-time of offenders and targets in the absence of a capable guardian (Brantingham and Brantingham, 1981; Cohen and Felson, 1979; Eck, 2010). As a result of this way of thinking about how crimes are committed two key ideas emerge. One conclusion that is obtained is that for an offence to take place an offender, target/victim and lack of guardian must occur simultaneously in time and space. As such several authors (Bowers and Guerette, 2014; Clarke and Eck, 2003; Cohen and Felson, 1979; Eck, 1994; Tilley, 2012) postulate that crime prevention requires one of three components: a handler for the offender, a manager for the place or a guardian for the target/victim. Identifying which combination of these preventative measures will be successful and what approaches to take is context dependent and what works in a particular setting may not be directly applicable elsewhere (Bowers and Guerette, 2014). In the work of Clarke and Newman (2006) ideas relating to the applicability of the routine activities theory to terrorism can be found, such as, terrorists using local knowledge to determine the best targets and locations for attacks.

1.3.3 Crime Pattern Theory

Combining the fundamental ideas of the rational choice perspective and routine activities theory leads to a macro level analysis of crime known as crime pattern theory. The main hypothesis of this theory is that crimes do not occur randomly or

uniformly but are actually the result of the combination of environmental and social factors (Brantingham and Brantingham, 1993a). In particular, crime forms patterns in space and/or time as a result of offenders aiming to maximise their cost-benefit analysis within their environment of operation built from their routine activities.

Pattern theory can also be extended to the problem of terrorism. Using the same logic that combines the rational choice perspective and routine activity theory of crime Lafree, Morris and Dugan (2010) finds strong evidence for a patterned nature of terrorism. Some examples of the terrorism patterns found include that 10 countries in the study accounted for 38% of all terrorist attacks and that 32 countries were found to have more than 75% of all attacks.

1.3.4 Crime Reduction

The motivation to formulate and apply the principles of crime science is that by uncovering the macro level behaviour of criminals there exists great potential for crime reduction with minimal resources. Specifically the rational choice perspective, routine activities theory and crime pattern analysis of crime identify proximal circumstances, crime opportunities and situational factors as causes of crime (Smith and Tilley, 2005). Therefore, by targeting these common areas which lead to crime it is theoretically possible to make large reductions in the levels of criminal activity. This has been shown to hold true in a number of examples including crime prevention through environmental design (Atlas, 2008; Crowe, 2000; Draper and Cadzow, 2004), situational crime prevention (Clarke, 1997; Heal, Laycock and Great Britain. Home Office. Research and Planning Unit, 1986; Smith and Cornish, 2003) and

problem-oriented policing (Goldstein, 1990; *Problem-oriented policing* 2003; Scott, 2000). Application of these sort of approaches is also beginning to occur in the field of terrorism studies (Atlas, 2003; Clarke and Newman, 2006; McGarrell, Freilich and Chermak, 2007).

1.3.5 Patterns in Terrorism

Crime science provides the fundamental theories which motivate the study of spatial and temporal patterns in terrorist activity. Some studies already exist looking for these patterns and in order to provide context for this thesis background to research in this area is now provided.

A widely studied area where temporal clustering is of interest is in the area of human behaviour. One simple model that can be used for human activities is a Poisson model assuming random distribution of events in time. However, in research conducted by Barabási (2005) a discussion is presented highlighting that in fact a non-Poisson model is more relevant for human behaviour modelling. Moreover, Barabási argues for a more clustered nature of human events with a more accurate model being heavy tailed with most events occurring closer together in time. This type of idea has begun to find traction in the field of terrorism where temporal patterns have been explored. In multiple studies (Bohorquez et al., 2009; Clauset and Gleditsch, 2012; Clauset and Woodard, 2013; Clauset, Young and Gleditsch, 2007; Johnson et al., 2013; Johnson et al., 2011; Picoli et al., 2014) temporal clustering has been found with inter-event times being shown to have a heavy tailed distribution as predicted by Barabási (2005).

In addition to studies considering the temporal patterns of terrorism there has also been advances in investigating spatial and spatio-temporal patterns. These types of studies stem from similar ideas in crime science concerning the observation of a contagion effect. For example, for residential burglaries Johnson and Bowers (2004) find that following a residential burglary there is an increased risk for homes with 300-400 metres of the initial event for 1-2 months. In the context of terrorist activity these sorts of patterns have been researched in numerous settings. In Iraq Townsley, Johnson and Ratcliffe (2008) and Braithwaite and Johnson (2012) have shown that the spatial and temporal components of insurgent activity are more clustered than a hypothesis assuming event independence. Similarly results are reported in the research of Lafree et al. (2012) studying the ETA (Euskadi ta Askatasuna) terrorist group in Spain. It was found that the ETA group operated during two distinct phases. During the first phase attacks were concentrated in the Basque region of Spain. However, after the organisation shifted its focus to a war of attrition the spatial clustering shifted outside of this region. It was also found that due to further distances when attacking outside the Basque region the temporal length between terrorist events increased. Spatial and spatio-temporal results suggesting a clustered nature of terrorism have been observed in numerous other countries including Afghanistan and Pakistan (O'Loughlin, Witmer and Linke, 2010; Zammit-Mangion et al., 2012), Israel (Berrebi and Lakdawalla, 2007; Kliot and Charney, 2006), the Northern Caucasus in Russia (O'Loughlin, Holland and Witmer, 2011; O'Loughlin and Witmer, 2012) and the US (Cothren et al., 2008; Smith, 2008).

1.4 Research Questions

Despite the existence of some research into patterns of terrorism there still remains a gap in the literature for developing mathematical models of terrorist events. More specifically, there is an opportunity for more research to be conducted into the spatial, temporal and spatio-temporal patterns behind terrorist attacks. In this thesis these opportunities will be seized upon by considering a range of mathematical models which can capture clustering in terrorism data. Therefore, the first research question that will hopefully be addressed is as follows

1. What does mathematical modelling reveal about the spatial, temporal and spatio-temporal characteristics of terrorist attacks?

An essential component of mathematical modelling is data. In particular, if the underlying data used in a study is inaccurate in some sense then the conclusions drawn from models may be misleading. Hence, alongside model selection it is also necessary to justify the data source employed in a study. In this thesis this issue is addressed by considering the difference in outcomes using human and computer created databases. This leads to the following research question

2. How does the method of data collection impact the mathematical models used to study terrorism and the what implications does this have on the validity of conclusions drawn from such models?

A principle objective in terrorism research is to derive outcomes with policy relevance in the field of counter-terrorism. As such creating tools which can be used by

academics and practitioners to gain insight about terrorist groups is important and incredibly useful. However, there is a lack of such tools in the literature so in this thesis a final question that will be researched is

3. Is it possible to detect in real-time the behaviour of terrorist groups by using open source data and mathematical models?

1.5 Case Studies

As with any mathematical modelling it is necessary to have a clear understanding of the focus of the problem. In this thesis the central case studies that will be examined revolve around two terrorist groups known as the Provisional Irish Republican Army (PIRA) and Al Shabaab (AS). For context a brief history of each of these organisations will now be presented.

1.5.1 Provisional Irish Republican Army

There is a plethora of historical accounts concerning the origin and evolution of the Provisional Irish Republican Army (PIRA or Provisional IRA) (Bell, 2000; Coogan, 2002; English, 2004; Moloney, 2003; Patterson, 1997; Taylor, 1998). A summary of these texts is provided here.

The IRA is an organisation born of the violent Easter Risings in 1916. Although reincarnated over many years with different names arguably its most prominent period was during “The Troubles” in Northern Ireland between 1969 until 1998. During this period many different groups existed including the Official IRA, the Irish

National Liberation Army, the Continuity IRA, the Real IRA and the Irish People's Liberation Organisation. However, it was the group known as the Provisional IRA which was the most prolific organisation.

The main objective of PIRA was to have Northern Ireland removed from the United Kingdom and instead have a 32 county Republic of Ireland. Initially PIRA aimed for a quick offensive against British Forces. However, in 1977 the group realised that a change of approach was necessary and instead PIRA began preparing for a war of attrition. This new objective was reflected in a publication of PIRA known as the "Green Book" (O'Brien, 1999) which called for a "war of attrition against enemy personnel which is aimed at causing as many casualties and deaths as possible so as to create a demand from their people at home for their withdrawal". According to the sociological theory PIRA went through five distinct phases during its armed struggle (Asal et al., 2013):

- **1969-1976 - Phase 1:** During this phase the organisation was arranged in a military style consisting of brigades, battalions and companies. Within this army structure each unit of the organisation was given responsibility for a specific geographical area of combat.
- **1977-1980 - Phase 2:** A cell-based structure was adopted. This approach was characterised by PIRA fracturing into small groups of members known as Active Service Units (ASUs) (Horgan and Taylor, 1997). The aim of this re-structuring was to improve the organisation's secrecy by making it harder to infiltrate. This change in structure was successful with 465 fewer charges for paramilitary activity within a year (Smith, 1997). During this phase new

leaders were also appointed for the organisation including Gerry Adams and Martin McGuinness (Moloney, 2003).

- **1981-1989 - Phase 3:** This period began with the Hunger Strikes by Provisional IRA members protesting against the conditions of their incarceration. A catalyst moment during this period was the death of a PIRA member known as Bobby Sands who had been elected to Westminster whilst in prison and died on hunger strike (English, 2004). This incident resulted in a rise of sympathy for PIRA and its political wing Sinn Féin. As a result the Republican campaign moved into the political arena through the Sinn Féin party who now had similar levels of prestige as their militant wing, PIRA.
- **1990-1994 - Phase 4:** Secret meetings occurred involving top ranking PIRA leaders who were negotiating a ceasefire with the British Government.
- **1995-1998 - Phase 5:** Finally the peace talks were announced and a ceasefire ratified in the Good Friday Agreement. For many this signalled the end of “The Troubles”.

One of the main strengths of PIRA was its ability to develop and deploy improvised explosive device (IED) attacks. In the “Green Book” it states that the group aimed to have a “bombing campaign aimed at making the enemy’s financial interests in our country unprofitable”. In an extensive account of the evolution of IED usage by the PIRA Oppenheimer (2009) illustrates that both the physical make-up of the IEDs and the tactics used to employ them evolved greatly within the organisation. For example, PIRA developed 15 different versions of mortar bombs with

Oppenheimer, who is a weapons specialist, concluding about the Mark 6 mortar that “in just over a year, the IRA mortar had developed from something relatively primitive to an advanced weapons series”. This sentiment was further echoed by Ryder (2005) who described the Mark 15 mortar bomb as “perhaps the ultimate improvised weapon.” Alongside this extensive development PIRA also improved their tactics (Sutton, 1994) such as using secondary IED devices following initial attacks to catch the British Forces off guard (Oppenheimer, 2009).

1.5.2 Al Shabaab

As with the PIRA the organisation in Somalia known as Al Shabaab (AS) was formed from a complex history of events. Here a brief overview is provided giving a summary of the accounts of Anderson and McKnight (2015a), Hansen (2013), Marchal (2009) and Wise (2011).

The group known as Al Shabaab was initially part of a collective called the Islamic Courts Union (ICU). After many years of fighting between warring clans in Somalia the ICU helped by AS became a dominant force taking control of the countries capital, Mogadishu, in 2006. Fearing a “jihad” the Ethiopian government sent troops into Somalia in order to contain the expansion of the ICU. The Ethiopian forces had some success in reversing the gains made by the ICU. With these gains the ICU leaders began to retreat whilst members of Al Shabaab began launching a guerrilla warfare campaign against the Ethiopian forces. Al Shabaab managed to frustrate the Ethiopian efforts and this success lead to AS becoming an independent group using IED attacks to counter the Ethiopian army. Al Shabaab’s abilities to launch attacks

contributed to a withdrawal of Ethiopian troops in 2009. A replacement force to counter AS was made up of Ugandan and Burundian peacekeeping forces as part of the African Union Mission to Somalia (AMISOM). Despite a new force Al Shabaab remained, and still remains, undefeated. Moreover, since 2010 the AS group has become more internationally focused launching its first attack outside of Somalia in Kampala, Uganda. Al Shabaab have also launched high profile international attacks against Kenya following an invasion of Kenyan armed forces into Southern Somalia in 2011 to counter AS gains (Anderson and McKnight, 2015b). High profile incidents in Kenya include the Westgate Mall in Nairobi in 2013 in which 67 people were killed (Williams, 2014) and an attack on the Garissa University in 2015 which resulted in 148 deaths (Lyons et al., 2015). The international dimension is further enhanced by links between the group and al Qaeda since 2008. Furthermore, this link with al Qaeda has led to advances in the tactics used by Al Shabaab in areas such as assassinations, suicide bombings and IEDs (Shinn, 2011). Over the course of AS history to date there have been three main leaders

- **2006-2008:** Aden Hashi Ayro (killed by a US airstrike).
- **2008-2014:** Ahmed Abdi Godane (killed by a US drone strike).
- **2014-Present:** Abu Ubeyda.

1.6 Research Methodology

Having discussed the theoretical foundations for this thesis and presented the research questions to be studied it is now necessary to discuss the methods to be used. The

two types of models which have been studied are the Hawkes point process and geographic profiling. The theoretical backgrounds for each of these models is now provided.

1.6.1 Hawkes Process

One type of mathematical model which has recently shown great potential for studying security problems is known as a Hawkes self-exciting point process model. This type of model was first discussed by Hawkes (1971) in the context of mathematical modelling of earthquake occurrences. The essential idea behind such models is to capture the influence of past events on future occurrences. In the context of earthquakes Ogata (1988) demonstrated that self-exciting models could be used to understand historical ties between earthquake magnitudes thus opening up the possibility for predicting major future events. Mathematically such models are said to capture the clustering of events in time and space where the occurrence rate of future events depends on the occurrence of historical events (Daley and Vere-Jones, 2003).

Despite being a very useful model one of the drawbacks of the Hawkes process was its inability to capture other influencing events. For example, in the study of earthquakes it may be useful to consider the individual contributions of main and aftershock earthquakes to the rate of future earthquake occurrence. An attempt at solving this problem was touched upon in the initial exposition of Hawkes (1971) with some work done to develop a multivariate model. Further developments to the work of Hawkes were undertaken by Adamopoulos (1975) who expanded the analytical foundations of the multivariate model. With a multivariate Hawkes process, in

addition to events being self-exciting, other external factors may also stimulate future events through a method of mutual excitation. As was the case for the univariate Hawkes process the multivariate model has also been demonstrated successfully in the context of earthquake occurrences by Ogata, Akaike and Katsura (1982).

It is by analogy of past events stimulating future ones that has led to the application of the Hawkes process to crime and security problems. This idea stems from observations similar to the burglary example of Johnson and Bowers (2004) where following a residential burglary homes within 300-400 metres were at an increased risk of burglary for 1-2 months. This shows that burglary risk decays in space and time a result found in numerous studies (Bernasco, Johnson and Ruiter, 2015; Sagovsky and Johnson, 2007; Townsley, Homel and Chaseling, 2000).

A batch of point process research which has proven particularly successful is related to gang violence. In a student research paper by Egesdal et al. (2010) the Hawkes process model is used to examine the retaliatory nature of gang violence in Los Angeles. The authors discovered from this investigation that the historical dependence of the Hawkes process outperformed a simple Poisson process. A Hawkes process was found to capture the real world patterns of violence between pairs of gangs. Moreover, this type of model, similar to findings in earthquake research, was used as a predictive tool to study future violence. One of the limitations in the work of Egesdal et al. (2010) was that only a univariate Hawkes process was employed. In this case univariate means that only violence of one gang directed towards another was studied neglecting the influence that fighting between multiple gangs had on any one groups attacks. Later this research was developed using a multivariate

Hawkes process. By a similar analogy of multiple earthquake interactions Short et al. (2014) considered the use of the multivariate Hawkes process in the context of multi gang violence. The additional complexity of a multidimensional model was rewarded with further insights into the nature of gang violence in LA leading to better policy recommendations for tackling this crime problem.

Despite these initial successes of the Hawkes self-exciting models gaps remain in their application to security and crime science problems. One field in particular which may have much to benefit from point process models is terrorism. By analogy to the attacks and counter-attacks between gangs and those between security forces and terrorist groups intuition suggests that there may be some new information to be gained about terrorism by applying similar methods to those just discussed. Some recent attempts in this direction have been made. In research conducted by Lewis et al. (2012) violent civilian deaths in Iraq between 2003 and 2007 were found to be well modelled by a Hawkes process. In particular, Lewis et al. found that as much as 37-50% of all violent events were the result of self-excitation. Moreover, the self-excitations that were observed were found to extend between two to six months. From this approach precise policy relevant suggestions were obtainable indicating ways in which a serious security problem like violence against civilians could be negated. Another example looking at terrorist event data between 1970-1993 in Northern Ireland by Mohler (2013) showed that the data did indeed exhibit the property of historical dependence. Again these sort of findings can be used to stimulate new ideas to approach the problem of tackling terrorism. Many other studies with these and similar types of observations have been conducted including

cases in Afghanistan (Zammit-Mangion et al., 2012), Indonesia (Porter and White, 2012), Israel and Northern Ireland (Mohler, 2013) and the Philippines and Thailand (White, Porter and Mazerolle, 2012). The result of these studies is that point process models which capture temporal patterns provide a better fit to terrorism data than simple Poisson process models. This indicates that there is much to be gained from using point process models to study temporal patterns in terrorist activity. Moreover, another outcome of these types of studies is that even in the very difficult security environments such as conflicts or illegal terrorist organisations it is possible to use event data to understand underlying dynamics causing events to evolve in time and space.

Hawkes point process models will be used in this thesis to contribute to the research questions via the two case studies of the Provisional IRA and Al Shabaab. In particular, the following hypotheses will be tested

Hypothesis 1 *Timings of attacks by both the Provisional IRA and Al Shabaab are past dependent, whereby, previous attacks can trigger future attacks similar to earthquakes and aftershocks.*

Hypothesis 2 *The temporal patterns of attacks by the Provisional IRA differ in each of the organisation's phases corresponding to structural and tactical shifts identified by social scientists.*

In addition to modelling historical dependence in timeseries there can also be a need to capture other dynamics. One particular example is the relationship between crime occurrence and weather patterns.

As described by Cohen and Felson (1979) crime occurs as a result of opportunities governed by environmental conditions. For example, an offender may be more likely to commit a crime when temperatures are warmer compared to times when there is rain. This understanding of crime through routine activities leads to the idea that there is a great potential for reduction in criminal events by focusing on environmental factors. For example, Cohn (1990) finds assaults, burglary, collective violence, domestic violence and rape are positively correlated with temperature. Similarly, Pakiam and Lim (1984) reports that in Singapore crime against the person increases with increasing temperature and climate comfort measures. Moreover, Ranson (2014) elucidates a casual link between rises in multiple crime categories and climate changes.

Extending Cohen and Felson's routine activity approach to studying crime comparable results have been observed in terrorist attacks (Clauset and Woodard, 2013; Townsley, Johnson and Ratcliffe, 2008; Zammit-Mangion et al., 2012). Also by similarity to studies of criminal events terrorism has been shown to have cyclical dynamics. Enders, Parise and Sandler (1992) demonstrate that terrorist events, such as, bombings, hostage events and assassinations have cycles ranging between 21 to 54 months. Furthermore, threat of attacks were found to have a seasonal trend of 11 months, which, the authors attribute to onset of tourism periods. More evidence of terrorism cycles was observed by Weimann and Brosius (1988) where contagion dynamics as well as a constant 1 month periodicity were modelled with a first-order moving average.

To extend the literature on cyclical dynamics in terrorist attacks, and their links

to weather patterns, in this thesis rainfall patterns will be incorporated into the Hawkes process via a seasonal background rate. An application of this extended Hawkes model will be to study the relationship between AS attacks and rainfall in Somalia. In particular, Muchiri (2007) describes rainfall as the “defining characteristic of the climate” in Somalia with two main rainfall seasons. The first season is known as Gu and occurs between April and June. A second rainy seasons is named Deyr and commonly appears between October and December.

As with the routine activities description of crime it may occur that AS attacks are influenced by the onset of rainy seasons. For example, rain may make mobilising attacks more difficult if flooding occurs. On the other hand, the availability of water sources may provide AS with the ability to sustain attacks due to easy access to resources. The hypothesis that will be tested to determine if rainfall influences AS attacks is

Hypothesis 3 *The temporal patterns of attacks by Al Shabaab have a 3 month cycle corresponding to wet and dry seasons in Somalia.*

The previously stated Hypotheses 1-3 concern the ability to model various terrorist dynamics using the Hawkes process model. In real world situations where counter-terrorism practitioners are attempting to disrupt and prevent terrorist attacks a real-time analysis is essential. Therefore, combining the research outcomes of using the Hawkes process model the following hypothesis will be addressed

Hypothesis 4 *A Hawkes process model can be used to produce a real-time analysis tool to study terrorist groups.*

1.6.2 Geographic Profiling

Another area of mathematical modelling which has shown potential in the field of crime and security studies is geographical profiling (GP). As has already been discussed illegal criminal activities result from legal activities of everyday life (Brantingham and Brantingham, 1993a; Cohen and Felson, 1979) and that crimes occur when three environmental aspects, offender, target and place converge (Brantingham and Brantingham, 1981; Cohen and Felson, 1979; Eck, 2010). This motivates the idea that by studying where in space a series of crimes are committed there is potential to hone in on an area where the offender lives. The mathematical formulation of this concept is GP. The idea of GP is to incorporate sociological and criminological theory surrounding offender distance to crime into a mathematical model that allows predictions of most likely offender locations (Rossmo, 2000). In particular, by studying the spatial distribution of a series of crimes GP can be used to generate a risk surface which when overlaid onto the map of a geographical area can indicate which specific locations have the highest probability of being linked to a suspect. Such an idea has important implications for law enforcement as it allows limited resources to be targeted in the most beneficial manner. As discussed in the research of Verity et al. (2014) GP has been adopted by a variety of law enforcement agencies as a tool for determining the geographic location of criminals, such as, the Royal Canadian Mounted Police, the Los Angeles Police Department and the National Crime Agency. Moreover, GP techniques have been incorporated into several software packages for use in criminal investigations (Canter et al., 2000; Levine, 2006; Miller, 2003).

One of the main drawbacks of the GP model was that it only dealt with the case of searching for a single source (O’Leary, 2009; O’Leary, 2010; O’Leary, 2012) - the offender’s home. But in the work of Verity et al. (2014) the authors develop the GP approach by incorporating a Dirichlet Process Mixture Model (DPMM). The DPMM makes the extension of allowing multiple, and potentially unknown number, of sources to be found. Moreover, Verity et al. (2014) also provide a Markov Chain Monte Carlo (MCMC) method for applying the DPMM in situations with large datasets when the problem is otherwise numerically difficult. The mathematical foundations of the DPMM already exist (Green and Richardson, 2001) with the essential idea being that the model clusters data points without requiring knowledge of how many clusters exist a priori. This type of model has been highly effective in the field of biology where it has been used to discover multiple sources of malaria outbreaks (Verity et al., 2014) however it has yet to be used in a major study in crime science. Moreover, the current DPMM approach does not include a temporal component. Therefore, there is no way at present to use the DPMM to consider how the spatial patterns found by GP vary in time.

In this thesis the version of GP developed by Verity et al. (2014) will be used to study spatial and spatio-temporal patterns of improvised explosive device (IED) usage by the Provisional IRA. Specifically, the following hypotheses will be tested

Hypothesis 5 *The DPMM of GP can be used to determine the locations of IED making factories from the location of IED attacks.*

Hypothesis 6 *The DPMM can be used to assess the effectiveness of counter-terrorism strategies which aim to shut down IED factories.*

Hypothesis 7 *Incorporating a temporal component in the DPMM will produce new insights into the spatio-temporal patterns of IED attacks by the Provisional IRA.*

1.7 Data Issues

A major issue that has implications for the types mathematical models just discussed, and other modelling both within the confines of academic research and beyond, is obtaining useful and reliable data. This issue is becoming more significant as a plethora of digital data becomes available in an era of “big data”. A universally accepted definition of big data is difficult to find (Ward and Barker, 2013) but four principle components, known as the “four v’s” can be used as a useful way to characterise a big data problem - volume, variety, velocity and veracity (IBM, 2015). The first three of these v’s deal with magnitude covering the required memory for the data, multiple data sources and rate of data production. For research purposes the final v is of particular importance covering the need for data validity to be confirmed. In data analysis false conclusions drawn from inaccurate data has serious implications and thus there is a need to cross-reference data sources and collection methods in order to have confidence in research deductions.

A method which has great potential to deal with the problems of big data is found in a field of research known as data mining (Han, Kamber and Pei, 2012; Larose and Larose, 2014; Wu et al., 2014). Data mining can be defined as “a set of mechanisms and techniques, realized in software, to extract hidden information from data” (Coenen, 2011). Since human capacity is limited the aim of data mining is to

use computers in order to analyse large datasets. Furthermore, data mining techniques are often combined with information retrieval (IR) tools. IR covers methods of obtaining datasets. Finding ways to automate the IR process enables efficient data analysis frameworks to be constructed. This can be incredibly useful to researchers in a time when digital records of events are beginning to flourish.

Another area of research which can assist in the management of big data is known as machine learning (ML). ML is an umbrella term which covers tools and techniques developed to enable computers to derive value from data via learned structures and patterns (Alpaydin, 2014). The two main branches of ML techniques are supervised and unsupervised learning (Alpaydin, 2014; Masashi, 2016; Müller and Guido, 2017). Supervised ML provides inputs and outputs for an algorithm to learn. On the other hand, unsupervised ML provides only inputs to the algorithm. Using ML automatic tools can be constructed to conduct data analysis. Therefore, ML can complement and enhance data mining and IR techniques and vice versa.

One academic area in particular that can gain from all the above methods is terrorism studies. As briefly mentioned earlier terrorist groups rely heavily on secrecy in order to operate undetected thus making mathematical modelling of terrorist groups inherently difficult (Silke, 2001). However, by applying data mining, IR and ML tools to the big data available from of open sources that can be found relating to terrorism, such as, news articles covering terrorism events, it is possible to generate interesting research datasets and discover new insights into terrorism dynamics.

Utilising the tools and techniques of big data, data mining, IR and ML to obtain and analyse data allows researchers to compare analytical results arising from a

variety of data sources. In particular, comparisons can be made between computer generated datasets and those of more traditional human built databases such as the Global Terrorism Database (Lafree and Dugan, 2007) and the Armed Conflict Location and Event Data (ACLED) Project (Raleigh et al., 2010). In this thesis this idea will be taken forward by comparing the terrorism events obtained through the different data collection methods of an online automated information retrieval tool and the ACLED database. In particular, attacks conducted by the Al Shabaab group will be considered. As well as contributing to the second of the research questions discussed earlier studying the impact of different data sources will also provide insight into the best approach for creating a real time terrorism analysis tool. In particular, through comparison of an IR and human collated dataset the strengths and weaknesses of each should be revealed.

In this thesis the issue of examining data collection methods will be covered in the final research chapter. This is in opposition to the more natural approach of discussing data retrieval prior to any modelling or analysis. The motivating reason for this presentation stems from the exploratory nature of data collection techniques developed in this thesis. In particular, the aim will be to demonstrate new approaches that can be devised based on modern technological capabilities to obtain datasets. However, since these approaches are still in early development it would be considered unwise to rely on their usage throughout the entirety of the research in this thesis without more in-depth studies.

The central hypothesis for the collection methods comparison is

Hypothesis 8 *An automated information retrieval tool can be constructed to yield*

terrorism data comparable to a manually collated database across temporal, spatial and event description categories.

1.8 Dissertation Structure

With the models and case study groups defined the ultimate aim of this dissertation is to contribute to the available tools and techniques for academics and practitioners to use to study terrorism and terrorist groups. In particular, ranging from older terrorist groups such as the PIRA in Northern Ireland to more modern day examples such as the AS movement in Somalia one of the primary issues faced by law enforcement and intelligence agencies is how to understand and predict a group's actions. It's clear that there are many interactive methods by which terrorist groups can be monitored such as a variety of surveillance techniques and these are irreplaceable sources of information. However, it would be unwise to ignore the fact that open source data concerning the actions of terrorist groups can easily be used to shed a lot of light on what an organisation is doing. Such a field of study is beginning to bloom but a systematic way of looking at what such open source data actually shows is desired.

The layout of this dissertation will be in five chapters. Chapter 1 is the present chapter which is the literature review for this thesis. In Chapter 2 a case study of the PIRA is made. In this chapter the Hawkes process model and relevant computational details will be provided. Moreover, two techniques will be discussed which illustrate how the Hawkes process model parameters can be used to determine major changes in a terrorist group's structure based on a change point analysis of event times.

Having analysed a Hawkes process with just a constant background rate in Chapter 2 in Chapter 3 attention will be shifted to a more responsive underlying rate of events via the introduction of a seasonally varying Hawkes process. In addition, in Chapter 3 a modern and ongoing terrorist group in the form of Al Shabaab will be used as a new case study to test the wider applicability of the Hawkes process model. Next in Chapter 4 the method of geographic profiling will be introduced in more detail as will the computational tools necessary to use this type of model, such as, the Dirichlet Process Mixture Model and Markov Chain Monte Carlo method. This chapter will then move on to examine a use of GP in the context of studying counter-terrorism strategies by looking at the consequences of bomb factory closure on a terrorist group's attack profile. Next in Chapter 5 an automated data sourcing technique will be constructed which provides a real time data gathering tool. With data derived from using this tool it will be possible to explore issues surrounding the research impact of automated versus manual data collection methods. Finally, in Chapter 6 all of the research findings from this thesis will be summarised and possible areas of future research will be presented.

CHAPTER 2

HAWKES PROCESS MODELLING OF THE PIRA

2.1 Introduction

The Hawkes self-exciting point process model is used to model events where past occurrences can increase the rate of future occurrences (Hawkes, 1971). This structure of the Hawkes process made it useful in the study of earthquake modelling whereby the relationship between main shocks and aftershocks could be captured mathematically (Ogata, 1988). Picking up on an analogy between this type of behaviour and the retaliatory nature of gang violence the Hawkes process was then used to model gang attacks in Los Angeles (Egesdal et al., 2010; Short et al., 2014). In this chapter a similar theme will be continued by considering the self-exciting nature of terrorist attacks. The mathematical details of the Hawkes process and parameter estimation methods will be provided as well as details of how the model can be applied to terrorist attacks. With the mathematical formulations established a case study is then given considering improvised explosive device (IED) attacks by the Provisional Irish Republican Army in Northern Ireland. For this case study new methods will be introduced showing how the Hawkes process can be used to determine temporal evolutions of the PIRA over the time of its existence.

2.2 Mathematical Formulation of Point Process Models

In general terms a temporal point process is mathematical construct which is used to capture the properties of points in time. A common method to move from this

abstract idea to an applied setting is to study point processes via their conditional intensity function. In the work of Daley and Vere-Jones (2003) all the necessary definitions and details to construct the intensity function of a point process are provided and also summarised here.

The formulation begins by defining the counting process at time t denoted $N(t)$ as

$$N(t) = \sum_{t_i < t} \mathbf{1}_{t_i}([0, t)), \quad (2.1)$$

where $\{t_i\}$ are event times of the phenomenon being studied and $\mathbf{1}_{t_i}([0, t))$ is the indicator function which is equal to 1 if $t_i \in [0, t)$ and 0 otherwise. It is also necessary to define the history of events up to time t as

$$\mathcal{H}(t) = \{t_i | t_i < t\}. \quad (2.2)$$

The conditional intensity function $\lambda(t)$ associated to the counting process N and dependent on the history \mathcal{H} can now be described as

$$\lambda(t|\mathcal{H}(t)) = \lim_{\delta t \rightarrow 0} \frac{\mathbb{E}(N(t + \delta t) - N(t) | \mathcal{H}(t))}{\delta t}. \quad (2.3)$$

Qualitatively, λ is the expected number of events that occur at each time t . Analytically this model is unique if $N(t)$ is simple and finite. In this case simple means that all times t_i in $N(t)$ are unique meaning $t_i \neq t_j$ for $i \neq j$.

2.2.1 Poisson Process Model

Before introducing the Hawkes process it is initially beneficial to consider a simpler model known as a Poisson process (Ross, 2010). The Poisson process works by assuming time intervals of a fixed length have a constant probability of an event. Moreover, events described by a Poisson process are independent of the history of past events. For the case of a Poisson process the conditional intensity is given by a constant function

$$\lambda(t) = \mu \tag{2.4}$$

where μ is a positive constant.

2.2.2 Hawkes Process Model

The simplicity of the Poisson process does not allow for more complex dynamics to be captured. This last point is particularly prominent when the data being considered may have a past dependent nature. One possible solution to overcome this problem is to use a Hawkes self-exciting point process model. For a dataset satisfying such a model this means that a given event raises the chances of another event in the future.

For a given set of unique event times $\{t_i\}$ the conditional intensity function of the Hawkes self-exciting process is defined as (Hawkes, 1971)

$$\lambda(t) = \mu + k_0 \sum_{t > t_i} g(t - t_i), \tag{2.5}$$

where the response function g is defined as $g(t) = \omega e^{-\omega t}$ in this thesis. Other forms

for g are also possible (Hardiman, Bercot and Bouchaud, 2013; Mitchell and Cates, 2009; Ogata, 1999; Rambaldi, Pennesi and Lillo, 2015; Wang, Bebbington and Harte, 2010). Here the form of g employed by Egedal et al. (2010) and Lewis et al. (2012) is used based on the success of (2.5) to model violent conflicts. The right hand side of this model can be understood in terms of the three parameters (μ, k_0, ω) . The first is the background rate μ which simply describes the average rate of event occurrence. Next is the jump factor k_0 . This component indicates the rate increase of events following a past event with large values implying the underlying process is highly reactive to new event occurrences. Finally the value of ω controls the decay rate after a rise in the event rate. As this value occurs in the exponent the higher its value the shorter the timespan of influence an event has on future events. Moreover, as Lewis et al. (2012) explain the inverse parameter ω^{-1} describes the average time period over which an increased rate of events occurs.

Some assumptions need to be made for the model and its parameters to be correct mathematically and make sense in a real-world setting. Since the Hawkes process is dependent on the infinite past, events outside of the observation period could lead to incorrect results. Therefore, Rasmussen (2013) explains that the event times $\{t_i\}$ should be measured from time zero. Practically determining the start of events can be difficult and this issue will be addressed in more detail in Section 2.5.

For the conditional intensity function to be unique the event times must also be unique with $t_i \neq t_j$ for $i \neq j$ (Daley and Vere-Jones, 2003). How this assumption can best be met is context dependent and for each dataset introduced in this dissertation, and studied using the Hawkes process, a description of the data cleaning performed

will be provided. Essentially all methods of handling this assumption rely on some sort of timestamp randomisation as in Bowsler (2007). For event times which are measured as positive integers the simplest form of randomisation is to differentiate between two identical timestamps by adding a random number taken from a uniform distribution over the interval $(0, 1)$ as applied in this thesis.

In the work of Lewis et al. (2012) an assumption is made that the parameter values should all be positive. This assumption ensures that the interpretations of the model parameters provided earlier make sense in a real-world application. It is immediate that the background rate μ should satisfy this criteria since the lowest number of events is zero. It is possible that the jump factor k_0 can be negative producing a self-inhibiting effect as discussed in the work of Reynaud-Bouret and Schbath (2010), however, this case will not be considered in this thesis. The Hawkes process is such that following an event occurrence there is an increase in the rate of new events. Since it is unreasonable to expect such an increase to persist over an indefinite period of time the condition that $\omega \geq 0$ also makes sense.

According to the work of Hawkes and Oakes (1974) if $\mu > 0$ and

$$0 < \int_0^{\infty} g(t)dt < 1 \quad (2.6)$$

there exists a unique Hawkes process for the event times $\{t_i\}$. For the case of the self-exciting response function this condition becomes

$$\int_0^{\infty} \omega e^{-\omega t} dt = 1. \quad (2.7)$$

Therefore, the Hawkes process is well-defined if, and only if, $0 < k_0 < 1$.

Another interesting result in the work of Hawkes and Oakes is to express the Hawkes process as a branching ratio. If the above conditions on the response function hold then the process is stationary. This means that given a long history the process is time invariant. Mathematically this is expressed as

$$\mathbb{E}(\lambda) = \frac{\mu}{1 - \int_0^\infty g(t)dt}. \quad (2.8)$$

With this branching ratio interpretation the Hawkes process can be described as having parent events occurring at rate μ each with a probability of having offspring determined by the response function g (Rasmussen, 2013). Each offspring also has a chance of further offspring. However, the assumption (2.6) ensures that the process is non-explosive.

2.3 Parameter Estimation

To find the parameters of the intensity function for the Hawkes process a technique known as maximum likelihood estimation (MLE) can be employed (Ozaki, 1979). The aim of this method is to find parameters which maximise the loglikelihood function. For a set of event times $\{t_i\}_{i=1}^N$ the loglikelihood is given by (Rubin, 1972)

$$\log L(\{t_i\}; \mu, k_0, \omega) = \sum_{i=1}^N \log(\lambda(t_i)) - \int_0^T \lambda(t) dt, \quad (2.9)$$

where T is the end time of the observation period commonly taken to be the last event time $T = t_N$ (Ozaki, 1979).

In the work of Ozaki (1979) and Ogata (1981) it has been shown that the parameter set $\{\mu, k_0, \omega\}$ which maximises (2.9) is a good approximation to the true parameters of the underlying Hawkes process. Heuristically, the loglikelihood function can be considered as a comparison method. In this case the comparison is between the value of the intensity function at event times, the first term, and the value of the function at all times in the time interval considered, the second term. Therefore, when maximising the loglikelihood the parameters which give the best representation of the actual event data are found.

The first term on the right hand side of this equation can be substituted directly from (2.5). For the second term the integral can be simplified as follows

$$\int_0^T \lambda(t) dt = \int_0^T \mu + k_0 \sum_{t_i < t} \omega e^{-\omega(t-t_i)} dt \quad (2.10)$$

$$= \mu T + k_0 \sum_i \int_0^T \omega e^{-\omega(t-t_i)} \mathbb{1}_{\{t > t_i\}} dt \quad (2.11)$$

$$= \mu T + k_0 \sum_i \int_{t_i}^T \omega e^{-\omega(t-t_i)} dt \quad (2.12)$$

$$= \mu T + k_0 \sum_i \left[-e^{-\omega(t-t_i)} \right]_{t_i}^T \quad (2.13)$$

$$= \mu T + k_0 \sum_i \left[1 - e^{-\omega(T-t_i)} \right]. \quad (2.14)$$

Substituting this result in (2.9) yields the full form for the loglikelihood as

$$\log L = \sum_{i=1}^N \left[\log \left(\mu + k_0 \sum_{t_i > t_j} \omega e^{-\omega(t_i - t_j)} \right) + k_0 (e^{-\omega(T - t_i)} - 1) \right] - \mu T. \quad (2.15)$$

As was noted by Liniger (2009) this form of the loglikelihood can be improved to get a faster computational time by observing that

$$\lambda(t_i) = \mu + k_0 \sum_{t_i > t_j} \omega e^{-\omega(t_i - t_j)} \quad (2.16)$$

$$= \mu + k_0 \left[\sum_{t_{i-1} \leq t_j < t_i} \omega e^{-\omega(t_i - t_j)} + \sum_{t_j < t_{i-1}} \omega e^{-\omega(t_i - t_j)} \right] \quad (2.17)$$

$$= \mu + k_0 \left[\omega e^{-\omega(t_i - t_{i-1})} + e^{-\omega(t_i - t_{i-1})} \sum_{t_j < t_{i-1}} \omega e^{-\omega(t_{i-1} - t_j)} \right] \quad (2.18)$$

$$= \mu + k_0 \omega e^{-\omega(t_i - t_{i-1})} + (\lambda(t_{i-1}) - \mu) e^{-\omega(t_i - t_{i-1})}. \quad (2.19)$$

With this recursive method the first term on the right hand side of (2.15) can be calculated more efficiently thus improving computational speed.

There are a variety of methods by which the maximisation of this function can be achieved. For this dissertation the Python programming language has been used with the Scipy Optimize (*SciPy Optimize* 2015) routines being employed to find the optimal parameter set. It should be mentioned at this point that the routines provided in Scipy Optimize search for the minimum of a function and so the equivalent mathematical problem of finding the parameters which minimise $-\log L$ was undertaken. Both the ‘‘Nelder-Mead’’ and ‘‘differential evolution’’ methods were found to

work well for finding the Hawkes process parameters.

The Nelder-Mead method of optimisation relies on a downhill simplex approach (Nelder and Mead, 1965). For this method a moving polytope is considered. At each stage of the minimisation the function is evaluated on the polytope with the shape shifting according to where the lowest function value is obtained. This process continues until the polytope contracts sufficiently to declare convergence of the algorithm. An advantage of the Nelder-Mead algorithm is that it relies only on function calculations and not derivatives which can be complicated or even impossible to compute. However, a drawback of the Nelder-Mead method is that it is a local minimisation technique. Therefore the algorithm must be run multiple times starting at different points to find the optimal solution.

Published by Storn and Price (1997) the differential evolution method of optimisation uses a set condition to compare possible minimising values and find the optimum solution. In particular, a set of possible solutions is created and the algorithm moves around the function domain by combining these possible solutions. If one of these combinations is found to provide a lower function value then it is included in the set of possible solutions. This process continues until convergence is reached. Similar to the Nelder-Mead method the differential evolution approach relies only on function evaluations and hence there is no need to try and compute derivatives. The differential evolution approach can search a large space of possible solutions with the hope of finding a global minimum although this is not guaranteed. The differential evolution method in Python also requires user specified bounds on each parameter. In this thesis the bounds that were set were $\mu \in [0, 1]$, $k_0 \in [0, 1]$, $\omega \in [0, 10]$. The

bound on μ follows from the fact that the event data used was measured according to the day of event so the background rate should not be more than 1 per day. As discussed earlier a mathematical condition exists on k_0 limiting it to the range $[0, 1)$ to ensure the model is well-defined. Since ω^{-1} is the average length of time over which a series of attacks decays for data on the scale of day of event, as in this thesis (see Section 2.5.1), we require $\omega \leq 1$. However, to ensure that the decay in the Hawkes process is exponential a cut-off of 10 was decided in the search space of the parameter ω since higher values would imply the response function is sufficiently small to be deemed 0. In cases where the response function is deemed to be 0 the conditional intensity function becomes a Poisson process with a constant background rate. The observation was also made that for datasets with more than 500 points the differential evolution algorithm found the minimising parameters of the loglikelihood in one run. With datasets less than 500 entries multiple algorithm runs were undertaken. This final point was especially relevant when studying change points (explained in Sections 2.5.2 - 4) where there was a need to greatly reduce code runtime.

2.4 Model Goodness of Fit

To test whether the Hawkes process with parameters obtained via the MLE is a good fit to the original data a method known as residual analysis can be employed (Brown et al., 2002). The aim of this test is to look for differences between the fitted model and underlying data dynamics. The test relies on the following logic. Assume that a set of event times $\{t_i\}$ are from a Hawkes process with rate function λ and compute

the residuals using the formula

$$\tau_i = \int_0^{t_i} \lambda(t) dt \quad (2.20)$$

for each i . The $\{\tau_i\}$ form the residual process. These residuals should be distributed as a stationary Poisson process with unit rate (Papangelou, 1972). From this it can be deduced that

$$Y_i = \tau_i - \tau_{i-1} \quad (2.21)$$

$$= \int_0^{t_i} \lambda(t) dt - \int_0^{t_{i-1}} \lambda(t) dt \quad (2.22)$$

$$= \int_{t_{i-1}}^{t_i} \lambda(t) dt \quad (2.23)$$

$$= \int_{t_{i-1}}^{t_i} \mu + k_0 \sum_{t_j < t} \omega e^{-\omega(t-t_j)} dt \quad (2.24)$$

$$= \mu(t_i - t_{i-1}) + k_0 \int_{t_{i-1}}^{t_i} \sum_{t_j < t} \omega e^{-\omega(t-t_j)} dt \quad (2.25)$$

$$= \mu(t_i - t_{i-1}) + k_0 \sum_{j=1}^{i-1} \int_{t_{i-1}}^{t_i} \omega e^{-\omega(t-t_j)} dt \quad (2.26)$$

$$= \mu(t_i - t_{i-1}) + k_0 \sum_{j=1}^{i-1} [-e^{-\omega(t-t_j)}]_{t_{i-1}}^{t_i} \quad (2.27)$$

$$= \mu(t_i - t_{i-1}) + k_0 \sum_{j=1}^{i-1} [e^{-\omega(t_{i-1}-t_j)} - e^{-\omega(t_i-t_j)}] \quad (2.28)$$

are exponentially distributed. This implies that

$$U_i = 1 - \exp^{-Y_i} \quad (2.29)$$

$$= 1 - \exp \left(- \left(\mu(t_i - t_{i-1}) + k_0 \sum_{j=1}^{i-1} [e^{-\omega(t_{i-1}-t_j)} - e^{-\omega(t_i-t_j)}] \right) \right) \quad (2.30)$$

are uniform random variables.

Hence evidence for the fit of the Hawkes process can be gained by checking whether the U_i come from a uniform distribution. One such test is the Kolmogorov-Smirnov (KS) test (Massey, 1951). This test works by comparing the value of a test statistic calculated via the formula (Zar, 2014)

$$D_n = \max_k \left(\left| U_k - \frac{k-1}{N} \right|, \left| \frac{k}{N} - U_k \right| \right) \quad (2.31)$$

to a critical value D_α . Critical values can be found in many different sources of statistical tables e.g. O'Connor and Kleyner (2012). Evidence suggesting that the Hawkes process does fit the data is found if $D_n < D_\alpha$.

Another technique which is also relevant to judging the fit of a Hawkes process is the Akaike Information Criterion (AIC) (Akaike, 1974). The AIC is a method of comparing models applied to the same dataset. A better fitting model will have a lower value for the AIC given by

$$\text{AIC} = 2k - 2 \log L, \quad (2.32)$$

where k is the number of model parameters and $\log L$ is the MLE. At this point it should be mentioned that the AIC is simply a comparative tool and not a significance test. Therefore, some simple rules of thumb are (Burnham and Anderson, 2002) that difference of AIC value of $0 - 2$ suggests little difference between models, $4 - 7$

moderate evidence for difference, > 10 strong evidence for model difference.

2.5 Applying the Hawkes Process

In this section the techniques that have been outlined at the start of this chapter will be applied to study IED attacks by the Provisional IRA.

2.5.1 IED Data

Before applying the Hawkes process model it is first necessary to have an overview of the data being used. For this thesis the data studied was collected as part of a research project by Asal et al. (2013). The database used consists of information about IED attacks committed by the PIRA between 1969-1998. Details of the event fields in this database that are relevant to this thesis are provided in Table 2.1.

Table 2.1: PIRA IED Dataset Event Fields

Field	Values
Date	Year (1970-1998)/Month (1-12)/Day (1-31)
Location	{Antrim, Armagh, Belfast, Derry, Down, Fermanagh, Tyrone}
Target	{Political, Military, Police, Paramilitary, Government, Transport, Civilian, Foreign}

Since the temporal scale of the data is day of event an important processing step that was undertaken with this data was to ensure that the event times were unique as discussed in Section 2.2.2. One way this could have been achieved would be to simply take the first occurrence of each timestamp and discard duplicates. However,

this approach risks removing a large amount of data. A better option that was employed for this thesis was to take account of the additional spatial information in the database. In particular, each IED event time also had details of the county where the event took place. Therefore, when two timestamps were equal and in the same county the event was recorded only once. On the other hand, if two timestamps were equal but the counties where the event occurred were different then the timestamps were distinguished via the addition of a random number from a uniform $(0, 1)$ distribution as in Bowsher (2007). For this purpose Belfast was considered as a separate county. This is justified by looking at the command and functional structure of the PIRA which reveals that at the county and Belfast levels IED attacks were fairly autonomous (Horgan and Taylor, 1997).

2.5.2 Change Point Detection

The foundation for the research conducted in this thesis is a masters degree submitted by Tench (2014). In this masters degree the IED event database was studied via the Hawkes process. The method used was to separate the database according to the five phases of PIRA activity outlined in Section 1.5.1. The Hawkes process was then applied directly to the data corresponding to each phase. One of the problems with this approach however was that no investigation was made concerning how to manage the past dependence of the Hawkes process. In particular, as pointed out by Rasmussen (2013) if events occur outside the dataset being studied the parameters of the Hawkes process may not truly reflect the underlying dynamics. Therefore, when dividing the dataset into the five phases it is necessary to study the influence of the

events in past phases on future ones. This is especially relevant because the sociological timings marking the start of each phase may not be the same as mathematically determined boundaries. This leads to the idea of change point detection studied in the remainder of this chapter where properties of the Hawkes process will be used to study mathematically when the changes in the phases of the PIRA occurred. Two novel approaches have been designed in this thesis to study these change points.

2.5.3 Change Point Analysis Based on Sociological Boundaries

The first approach that will be used to study change points in the IED data is one developed and published by the author of this thesis (Tench, Fry and Gill, 2016) (see Appendix A). The method used is to consider the sociological boundaries in pairs and observe how far the influence of one phase extends into the next. Specifically for $i \in \{1, 2, 3, 4, 5\}$ consider phases i and $i + 1$. Then step one is to compute the MLE parameters of phase $i + 1$. Next the final point of phase i is added to $i + 1$ and the MLE parameters of this new dataset is calculated. This previous step is then repeated with the last two points in phase i added to phase $i + 1$ and so on until the MLE parameters for the combined dataset of phases i and $i + 1$ are computed. To determine where a phase change has occurred the residual analysis test as described earlier is used. For each of the parameter sets obtained the KS test statistic D_n is calculated and compared to the KS critical value D_α for that particular dataset. For significance it is necessary that $D_n < D_\alpha$. Hence a greater positive difference $\Delta = D_\alpha - D_n$ provides more evidence that the model is a good

fit for the data. Therefore, the dataset yielding this greatest difference is viewed as the most accurate representation of the past dependence of phase $i + 1$ into phase i .

Using this KS test method the mathematical boundaries between the phases are found to be as in Table 2.2. Corresponding to the new boundaries are model parameters in each phase and details of their goodness of fit as provided in Table 2.3. From Table 2.3 it can be seen that the models found all have significant fits under the KS test at the 95% level except the model in phase 1. Also in Table 2.3 are the results of applying a simple Poisson process as a comparison to the more complex Hawkes process model. The AIC comparisons show that the Hawkes process outperforms the Poisson process in all five phases.

Further to the quantitative results in Figures 2.1-2.4 are graphs illustrating how the parameters $\{\mu, k_0, \omega\}$ and $\Delta = D_\alpha - D_n$ values change for each consecutive pair of phases. It is also interesting to note some of the trends which appear in these figures. The graphs in Figure 2.1 are relatively flat indicating that the dynamics present in phase 2 extend quite far into phase 1. However, in Figure 2.2 there is a less distinct pattern suggesting much greater variability in the dynamics in phases 2 and 3. In contrast in Figure 2.3 there is a trend appearing after the identified change point. This implies that a Hawkes process corresponding to a dataset using the mathematical boundary has similar characteristics to a Hawkes process found from the dataset of phases 3 and 4 combined. In particular, this would suggest that mathematically these phases are very similar and potentially could exist as a single phase. In Figure 2.4 there is again a less distinctive pattern suggesting quite different dynamics between phases 4 and 5.

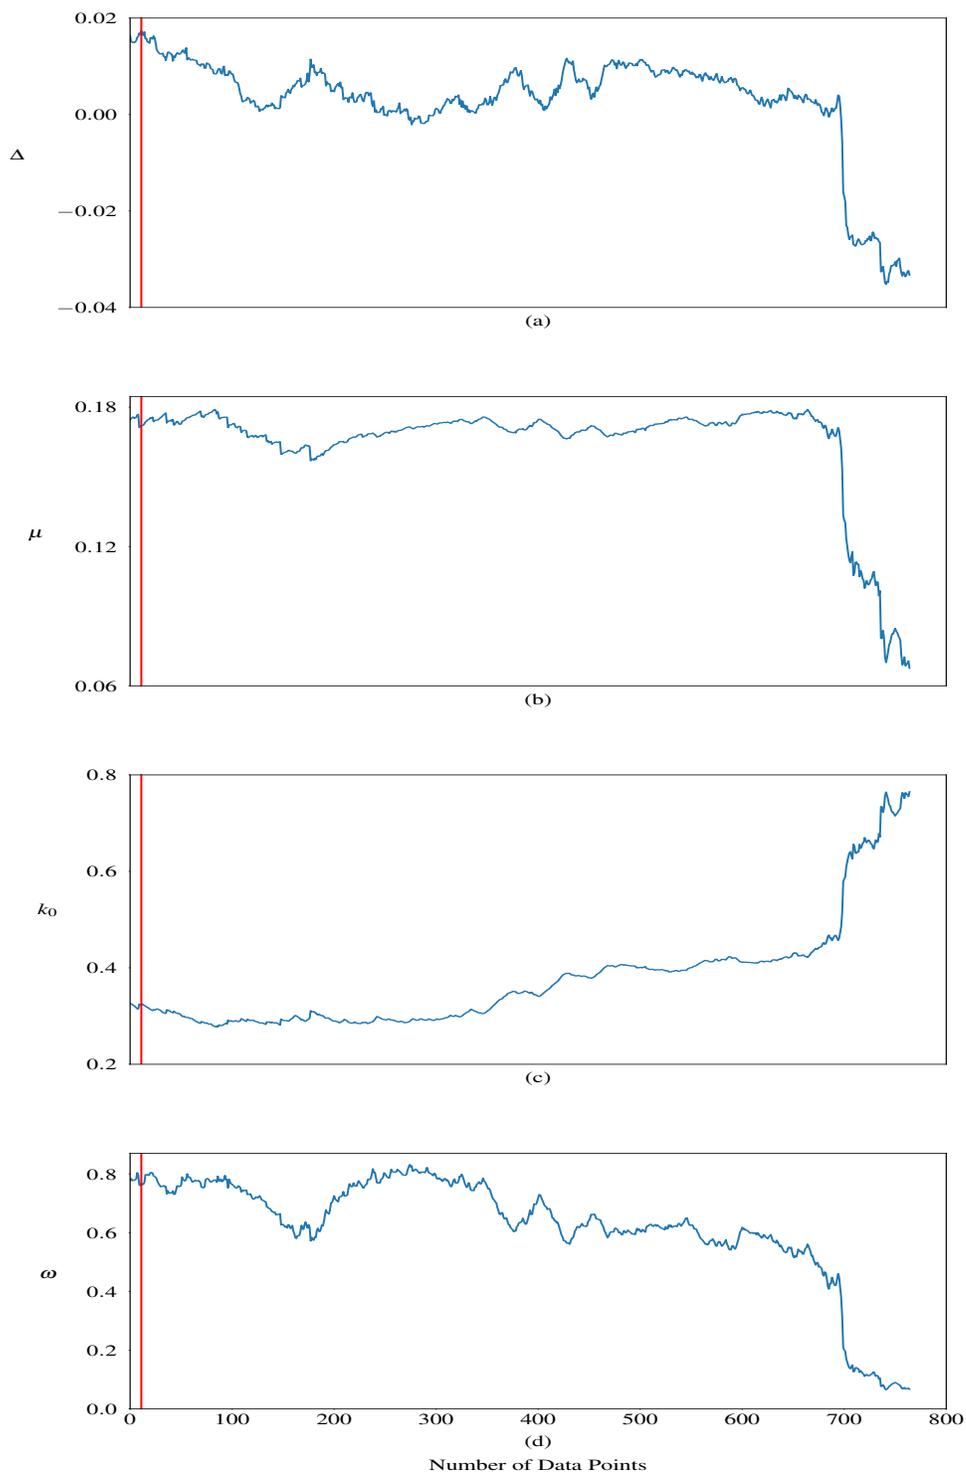


Figure 2.1: Figure (a) shows the values of $\Delta = D_\alpha - D_n$ for each point added from phase 1 to phase 2. Figures (b)-(d) show the corresponding changes in the MLE parameter values. The vertical line indicates the change point between the phases as determined by maximising $\Delta = D_\alpha - D_n$. The x-axis runs backwards in time representing each new pointed added from phase 1 to phase 2.

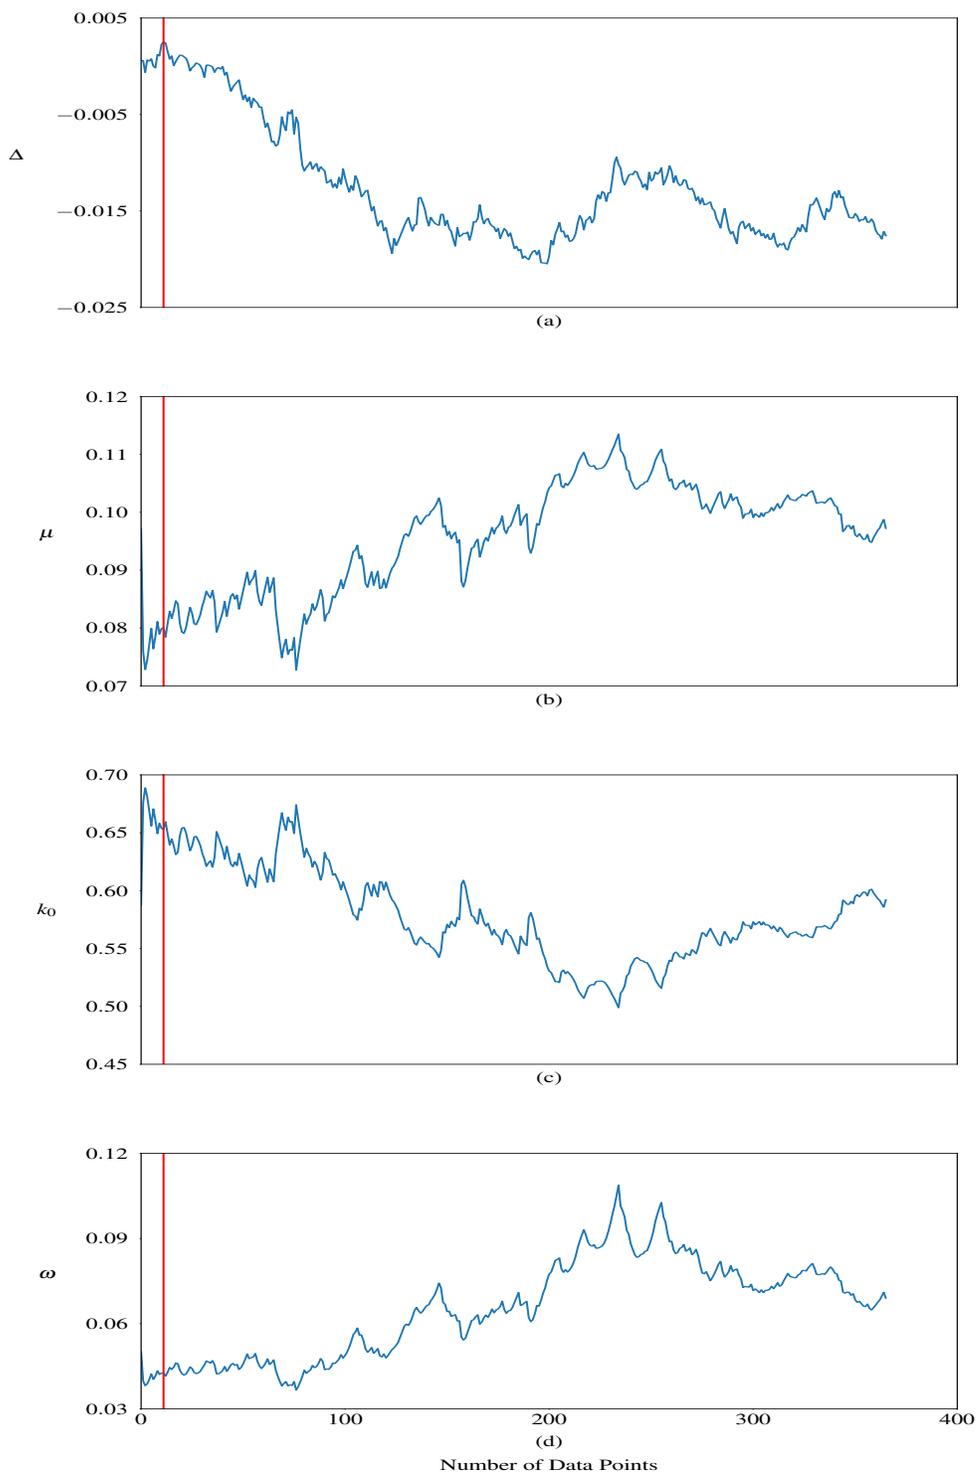


Figure 2.2: Figure (a) shows the values of $\Delta = D_\alpha - D_n$ for each point added from phase 2 to phase 3. Figures (b)-(d) show the corresponding changes in the MLE parameter values. The vertical line indicates the change point between the phases as determined by maximising $\Delta = D_\alpha - D_n$. The x-axis runs backwards in time representing each new pointed added from phase 2 to phase 3.

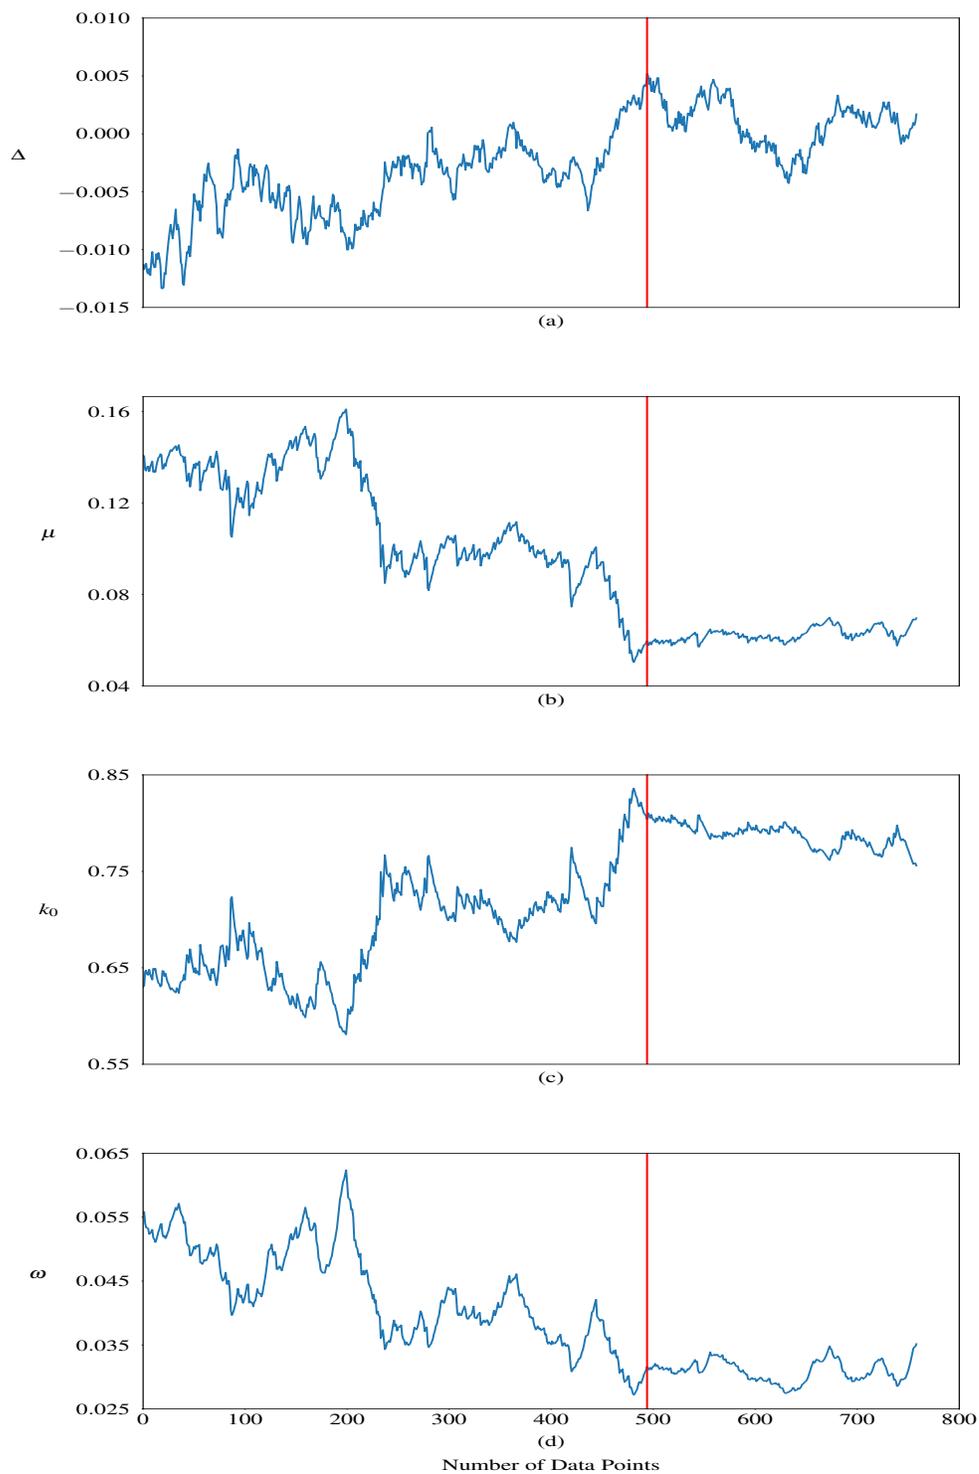


Figure 2.3: Figure (a) shows the values of $\Delta = D_\alpha - D_n$ for each point added from phase 3 to phase 4. Figures (b)-(d) show the corresponding changes in the MLE parameter values. The vertical line indicates the change point between the phases as determined by maximising $\Delta = D_\alpha - D_n$. The x-axis runs backwards in time representing each new pointed added from phase 3 to phase 4.

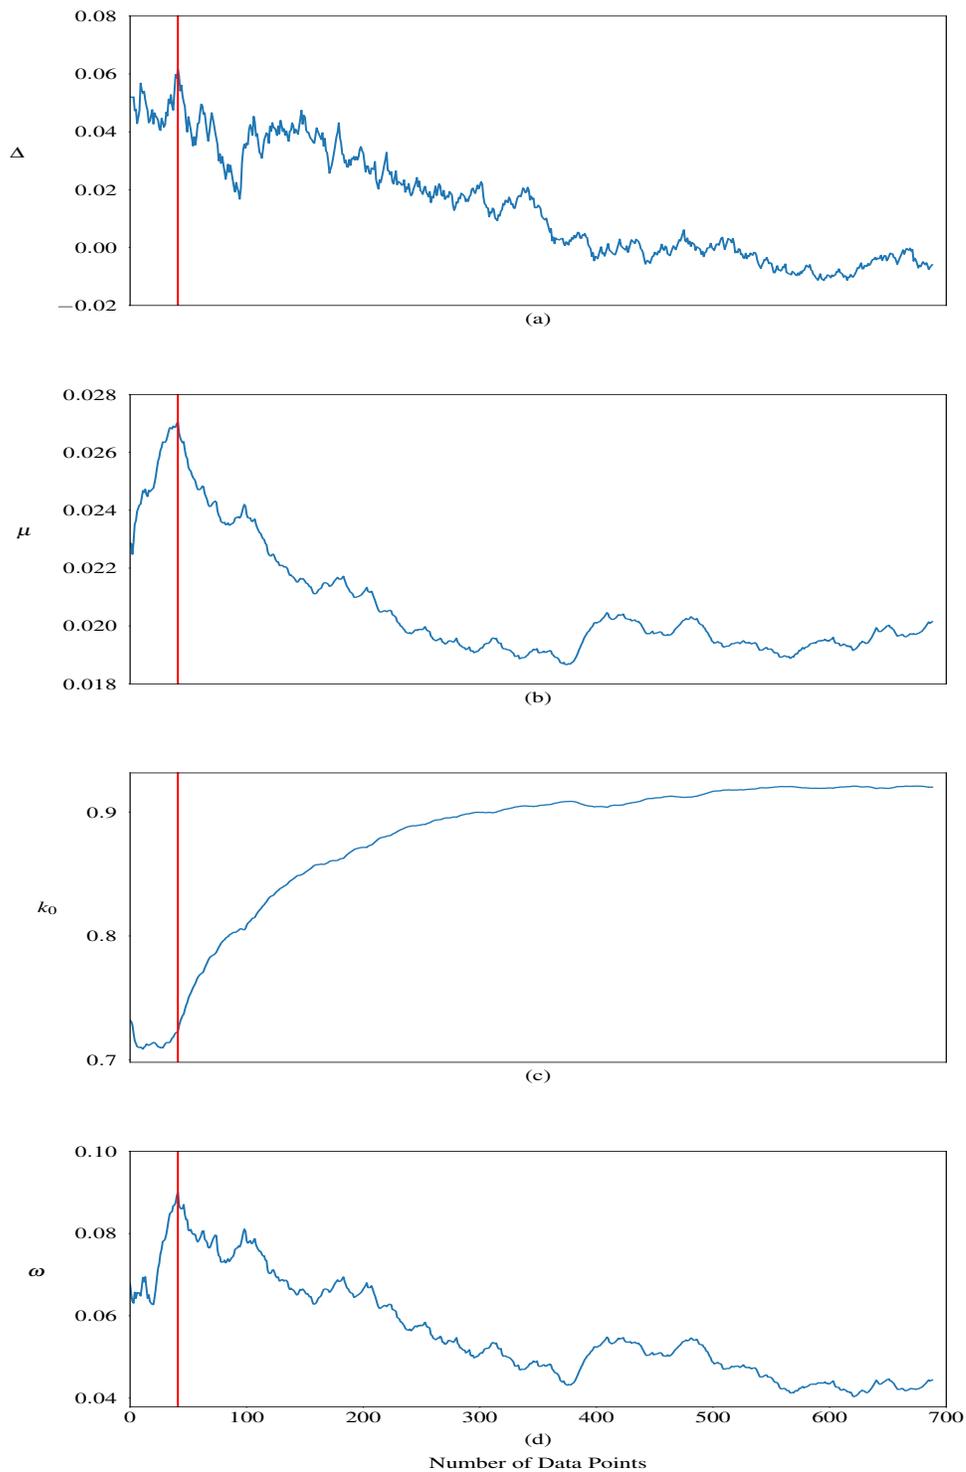


Figure 2.4: Figure (a) shows the values of $\Delta = D_\alpha - D_n$ for each point added from phase 4 to phase 5. Figures (b)-(d) show the corresponding changes in the MLE parameter values. The vertical line indicates the change point between the phases as determined by maximising $\Delta = D_\alpha - D_n$. The x-axis runs backwards in time representing each new pointed added from phase 4 to phase 5.

Table 2.2: Change Point Analysis Results with KS Test Method Boundaries

Phase	Number of Data Points Added	New Boundary	Original Boundary
1	-	27/01/1970	27/01/1970
2	11	24/10/1976	01/01/1977
3	11	22/08/1980	16/01/1981
4	494	11/04/1984	02/01/1990
5	41	22/05/1994	07/01/1995

Table 2.3: Parameters with KS Test Method Boundaries

		Phase 1	Phase 2	Phase 3	Phase 4	Phase 5
Model 0	μ	0.3020	0.2541	0.2250	0.3028	0.0957
	AIC	3359.6	1784.4	3834.3	5190.5	1079.7
Model 1	μ	0.0543	0.1721	0.0800	0.0597	0.0270
	k_0	0.8241	0.3233	0.6529	0.8040	0.7231
	ω	0.0542	0.7685	0.0426	0.0316	0.0901
	KS Test	0.0686	0.0528*	0.0465*	0.0343*	0.0455*
	KS Critical 95%	0.0492	0.0701	0.0490	0.0396	0.1072
	KS Critical 99%	0.0590	-	-	-	-
	AIC	3083.7	1717.6	3750.5	5004	987

* Significant at 95% level.

2.5.4 Change Point Analysis Independent of Sociological Boundaries

The previous section described a method which uses sociologically defined boundaries as a starting point to find mathematical boundaries working backwards between pairs of phases. A different approach to finding the mathematical change points for the IED

database is to consider the problem without using prior knowledge of the sociological bounds. Previously the change between phases 1 and 2 was found starting at the sociologically defined boundary for phase 2 and working backwards through phase 1. However, this condition may conceal the fact that the best boundary mathematically between the two phases occurs after the sociologically determined start of phase 2. In this case an algorithm is required which looks at all possible positions for the change point and returns the optimal value.

One possible method is to use a brute force approach whereby a test statistic is employed to examine the set of all possible change point combinations and determine which set is optimal (Ross, 2015). This approach can be undertaken using the AIC. After a specific change point is identified two datasets are created corresponding the data either side of the point. The Hawkes process model is then fitted to each of these datasets and the AIC for each model is computed leading to two values AIC1 and AIC2. The total AIC can then be defined as the sum $AIC1+AIC2$. The total AIC is then the test statistic through which different change points are compared. Since lower values of the AIC imply a better model fit, after checking each change point in this way the result yielding the lowest total AIC is declared as the best change point. Although this is an effective method for small datasets, with larger datasets it is likely that many possible combinations exist making the brute force method computationally impossible to use.

To overcome lengthy computation times in this thesis an optimisation algorithm was written in Python which searches the space of possible change points more effectively. This approach can be understood as having two iteration loops. Firstly

the routine runs an outer loop. In this case the “differential evolution” minimisation routine of the Python package Scipy looks to minimise the total AIC. Since the total AIC is the sum $AIC1+AIC2$ it is also necessary to compute the values of each of $AIC1$ and $AIC2$. Therefore, the current change point being tested by the outer loop is passed to the inner loop. In the inner loop the dataset is first divided into two separate sets corresponding to the data on either side of the change point. Again using the “differential evolution” package the Hawkes process model is fitted to each dataset. The inner loop then ends by calculating $AIC1$ and $AIC2$ for each of the Hawkes process models and returning the value of $AIC1+AIC2$ to the outer loop. Finding the lowest total AIC then occurs by repeating this process with the algorithm then returning the change point yielding the lowest value of $AIC1+AIC2$.

The method of optimising the total AIC was applied to the dataset consisting of all data from phases 1 and 2 combined. Applying the AIC method only to two phases was to allow for an analysis of the consistency of the results of the algorithm by running it 50 times. No other pair of phases could be considered since the total AIC would also depend on any previous phases due to the past dependence of the Hawkes process. The results of these code executions are displayed graphically in Figure 2.5. Graphs in Figure 2.5 are Gaussian kernel density estimates (KDE) of the distributions of the parameter and AIC values for the two Hawkes process models obtained after each code run. The Gaussian KDE is implemented in Python’s Scipy.Stats package (*Scipy.Stats* 2016) and essentially fits one or more Gaussian bell shaped curves to the data.

In general the graphs in Figure 2.5 suggest that the parameter values occur in

two sets. Specifically, considering Figure 2.5(d) each of the AIC distributions have two distinct peaks. However, it should also be observed that in Figure 2.5(c) some of the ω values for the model in phase 2 are greater than one. However, since ω^{-1} is the average time window over which a series of attacks occurs and the data is recorded on a daily temporal scale we require $\omega \leq 1$. Removing the cases where $\omega > 1$ and plotting the distributions again leads to the graphs in Figure 2.6.

The graphs in Figure 2.6 reveal a much narrower distribution for the parameter values and only one main peak for the AIC values. The action of restricting the results to the case where $\omega \leq 1$ is further justified by the fact that from all the code runs the model with the lowest total AIC satisfied this condition. The change point corresponding to the models with the lowest total AIC is given in Table 2.4 and the corresponding model parameters in Table 2.5. From Table 2.5 it can be seen that the models found also have significant fits under the KS test at the 99% level. Moreover, the Hawkes process models in both phases significantly outperform a simple Poisson process model under the AIC comparison test. Also, in this case the model in phase 1 has a significant fit whereas the phase 1 model under the KS test method of Section 2.5.3 did not. This provides further evidence for why it is advantageous to consider the change point problem from a different perspective.

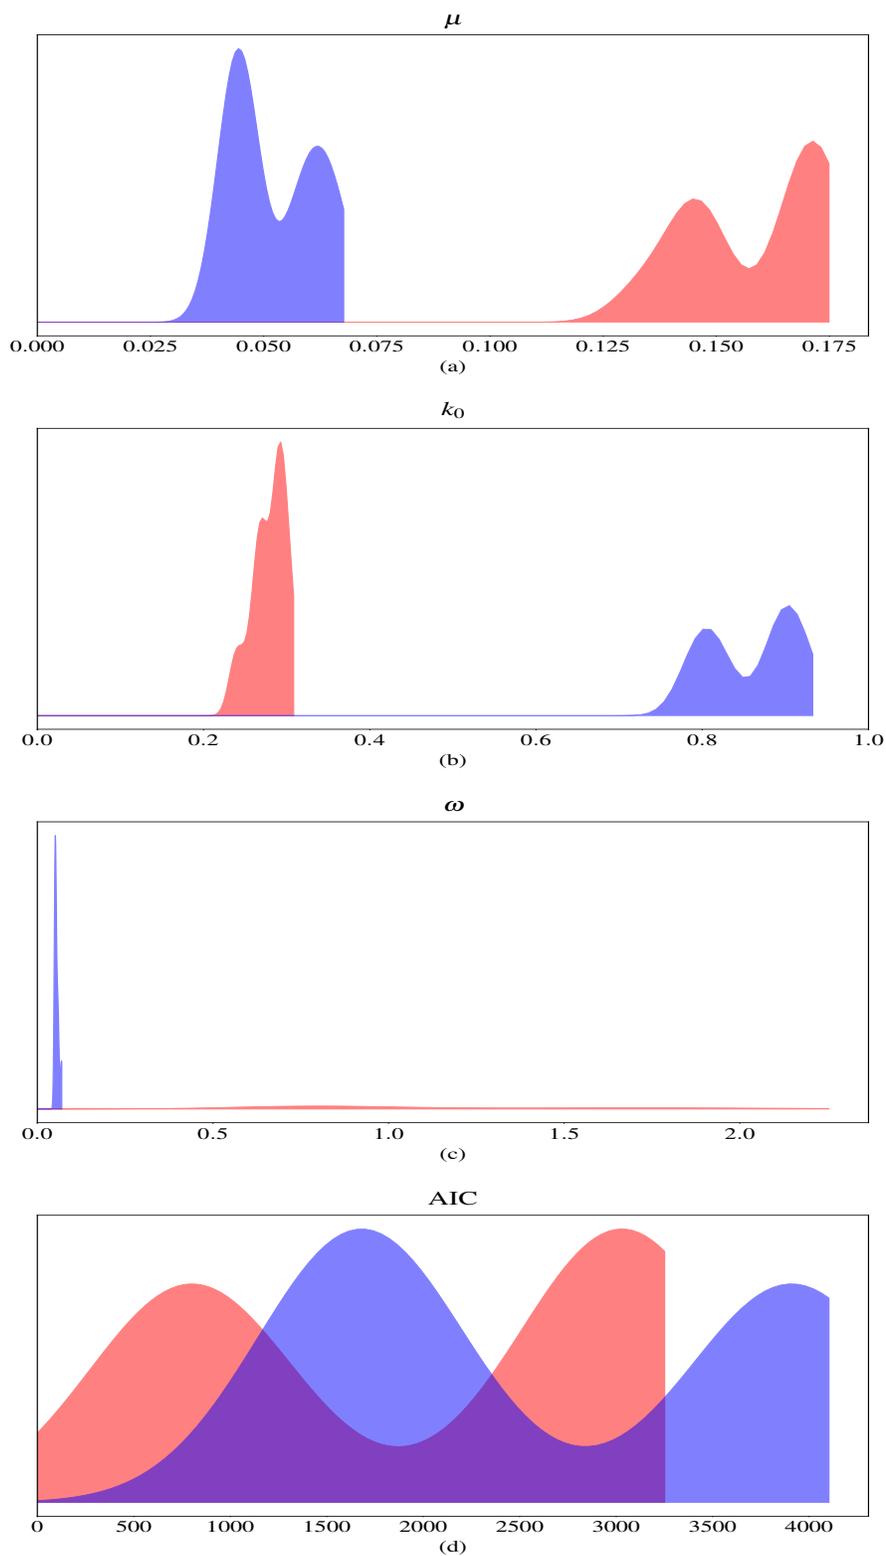


Figure 2.5: Graphs (a)-(c) show the distribution of the parameter values for the Hawkes process in phase 1 (blue) and phase 2 (red) after each attempt at minimising the total AIC. Graph (d) is the corresponding distribution of the AIC values for each of these Hawkes process models.

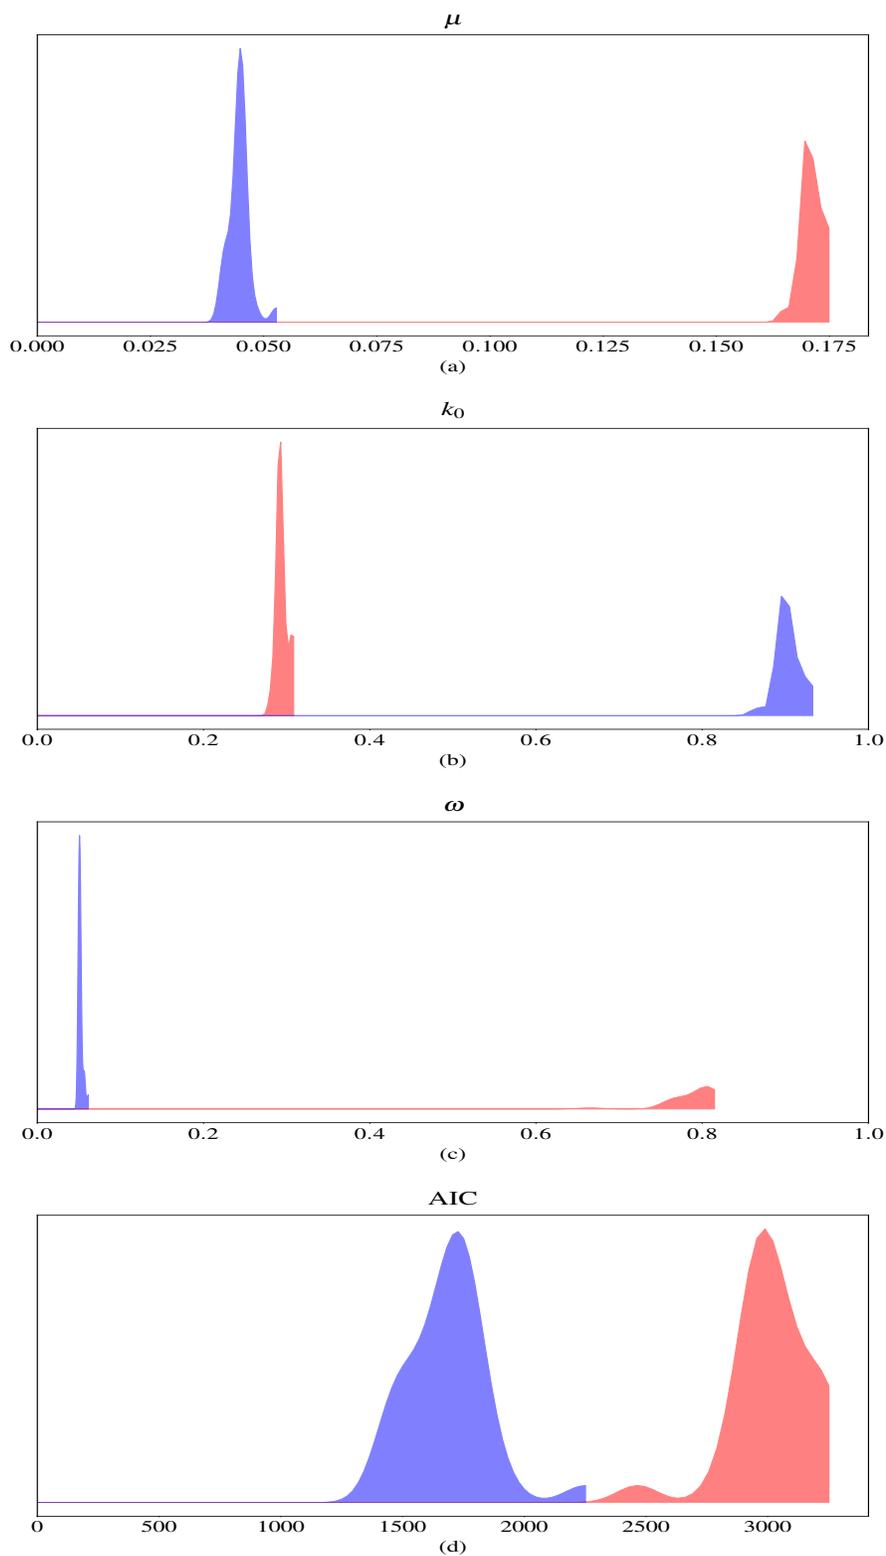


Figure 2.6: Graphs (a)-(c) show the distribution of the parameter values for the Hawkes process in phase 1 (blue) and phase 2 (red) after each attempt at minimising the total AIC and restricting the results to the case $\omega \leq 1$. Graph (d) is the corresponding distribution of the AIC values for each of these Hawkes process models.

Table 2.4: Change Point Analysis Results with AIC Test Method Boundaries

Phase	Number of Data Points Added	New Boundary	Original Boundary
1	-287	27/01/1970	27/01/1970
2	287	23/06/1973	01/01/1977

Table 2.5: Parameters with AIC Test Method Boundaries

		Phase 1	Phase 2
Model 0	μ	0.3837	0.2426
	AIC	1869.7	3153.1
Model 1	μ	0.0439	0.1717
	k_0	0.9074	0.2924
	ω	0.0497	0.8069
	KS Test	0.0724**	0.0551**
	KS Critical 95%	0.0623	0.0533
	KS Critical 99%	0.0746	0.0638
	AIC	1657.9	1717.6

** Significant at 99% level.

For the AIC method the validity of the minimisation routine needs to be checked. One of the ways that the models obtained can be checked is to use bootstrap simulations. The aim of this technique is to study the uncertainty associated with the model parameter values. To achieve this the parameter values are used to simulate event times from a Hawkes process. MLE is then applied to the simulated times to obtain their corresponding Hawkes process parameter values. Finally the distribution of the simulated parameter values are analysed. The bootstrap methodology can be summarised as follows (Embrechts, Liniger and Lin, 2011)

-
1. Simulate the Hawkes process 1000 times over the interval $[0, T]$ where T is the time of the original final observation. Each run of the simulation is required to produce a set of event times $\{t_1, \dots, t_N\}$ where N is the length of the original dataset.
 2. Compute the MLE parameters for each simulated dataset.
 3. Analyse the statistical properties of the resulting parameters.

A common simulation technique used for Hawkes process models is Ogata's Thinning Algorithm (Lewis and Shedler, 1979; Ogata, 1981). The basic idea behind this algorithm is to simulate a homogeneous Poisson process and decide whether or not to keep a point based on the Hawkes process intensity function λ . The algorithm below is based off of the work presented in Ogata (1981). To make the algorithm realistic the same number of points as in the original dataset are generated.

The final result of this algorithm is a set of simulated event times $\{t_1, \dots, t_N\}$ with N being the length of the original dataset.

Statistical analyses of the results of applying the bootstrap method to the models in phases 1 and 2 are provided in Tables 2.6 and 2.7. In Figure 2.7 the distributions of the bootstrap parameters are provided.

These observations provide strong evidence that the Hawkes processes found for each phase are well-defined. From the descriptive statistics we see that the means of the simulated parameter values are close to the actual parameters and the associated standard deviations are relatively small. The ranges of the simulated parameters suggest some extreme values were found but, as can be seen from Figure 2.7, these

```

Set  $\lambda^* = \mu$  and  $s = 0$ .
Generate a random number  $U$  from a uniform( $[0, 1]$ ) distribution and set  $u = -\log\left(\frac{1}{\lambda^*}\right)U$ .
if  $u \leq T$  then
  Set  $t_1 = u$ .
else
  Stop.
Set  $n = 1$  and  $\lambda^* = \lambda(t_n|t_1, \dots, t_{n-1})$ .
while  $n < N - 1$  do
  Generate a random number  $U$  from a uniform( $[0, 1]$ ) distribution and set  $u = -\log\left(\frac{1}{\lambda^*}\right)U$ .
  Set  $s = s + u$ .
  if  $s > T$  then
    Stop.
  else
    Generate a random number  $U$  from a uniform( $[0, 1]$ ) distribution.
    if  $U \leq \frac{\lambda(s|t_1, \dots, t_{n-1})}{\lambda^*}$  then
      Set  $n = n + 1$ ,  $t_n = s$  and  $\lambda^* = \lambda(t_n|t_1, \dots, t_{n-1}) + k_0$ 
    else
      Set  $\lambda^* = \lambda(s|t_1, \dots, t_{n-1})$ .

```

extrema are rare. The graphs in Figure 2.7 also demonstrate graphically that the simulation results are densely clustered around the original MLE parameters. This suggests a low level of uncertainty regarding the parameter values of the original Hawkes process models.

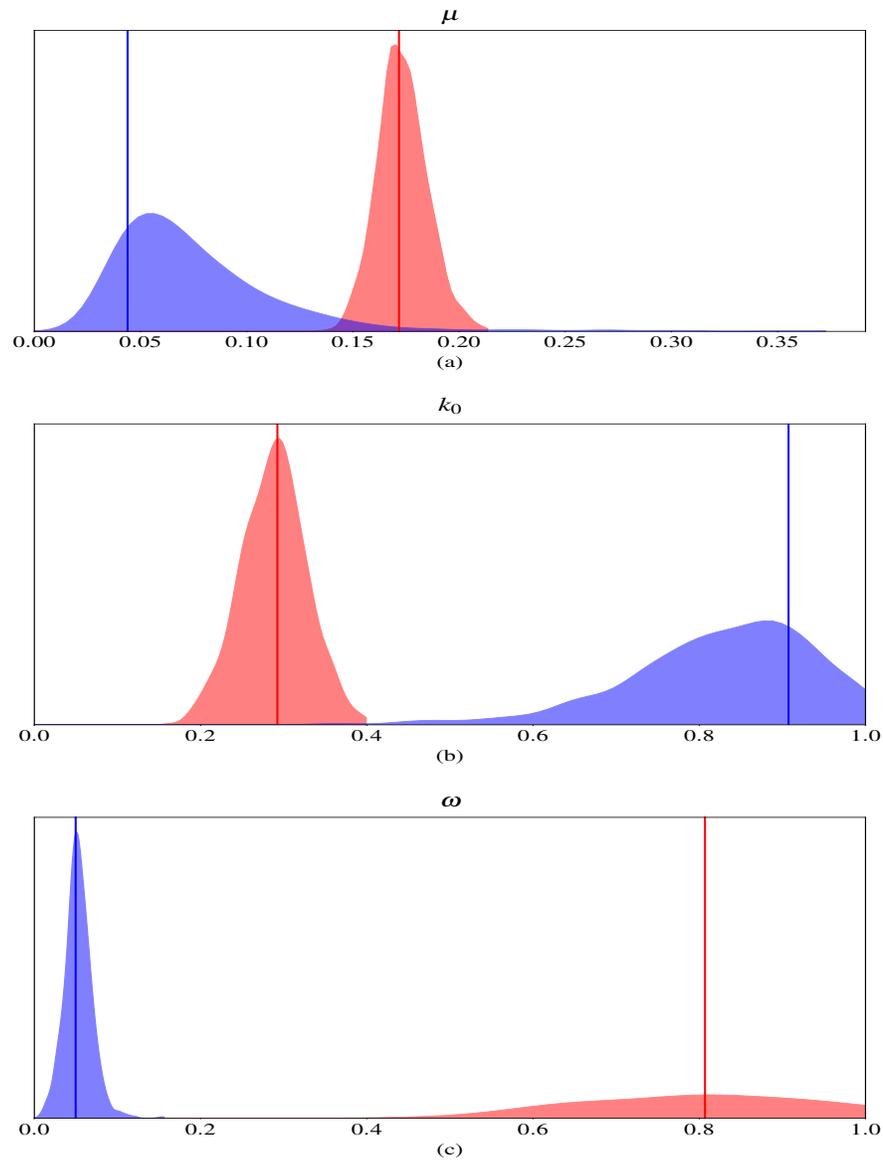


Figure 2.7: Graphs (a)-(c) show the distribution of the simulated parameter values for the Hawkes process in phase 1 (blue) and phase 2 (red). The parameters have been simulated from the Hawkes process models in phases 1 and 2 as defined by the AIC method for determining boundaries. The vertical lines represent the original model parameters.

Table 2.6: Phase 1 Bootstrap Statistics

Parameter	Mean	Standard Deviation	Range
μ	0.0757	0.0013	0.3558
k_0	0.8272	0.0035	0.6493
ω	0.0531	0.0167	0.1511

Table 2.7: Phase 2 Bootstrap Statistics

Parameter	Mean	Standard Deviation	Range
μ	0.1734	0.0116	0.0756
k_0	0.2882	0.0391	0.2353
ω	0.8393	0.1862	1.5967

2.6 Interpreting Phase Boundaries

In Sections 2.5.3 and 2.5.4 two methods to determine change points in temporal series of terrorist attacks have been constructed and applied to PIRA IED data. The results of these analyses have indicated that there exists a difference between the phase changes of PIRA found qualitatively in social science research and the boundaries found using quantitative methods. In particular, Figure 2.8 demonstrates the change in the distribution of events between the phases for the original, KS and AIC boundaries. The discrepancies between the qualitative and quantitative boundaries

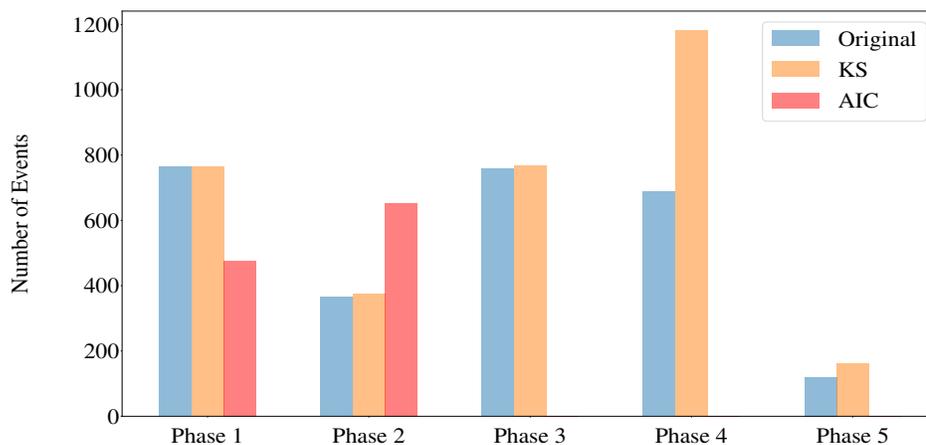


Figure 2.8: Distribution of IED events between phases under the original phase boundaries (blue), KS method boundaries (orange) and AIC boundary (red).

warrants further consideration to understand the origins of these differences.

For the KS method of determining phase boundaries in Section 2.5.3 all of the phases of PIRA were analysed. It was discovered that phases 2 and 3 only required 11 points from their prior phases to yield the best model fits. In addition, these extra points only moved the temporal boundaries by 2 months for phase 2 and 4 months for phase 3. From these observations we can conclude that the social science boundaries observed for these two phases are approximately consistent with the quantitative findings. Furthermore, since it is unreasonable to expect that an organisation like PIRA can shift its entire operations and organisational structure instantly the delays found in Section 2.5.3 seem reasonable.

The main differences between the research findings of Section 2.5.3 and the social science phase boundaries occur in phases 4 and 5. In phase 4 an additional 494 points were required to find the best fitting model. These extra points were equivalent to

shifting the start of phase 4 backwards by almost 6 years. During the literature review at the beginning of this thesis an overview of the evolution of PIRA was provided in Section 1.5.1. Here it was described that phase 4 marked a shift in PIRA towards political engagement with the British Government via secret channels of communication. Phase 4 was also preceded by phase 3 which Moloney (2003) describes as one of escalating violence similar to the “Tet Offensive” launched by the People’s Army of Vietnam in 1968. As stated in the observations about the change in phase boundaries for phases 2 and 3 organisational changes are unlikely to occur instantly. Moreover, since the negotiations in phase 4 were secret the wider PIRA cadre would likely not have known of their existence. Therefore, the large changes observed in the phase 4 boundaries indicate that we should consider phase changes as fluid. In particular, phases 3 and 4 could have been occurring in parallel and delineating them via a sequential boundary is inaccurate.

A similar observation for the phase boundary in phase 4 explains the result found for the boundary between phases 4 and 5. Specifically, phase 5 required 41 data entries from phase 4 corresponding to a shift in temporal boundary of close to 1 year. In Section 1.5.1 it was stated that the secret peace talks between PIRA and the British Government were announced and ratified in the Good Friday Agreement in phase 5. As before here we observe the fluid nature of temporal boundaries and their ability to occur in parallel. As noted by Coogan (2002) PIRA had the ability to “turn ... bombing[s] on and off like a tap”. Therefore, as ceasefire terms came to conclusion PIRA may have selectively used IED attacks to negotiate better terms. In addition, some PIRA members may have been reluctant to accept an end

to violence given that recruits were instructed to wage a “war of attrition against enemy personnel which is aimed at causing as many casualties and deaths as possible so as to create a demand from their people at home for their withdrawal” (O’Brien, 1999).

For the second approach to finding phase changes in Section 2.5.4 an AIC method was introduced. For this method only the phase boundary between phases 1 and 2 was considered due to computational restraints. From the AIC method it was found that 287 points needed to be added from phase 1 to phase 2. This result corresponded to a shift backwards in the temporal boundary, from the social science boundary, by 3.5 years. In contrast to the KS method, when using the AIC method a tradeoff is made to decide the best boundary for both phases simultaneously. As stated in the overview of PIRA in Section 1.5.1 phase 1 was characterised by a military structure. However, this organisational approach was susceptible to infiltration by British security forces. Therefore, in phase 2 PIRA split into a cell-based structure consisting of small groups known as Active Service Units. This new structure aided the PIRA in keeping its operations secret and protecting its member’s identities. The results of the AIC phase change analysis suggest that PIRA members may have been making a shift to increased secrecy at an earlier time. In particular, changing the entire structure of an organisation like PIRA would not be instantaneous. Moreover, PIRA may have wanted to move to a new structure slowly to avoid detection of the shift by British security forces. Therefore, it can be observed that a similar observation is made to the findings found for the KS test method. Specifically, terrorist groups evolve over time in response to external stimuli. Hence, a shift in

understanding phase changes is required to consider them as gradual processes which can overlap and occur simultaneously.

2.7 Interpreting Model Parameters

Mathematical modelling in general is motivated by the opportunities it provides to yield insights about real-world dynamics. One of the ways insight can be obtained for the Hawkes process is through studying the MLE parameters of the model. In particular, the intensity function of the Hawkes process is specified via three parameters $\{\mu, k_0, \omega\}$. Each of the parameter values obtained via MLE can be interpreted in terms of the events being modelled. In the context of this thesis it is of interest to study what the parameters imply about the timings of IED attacks by PIRA and how these timings relate to the tactical and structural evolution of the organisation. First the parameters obtained using the KS method for boundary determination will be discussed. These parameters will then be compared to the ones obtained via the AIC method for boundary determination.

2.7.1 KS Method Parameters

To interpret the meaning of the value of ω it is first useful to recall that the inverse ω^{-1} indicates the average length of time over which a series of attacks decays (Lewis et al., 2012). In phase 1 the inverse of ω satisfies $\frac{1}{0.0542} = 18.5$ days. Following this in phase 2 the value of ω increases producing a smaller average attack window of 1.3 days. This change in the average attack window can be related back to the

sociological and historical observations of the PIRA. In particular, it is discussed in the work of Asal et al. (2013) that in phase 1 of the PIRA the organisation adopted a militaristic structure. However, this particular organisational structure made the PIRA susceptible to infiltration by British security forces. Therefore, in phase 2 the PIRA re-organised their structure instead opting for a cellular strategy made up of Active Service Units (ASU) (Horgan and Taylor, 1997). This cell based approach saw the PIRA break into small groups meaning any potential informers could be quickly identified. The shift in structure was successful at limiting the security forces abilities to place informers within the PIRA with Smith (1997) reporting there were 465 fewer charges for paramilitary activity within a year. By comparing these two different organisational structures the dynamics which led to the observed changes in the value of ω can be explained. In particular, with a military hierarchy in phase 1 it may have been easier for the PIRA to carry out a sustained wave of attacks. Whilst in phase 2 with a more fragmented organisation sustained attacks may have been harder to achieve.

In phase 3 the average length of attacks becomes 23.5 days. This period is described by Moloney (2003) as one of escalating violence similar to the “Tet Offensive” launched by the People’s Army of Vietnam in 1968. This is reflected in the PIRA issued Green Book discussing the organisations aims and objectives. PIRA wanted a “bombing campaign aimed at making the enemy’s financial interests in our country unprofitable” and to wage a “war of attrition against enemy personnel which is aimed at causing as many casualties and deaths as possible so as to create a demand from their people at home for their withdrawal” (O’Brien, 1999). These objectives

are being reflected in the longer length of an average attack window.

For the final two phases of the organisation the trend is for ω to increase resulting in shorter lengths of time for waves of attacks. For phase 4 the average attack window is 31.6 days and this falls in phase 5 to 11.1 days. During the final two phases the leadership of PIRA was conducting secret negotiations with the British Government to end hostilities (Coogan, 2002; English, 2004; Moloney, 2003). However, during this time the PIRA wanted to demonstrate their attack capability in order to strengthen their negotiating position. As noted by Coogan (2002) PIRA had the ability to “turn ... bombing[s] on and off like a tap”. The PIRA’s flexibility in IED attacks is represented in the shortening of the average time window as the focus of bombings was as a bargaining tool rather than for a war of attrition.

Similar to the values of ω the parameter μ can also be related to actual observations about the PIRA. The parameter μ represents an underlying background rate at which new events occur. Moving from phase 1 to phase 2 the value of μ increases. Explained previously this change of phase saw the PIRA move from having a militaristic structure to a cellular based structure consisting of ASUs. Therefore, with less control within the fragmented organisation the ASUs may have been carrying out attacks more randomly resulting in an increase in the value of μ .

The final three phases show a declining trend in the values of μ . In phase 3 it has been discussed that at a strategic level PIRA moved towards a “Tet Offensive” campaign characterised by a renewed emphasis on bombings. Therefore, it is possible that to make this approach successful attacks were becoming more systematic as well as occurring in succession. Thus there would have been less randomness in the IED

attacks explaining the decline from phase 2 to phase 3 in the value of μ .

For phases 4 and 5 the leadership of the organisation was focused on negotiations with the British Government to end the conflict in Northern Ireland. Hence these two periods had an overall de-escalation of violence. This is precisely the trend found in the values of the parameter μ where the background rate of events decreases.

The values of the final parameter k_0 across the five phases of the PIRA also demonstrates trends related to actual events. This parameter describes the magnitude of increase in the intensity of events following an initial incident. The highest value of k_0 is found in phase 1. Since this was the military phase of PIRA in phase 1 coordinating attacks would have been easier for the organisation. Therefore, generating a sequence of IED attacks would have been more likely in this phase. In phase 2 however the cellular structure emerged. Hence coordinating multiple attacks may have been hindered which is demonstrated by a decline in the value of k_0 . In the next phase the “Tet Offensive” strategy was used and hence ASUs could have been instructed to launch more IED attacks. Thus a rise in the value of k_0 can be observed in phase 3. Finally in phases 4 and 5 the main focus of events was negotiations between PIRA and the British Government. During these phases the aim of IED attacks was to strengthen the position of PIRA at the negotiating table. The high values of k_0 in phases 4 and 5 could be representing PIRA using IEDs as a bargaining tool as described above.

2.7.2 AIC Method Parameters

To reduce computational time it was only possible to consider the AIC method for change point detection between phases 1 and 2. However, the resulting models, in Table 2.5, can still be compared to those from the KS method, in Table 2.3. For both phases 1 and 2 the magnitudes of the parameter values are the same for both methods. In particular, each of the parameter values resulting from the AIC method are close to those from the KS method. Moreover, the trends in the parameter values are the same. A similar qualitative assessment as was conducted for the KS method results holds for the AIC method results. This shows that the AIC method for determining change points also produces real-world interpretable parameters.

Also of interest is a comparison between where the division between phases occurs when trying to minimise the total AIC and that found earlier from the KS test method. Initially it can be seen that many more points are added to phase 2 for the AIC method than for the KS method. However, as just discussed, the parameters for the models obtained under the two different approaches are very similar. This observation suggests that the evolution of the PIRA was quite gradual since the underlying dynamics appear to be stable over a long period of time.

It is also of interest to observe that for both the KS method and AIC method the boundaries are different from those suggested by the sociological theory. This raises interesting research questions which can be pursued considering the possibility of some new sociological theory concerning the evolution of the PIRA.

2.8 Discussion

The chapter opened by introducing the Hawkes self-exciting point process model. To gain an initial understanding of how the model worked the mathematical background of the model was detailed. This background covered the basic components of the Hawkes model including the intensity function, maximum likelihood parameter estimation and a discussion of the computational nuances associated with the optimisation problem. The intensity function serves as the essential interface between the mathematical theory behind the point process and the real-world applications for which it can be used. In particular, as part of the mathematical foundation the specific form of the intensity function was provided and shown to rely on three parameters. The first of these three parameters was a background rate μ describing an average rate of event occurrence. Next a jump parameter, k_0 , is used to describe how following an incident there is a spike in the event rate occurrence due to self-excitations. Finally, since it is unrealistic that self-excitations last indefinitely a decaying kernel in the form of an exponential, with exponent ω , is used to dampen the increased event rate.

With the physical interpretation of the parameters explained the next step in this chapter was to introduce a case study in the form of the PIRA. As was already discussed earlier in the literature review taking a historical and sociological perspective on the PIRA's activities indicated five unique phases relating to shifts in the organisations structural and tactical modus operandi. The task undertaken in this chapter was to compare the mathematical boundaries for these phases to the five qualitative

descriptions. The mathematical approach to boundary detection followed two routes - a KS method and an AIC method.

The KS method designed used the idea of maximising the KS difference between the KS test statistic and critical value between pairs of phases. Each adjacent phase pair was scanned with a moving time window adding one data point from the previous to the next phase at each shift. With each new dataset thus obtained a Hawkes process was fitted to the time series and a corresponding KS test statistic and critical value were calculated. Finally by graphing all these differences it was possible to observe where the optimum change point should be placed. This method proved to be quick and straightforward to implement so that insights could be readily obtained.

A different method that was developed to handle the change point analysis problem was to study the total AIC between two phases. The aim of this technique is to find the break point which produces the lowest possible AIC for each phase combined without relying on the sociologically defined boundary. A problem that was discovered was the computational inefficiency of this approach. To overcome long code runtimes a strategy was developed involving a two loop procedure. In the outer loop a minimisation algorithm was being applied to the total AIC corresponding to the break point inbetween the two phases. Since this total AIC was a function of the sum of the two individual AIC values of the Hawkes process in each phase it was then necessary to have a second minimising loop. In this inner loop the MLE values for the Hawkes processes in each phase were computed in order to define the intensity functions. With the intensity functions defined the AIC of the Hawkes processes in each phase were found and the sum returned to the outer loop. This approach to

the change point provided the advantage of not having to rely on the sociological boundaries thus enabling the best change point among all possible change points to be detected.

Tying this chapter together the final section aimed to interpret the different mathematical findings from the KS and AIC approaches to change point detection in the context of the PIRA and “The Troubles” in NI. It was found that the Hawkes process MLE parameters from each technique were close in value with similar trends. This observation provided confidence that both methods were providing valid results. This confidence was further reinforced by converting the quantitative results into a qualitative description of the PIRA. By comparing the real-world interpretations of the MLE parameter estimates to historical accounts and sociological theory of the PIRA it was found that the Hawkes process models were indeed picking out important shifts within the organisation. Moreover, the differences between the phase points obtained here and those previously used by social scientists indicate that the Hawkes process is also offering new insights opening up opportunities for further research into understanding the origin of these differences. Such additional research may prove fruitful in uncovering as yet undetected subtleties concerning the PIRA and the political climate in the time it operated.

The results of this chapter provide evidence for Hypotheses 1-2. In particular, developing unique change point detection methods using the Hawkes process model has shown PIRA attacks are historically dependent. Moreover, via the interpretation of the parameters of the Hawkes models it has been shown that the phases of the PIRA were quantitatively distinct. Both of these outcomes illustrate that the research

conducted in this chapter provides useful advancements to knowledge for studying terrorism. Also, employing the change point detection frameworks that have been described important insights into the evolution of terrorist groups can be identified. Specifically, academic researchers in the field of social science can benefit from these types of methods to obtain deeper understandings of when and how terrorist groups change structurally and operationally. Furthermore, practitioners can use the models developed in this chapter to study terrorism in real time and quickly identify important shifts. Therefore, the research in this chapter also provides evidence for the validity of Hypothesis 4.

To extend the research presented in this chapter further use of spatial data could be made inside of the Hawkes process. In particular, it may prove fruitful to study change point detection dependent on both time and space. A possible method to achieve this objective would be to include a spatio-temporal kernel inside the Hawkes process model. This type of Hawkes model could then be used inside the change point detection framework developed and analysed in this chapter. By explicitly studying spatial components of PIRA attacks new insights may be revealed pertaining to their use of geography during “The Troubles”. In addition to deeper academic understandings of PIRA having models including spatial patterns could assist in identifying hotspots of terrorist activity. Moreover, improved understanding of spatial change points would help guide practitioners in developing evidence based counter-terrorism approaches.

CHAPTER 3

SEASONAL HAWKES PROCESS MODELLING OF AL SHABAAB

In 1979 Cohen and Felson published their paper describing research formulating the routine activities of crime (Cohen and Felson, 1979). According to their work they claimed crime occurs as the result of opportunities governed by environmental conditions. Through analysing the relationship between different environmental factors and criminal events it is possible that many new crime reduction techniques can be discovered.

Research since has focused on the influence of environmental factors on crime, particularly, the effect of weather. Multiple articles have been published exploring this link. For example, the research of Cohn (1990) concludes that assaults, burglary, collective violence, domestic violence and rape are positively correlated with temperature with a breakdown in these relationships occurring around 29°C. Similar results have been reported by Pakiam and Lim (1984) whose study of crime and weather in Singapore found that crimes against the person increase with increasing temperature and climate comfort measures. Extrapolating the effect of weather conditions on crime Ranson (2014) estimates rises in multiple crime categories as a result of climate changes.

It has already been shown in several studies that a theory akin to Cohen and Felson's routine activity approach holds true for terrorist attacks (Clauset and Woodard, 2013; Townsley, Johnson and Ratcliffe, 2008; Zammit-Mangion et al., 2012). Similarly, some studies have emerged illustrating repeating trends in the frequency of terrorist activity. In the research of Enders, Parise and Sandler (1992) it has been shown that terrorist events, such as, bombings, hostage events and assassinations have cyclical patterns ranging from 21 to 54 months. Moreover, threats of attacks

were found to have a seasonal trend of 11 months which the authors associate to the onset of tourism periods. Also finding cyclical trends in terrorist attacks Weimann and Brosius (1988) finds a contagion dynamic as well as a constant one month periodicity modelled with a first-order moving average.

However, a gap in terrorism literature remains to utilise advanced mathematical techniques to capture seasonal variations in terrorist attacks. In particular, it has been discussed that terrorist events have both cyclical and contagion dynamics. But there exists a disconnect between the models used to capture each of these components. By developing a more encompassing modelling framework which can simultaneously describe the different influences on terrorist activity it may be possible to discover more effective counter-terrorism policies. In this chapter a possible solution to this problem will be presented in the form of a Hawkes point process model with a seasonal background rate. To illustrate the usage of this type of model a case study of Al Shabaab (AS) terrorist incidents in Somalia will be examined.

3.1 Datasets

3.1.1 Al Shabaab Event Data

The data used in this chapter is sourced from the Armed Conflict Location and Event Database (*ACLED* 2016). *ACLED* is a project run from the University of Sussex and the database aims to provide up-to-date disaggregated data for conflicts in Africa as well as South and South-East Asia. The data is sorted according to date (on a daily scale), location, type of event (battles, civilian killings, riots, protests

and recruitment activity), event actors (rebels, governments, militias, armed groups, protestors and civilians), information on changes in territory control and fatalities. The database is formed using information from a variety of sources including local and international news sources and humanitarian agency reports. An outline of the two main fields used in this chapter is provided in Table 3.1

Table 3.1: ACLED Data Fields

Field	Data
EVENT_DATE	Day of the recorded event
EVENT_TYPE	Battle-{"Government regains territory", "No change of territory", "Non-state actor overtakes territory"}, Headquarters or base established, Non-violent transfer of territory, Remote violence, Riots/Protests, Strategic development, Violence against civilians

In the research of Maszka (2017) the strategic developments of AS are examined. Specifically four distinct phases of the group are identified. These phases are from the group's inception to December 2007, January 2008 to April 2008, May 2008 to July 2011 and August 2011 to the present time. Within these transitions the group also had three leadership changes according to the following timeline

- **2006-2008:** Aden Hashi Ayro (killed by a US airstrike).
- **2008-2014:** Ahmed Abdi Godane (killed by a US drone strike).
- **2014 - Present:** Abu Ubeyda.

Multiple historical and social science studies have been conducted into the evolution of the AS group (Anderson and McKnight, 2015a; Hansen, 2013; Marchal,

2009; Maszka, 2017; Wise, 2011). During the early stages in the formation of AS the group's main objective was resisting a perceived occupation of Somalia by Ethiopian troops. After the withdrawal of Ethiopian forces in 2009 AS struggled to retain local support and began an international recruitment campaign. This shift from a local struggle to a global movement was also accompanied by Godane's desire to link AS to the international terrorist group known as al Qaeda. In 2010 the African Union Mission to Somalia (AMISOM) consisting of Ugandan and Burundian troops moved to counter the growing threat presented by AS. However, AS persisted and in 2011 Kenyan troops crossed into South Somalia to provide a bulwark against the risks of AS attacking inside the Kenyan border. Retribution for these moves to contain and eliminate AS was seen in high profile attacks, such as, an attack on the Westgate Mall in Nairobi in 2013 (Williams, 2014) and another on the Garissa University in 2015 (Lyons et al., 2015).

Examining the historical roots of AS indicate that the group was heavily focused on fighting various governmental forces. However, as Maszka (2017) discusses under the change of leadership to Godane the AS group underwent a tactical shift. In particular, the usage of improvised explosive devices (IEDs) became more widespread. Alongside the use of IEDs Godane's appeal for foreign fighters also lead to the employment of suicide attacks in Somalia. These developments significantly increased events involving violence against civilians.

Since attacks against governmental forces and those against civilians were separate developments within the history of AS it was decided to model these two categories of events separately. Thus the ACLED data was divided according to the

following criteria

- Battle Territory (BT) - Covers pre-existing categories “Battle-Government regains territory”, “Battle-No change of territory”, “Battle-Non-state actor overtakes territory”.
- Violence Against Civilians (VAC) - Covers pre-existing category “Violence against civilians”.

In total the ACLED database, at the time of production of this chapter in 2016, contains a total of 5935 events involving the Al Shabaab militant group. The BT category of events includes 4089 recorded incidents and the VAC database has 670 events. Therefore, these two databases combined cover 80% of all the available Al Shabaab data and hence provide good coverage of the group’s activities for numerical study.

It should also be noted from Chapter 2 that one of the assumptions that must be met to use the Hawkes process model is that the timestamps of events must be unique so that $t_i \neq t_j$ for $i \neq j$ (Daley and Vere-Jones, 2003). For the study of Al Shabaab attacks in this chapter the events in the ACLED database are recorded according to day of the event. Therefore, to ensure the uniqueness of event times multiple events on a single day are counted as only one event. Justification for such an approach to analysing the Al Shabaab data comes from studies of the groups command structure which shows it operates in a top down fashion (Agbibo, 2014; Marchal, 2009) implying that multiple same day attacks can be considered as stemming from a single decision.

Due to the uniqueness assumption of the Hawkes model dividing the AS data into BT and VAC categories also ensures more data is retained. In particular, if the data was analysed in aggregate form then important information may have been lost due to the removal of repeated timestamps. However, under the analysis framework presented in this chapter this problem is avoided due to the data disaggregation used.

Before removing duplicate entries the BT category of events has 4089 records. After ensuring unique timestamps the BT dataset has 1751 events. Performing the same procedure for the VAC dataset leads to a reduction from 670 to 547 datapoints. These reduced datasets are used in the remainder of this chapter.

Since the AS group's tactics evolved in four phases a further disaggregation of the data can also be undertaken into the four datasets as illustrated in Tables 3.2 - 3.3. Observe that the dates demarcating the phase boundaries are adjusted based on the actual data. In particular, the first and last dates recorded in the BT and VAC datasets that fall inside each of the phases found by Maszka (2017) are used to define the phase edges. The motivation for this presentation of the phases is to ensure that the analyses conducted in this chapter capture the actual data dynamics. It should also be noted that phases 1 and 2 have very few datapoints. Therefore, the analyses in this chapter will focus only on phases 3 and 4.

Table 3.2: Number of Entries in BT Datasets

Dataset Name	Phase 1 (Aug 2006 - Nov 2007)	Phase 2 (Jan 2008 - Apr 2008)	Phase 3 (May 2008 - Jul 2011)	Phase 4 (Aug 2011 - Dec 2015)
Battle Territory	5	6	407	1333

Combining all of the previous disaggregations yields four datasets - BT (phases

Table 3.3: Number of Entries in VAC Datasets

Dataset Name	Phase 1 (Aug 2007)	Phase 2 (Jan 2008 - Apr 2008)	Phase 3 (Jun 2008 - Jul 2011)	Phase 4 (Aug 2011 - Dec 2015)
Violence Against Civilians	1	3	51	492

3 and 4) and VAC (phases 3 and 4). These datasets were used for the studies in the remainder of this chapter.

3.1.2 Weather Data

As discussed in the introduction to this chapter research has shown that terrorist attacks are often linked with seasonal trends. One of the main seasonal trends in Somalia is the onset of wet seasons. In the report of (Muchiri, 2007) rainfall is described as the “defining characteristic of the climate” in Somalia with two main rainfall seasons. The first rainfall season in Somalia is known as the Gu (April-June) resulting from the northward movement of weather fronts whilst the second is called Deyr (October-December) resulting from a southerly shift. An established monitoring body of weather patterns in Somalia is known as the Somalia Water and Land Information Management (SWALIM) which is managed by the Food and Agricultural Organisation (FAO) of the United Nations (*FAO SWALIM* 2016). All the data collected as part of the SWALIM project is openly accessible via the organisation’s website. Of particular interest in this thesis is rainfall data. The SWALIM database offers rainfall data from a variety of collection stations across the country. Since most of the Al Shabaab attack data in the ACLED database records events in the south of

Somalia (89% of all data points) data for the central southern data collection point in Diinsor in the Bay area was used in this thesis. This area is highlighted in Figure 3.1 (*FAO SWALIM* 2016). The rainfall data downloaded for use in this chapter consists of monthly volume counts spanning the years 2008 to 2015 providing full coverage of the period of AS attacks being studied. In Tables 3.4 - 3.5 an overview is provided of the rainfall datasets used in this chapter. Note that these rainfall datasets have start and end dates corresponding to those used for the BT and VAC datasets. This ensures consistency in the analyses of dynamics in the rainfall and event data.

Table 3.4: Number of Entries in BT Weather Datasets

Dataset Name	Phase 3 (May 2008 - July 2011)	Phase 4 (August 2011 - December 2015)
BT Weather	39	53

Table 3.5: Number of Entries in VAC Weather Datasets

Dataset Name	Phase 3 (June 2008 - July 2011)	Phase 4 (August 2011 - December 2015)
VAC Weather	38	53

3.2 Modelling Framework

The main objective of this chapter is to develop a model that can capture both seasonal trends and past dependent behaviour in terrorist activities. This task will be divided into two steps.

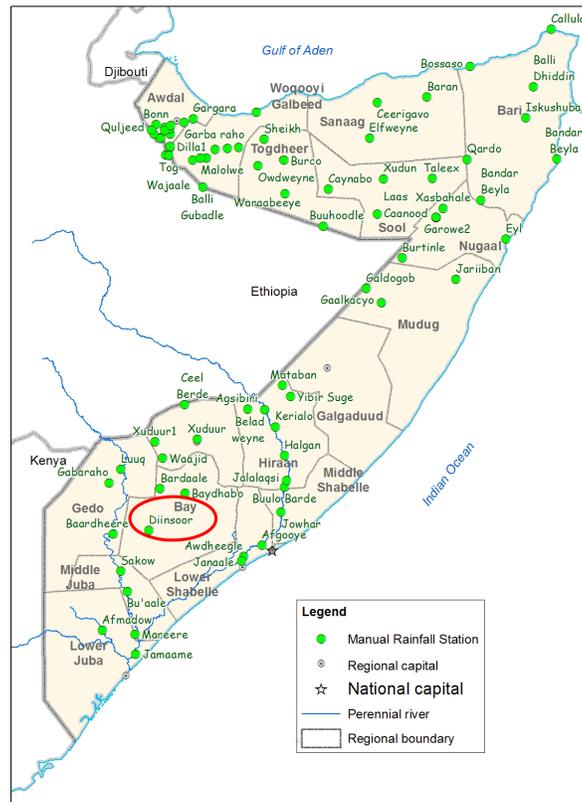


Figure 3.1: Map of Somalia. The area indicated in red indicates the location of the weather station from which the rainfall data used in this chapter was obtained.

Initially, in Section 3.3 the periodicities present in the AS event dataset and Somalia weather dataset will be obtained using Fourier analysis. By comparing these periods it will then be possible to determine if there exists a weather-related dynamic within AS attacks.

Having extracted the underlying periodicities the past dependent behaviour of the AS attacks will be captured using a Hawkes point process model in Section 3.4. Results will be presenting comparing the modelling of AS events using a Hawkes

model with a constant background rate and a Hawkes model using a seasonally varying background rate. The seasonally varying Hawkes process will utilise the periodicities found from the Fourier analysis described previously.

3.3 Identifying Seasonality

To justify the usage of a seasonally varying Hawkes process model it is necessary to demonstrate quantitatively that the AS attacks contain a seasonal dynamic. In particular, the focus of this chapter will be to study the influence of rainfall levels in Somalia on AS events. One method well suited to this task is the construction of a Fourier series. As described by Stade (2005) the aim of this approach is to generate a model using a summation of sinusoidal terms where each term has a distinct frequency. The inverse of these frequencies yields the periodicities of dynamics in the data being modelled. By obtaining a Fourier series model of both the Somalia rainfall datasets and AS event datasets it will be possible to discover their underlying cyclical trends as measured via periodicity. Directly comparing these periodicities then enables conclusions to be drawn concerning the link between weather trends and AS attacks.

Prior to presenting the analysis comparing the rain and event datasets the mathematical formulation of Fourier series will be presented in the following section. Since the rainfall data used is measured according to a monthly scale a comparison will be made to monthly AS data. The AS data is grouped into months via binning of the data with bin widths equal to 30 days.

3.3.1 Detrending Data

An important note concerning the use of Fourier analysis is that the dynamic of interest is data periodicity. Therefore, prior to constructing a Fourier series the data being studied should be detrended so that only cyclic repetition is present. Two common approaches of detrending are

1. subtract the mean of the data from the dataset,
2. remove a linear trend in the data using a least squares regression.

The exact method chosen to detrend a dataset is context dependent. Also it is important to remember that the detrended method has to be incorporated at the end of fitting the Fourier series. This can be achieved either by adding the mean as a constant or adding a linear function to correspond to the one found in the second detrend method.

To illustrate the context of data requiring detrending observe the histogram in Figure 3.2. This figure shows the monthly event counts corresponding to the “Battle for Territory” category in phase 3 of the AS evolution. It is clear from viewing this plot that there exists a linearly increasing trend across the dataset.

One method which can be used to detrend linearly increasing data is known as linear least squares regression. Assume a given set of data points $\{(t_n, f_n)\}_{n=1}^N$. To these datapoints we aim to fit a linear function $f(t) = \alpha t + \beta$ with parameters $p = \{\alpha, \beta\}$. This aim can be achieved by minimising the sum of squared differences

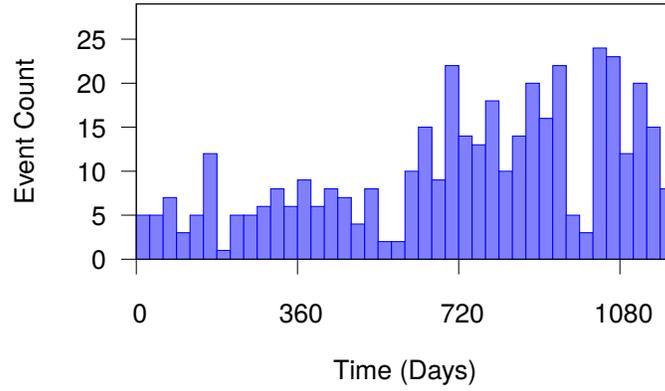


Figure 3.2: Battle territory phase 3 binned data histogram. Bin widths are set to 30 days to correspond to approximately one month of data.

given by the loss function (Hansen, Pereyra and Scherer, 2012)

$$\ell(p) = \sum_{n=1}^N (f_n - f(t_n; p))^2. \quad (3.1)$$

Since the aim is to minimise (3.1) we need to find the partial derivatives with respect to each of the two parameters and set the resulting equations equal to zero. Hence we obtain the set of simultaneous equations

$$\frac{\partial \ell}{\partial \alpha} = -2 \sum_n t_n (f_n - \alpha t_n - \beta) = 0, \quad (3.2)$$

$$\frac{\partial \ell}{\partial \beta} = -2 \sum_n (f_n - \alpha t_n - \beta) = 0. \quad (3.3)$$

The equations in (3.2) and (3.3) can be simplified to

$$\alpha \sum_n t_n^2 + \beta \sum_n t_n = \sum_n t_n f_n, \quad (3.4)$$

$$\alpha \sum_n t_n + \beta N = \sum_n f_n. \quad (3.5)$$

Solving the linear equations in (3.4) and (3.5) yields the following formulae for the parameter values

$$\alpha = \frac{N \sum_n t_n f_n - \sum_n f_n \sum_n t_n}{N \sum_n t_n^2 - (\sum_n t_n)^2}, \quad (3.6)$$

$$\beta = \frac{1}{N} \sum_n f_n - \alpha \frac{1}{N} \sum_n t_n. \quad (3.7)$$

Applying the detrending technique to the BT dataset in phase 3 gives the coefficients $\{\alpha, \beta\}$ in Table 3.6. Figure 3.3 illustrates how the histogram of event counts changes after detrending the data.

Table 3.6: Battle Territory Phase 3 Linear Detrend Function Coefficients

Coefficient	Value
α	0.0114
β	3.5329

3.3.2 Fourier Series

Literature describing the mathematical background and application details of Fourier series is abundant (Davis, 1989; Dyke, 2014; Pinkus and Zafrany, 1997; Stein and Shakarchi, 2003; Strang, 1986; Tolstov, 1977). Here a summary of this literature relevant to determining seasonal trends will be presented.

The basic premise of a Fourier series is to express a periodic function in terms of a summation of sines and cosines. For example, consider a function f with a period

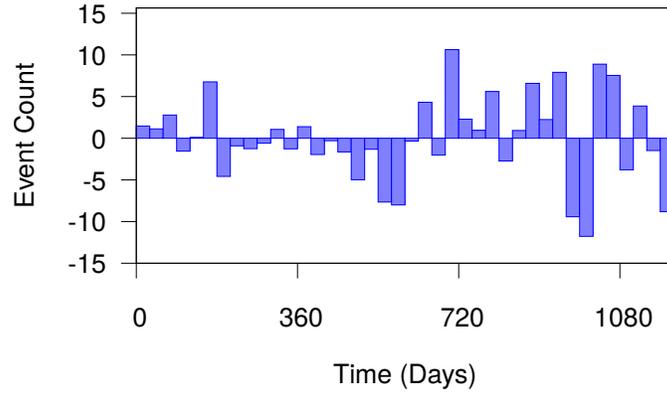


Figure 3.3: Battle territory phase 3 binned data histogram with a linear detrend. Bin widths are set to 30 days to correspond to approximately one month of data.

of $2T$ i.e. assume f satisfies the periodic property that

$$f(t) = f(t + 2T). \quad (3.8)$$

For the periodic function f the full Fourier series on the interval $0 < t < 2T$ is defined as

$$f(t) = \frac{a_0}{2} + \sum_{n=1}^{\infty} \left(a_n \cos\left(\frac{2\pi nt}{T}\right) + b_n \sin\left(\frac{2\pi nt}{T}\right) \right), \quad (3.9)$$

where the set of coefficients $\{a_i\}_{i=0}^{\infty}$ and $\{b_i\}_{i=1}^{\infty}$ are constants to be determined.

Another equivalent formula for expressing the Fourier series in (3.9) is found by mapping the problem into the complex domain. To achieve this mapping first recall De Moivre's formula (Abramowitz and Stegun, 1964)

$$e^{i\alpha} = \cos(\alpha) + i \sin(\alpha). \quad (3.10)$$

Now it can be observed that

$$c_k e^{ikt} + c_{-k} e^{-ikt} = c_k (\cos(kt) + i \sin(kt)) + c_{-k} (\cos(kt) - i \sin(kt)) \quad (3.11)$$

$$= (c_k + c_{-k}) \cos(kt) + i(c_k - c_{-k}) \sin(kt) \quad (3.12)$$

$$= a_k \cos(kt) + b_k \sin(kt), \quad (3.13)$$

where

$$a_k = c_k + c_{-k} \quad (3.14)$$

and

$$b_k = i(c_k - c_{-k}). \quad (3.15)$$

Rearranging the equations in (3.14) and (3.15) it can be deduced that

$$c_k = \frac{a_k - ib_k}{2} \quad (3.16)$$

and

$$c_{-k} = \frac{a_k + ib_k}{2}. \quad (3.17)$$

Now if we apply a mapping from $t \mapsto \frac{2\pi t}{T}$ the formula for the Fourier series (3.9) can be expressed as

$$f(t) = \sum_{n=-\infty}^{\infty} c_n e^{int}. \quad (3.18)$$

Now observe that for integers n, k such that $n \neq k$

$$\int_0^{2\pi} e^{int} e^{-ikt} dt = \int_0^{2\pi} e^{i(n-k)t} dt = \frac{1}{i(n-k)} e^{i(n-k)t} \Big|_0^{2\pi} = 0, \quad (3.19)$$

since $e^{i(n-k)0} = e^{i(n-k)2\pi} = 1$. In the case where $n = k$ we have that

$$\int_0^{2\pi} e^{int} e^{-ikt} dt = \int_0^{2\pi} e^0 dt = 2\pi. \quad (3.20)$$

Therefore, a formula for the coefficients c_k can be written as

$$c_k = \frac{1}{2\pi} \int_0^{2\pi} f(t) e^{-ikt} dt. \quad (3.21)$$

3.3.3 Discrete Fourier Transform

The formula for the c_k can be discretised using a numerical integration method known as the trapezoidal rule (Davis and Rabinowitz, 2007). The aim of this numerical approach is to divide the area of integration into rectangles, the area of each being easily computed, and then sum these areas to get an approximation to the original integral.

The interval $[0, 2\pi]$ can be subdivided into intervals of length

$$h = \frac{2\pi}{N} \quad (3.22)$$

with boundaries at points

$$t_j = jh, \quad j = 0, \dots, N - 1. \quad (3.23)$$

On this subdivision the trapezoidal rule produces the formula

$$\frac{1}{2\pi} \int_0^{2\pi} f(t)e^{-ikt} dt \approx \frac{h}{2\pi} \sum_{j=0}^{N-1} f_j e^{-ikt_j} := F_k \quad (3.24)$$

where $f(t_j) = f_j$. The formula in (3.24) is known as the discrete Fourier transform of the function f .

Some further simplification can be undertaken on the discrete Fourier transform which will be particularly useful when describing the calculation method known as the fast Fourier transform later on. Let

$$w = e^{ih} \quad (3.25)$$

then using the notation of complex conjugation

$$\bar{w} = e^{-ih}. \quad (3.26)$$

Now the discrete Fourier transform can be expressed as

$$F_k = \frac{h}{2\pi} \sum_{j=0}^{N-1} f_j e^{-ikt_j} = \frac{1}{N} \sum_{j=0}^{N-1} f_j (e^{-ih})^{jk} = \frac{1}{N} \sum_{j=0}^{N-1} f_j \bar{w}^{jk}. \quad (3.27)$$

It is also helpful to have the inverse of the discrete Fourier transform. Define

$$f_j = \sum_{k=0}^{N-1} F_k e^{ikt_j} = \sum_{k=0}^{N-1} F_k w^{jk}, \quad (3.28)$$

Now it can be seen that

$$f_j = \sum_{k=0}^{N-1} F_k w^{kj} = \sum_{k=0}^{N-1} \frac{1}{N} \sum_{m=0}^{N-1} f_m \overline{w}^{mk} w^{kj} = \frac{1}{N} \sum_{m=0}^{N-1} f_m \sum_{k=0}^{N-1} w^{(j-m)k}, \quad (3.29)$$

where,

$$\overline{w}^{mk} w^{kj} = e^{-ihmk} e^{ihkj} = e^{ih(j-m)k} = w^{(j-m)k}. \quad (3.30)$$

Since the sum $\sum_{k=1}^{N-1} w^{(j-m)k}$ is a geometric series it holds that for $j \neq m$

$$\sum_{k=0}^{N-1} w^{(j-m)k} = 1 + \sum_{k=1}^{N-1} w^{(j-m)k} \quad (3.31)$$

$$= 1 + w^{(j-m)} \left(\frac{1 - w^{(j-m)(N-1)}}{1 - w^{(j-m)}} \right) \quad (3.32)$$

$$= \frac{1 - w^{(j-m)N}}{1 - w^{(j-m)}} = 0 \quad (3.33)$$

where the fact that

$$w^{(j-m)N} = e^{ih(j-m)N} = e^{i2\pi(j-m)} = \cos(2\pi(j-m)) + i \sin(2\pi(j-m)) = 1 \quad (3.34)$$

has been used.

Notice that for the case where $j = m$ the sum on the left hand side of (3.31) is equal to N . This yields the result that the transformations defined above are indeed the inverse of one another. Importantly, having established formulae for the DFT and its inverse it is now possible to consider the problem of finding a Fourier series in terms of finding the values F_k .

3.3.4 Fast Fourier Transform

Although it is possible to compute the Fourier series coefficients from the discrete Fourier transform a significant reduction is possible to the required number of calculations. One approach which has been particularly successful at improving the speed of computing Fourier series is known as the Fast Fourier Transform (FFT). Algorithms to apply the FFT are numerous but the most commonly used approach is the Cooley-Tukey FFT (Cooley and Tukey, 1965). The essential idea of the Cooley-Tukey approach is to divide the discrete Fourier transform into two parts looking at odd and even terms.

Since the discrete Fourier transform can be viewed as a matrix-vector multiplication the order of complexity is simply given by $\mathcal{O}(N^2)$. However, if the total number of data points N is a power of 2 we can express

$$\begin{aligned}
 F_j &= \frac{1}{N} \sum_{k=0}^{N-1} f_k \bar{w}^{kj} \\
 &= \frac{1}{N} \sum_{\text{even } k} f_k \bar{w}^{kj} + \frac{1}{N} \sum_{\text{odd } k} f_k \bar{w}^{kj} \\
 &= \frac{1}{N} \sum_{k'=0}^{M-1} f_{2k'} (\bar{w}^2)^{k'j} + \frac{1}{N} \sum_{k''=0}^{M-1} f_{2k''+1} \bar{w}^{(2k''+1)j} \\
 &= \underbrace{\frac{1}{N} \sum_{k'=0}^{M-1} f_{2k'} (\bar{w}^2)^{k'j}}_{:= F'_j} + \bar{w}^j \underbrace{\frac{1}{N} \sum_{k''=0}^{M-1} f_{2k''+1} (\bar{w}^2)^{k''j}}_{:= F''_j}. \tag{3.35}
 \end{aligned}$$

We now have two series F'_j and F''_j which are discrete Fourier transforms but of length M which is half the initial length N . These smaller discrete Fourier transforms

can again be simplified in order to obtain a fast recursive formula known as the Fast Fourier Transform (FFT).

Recall that the index j represents the number of data points and that $j = 0, 1, 2, \dots, N - 1$. For $j = 0, \dots, M - 1$ we compute the discrete Fourier transform using the formula in (3.35). Then for $j = M, \dots, N - 1$ define

$$j' = j - M. \quad (3.36)$$

With this definition, and that $h = \frac{2\pi}{N}$, the following formulae are readily obtained.

$$\bar{w}^M = e^{-ihM} = e^{-ih\frac{N}{2}} = e^{-i\pi} = -1 \quad (3.37)$$

$$\bar{w}^N = \bar{w}^{2M} = 1 \quad (3.38)$$

$$\bar{w}^j = \bar{w}^{j'+M} = \bar{w}^{j'}\bar{w}^M = -\bar{w}^{j'} \quad (3.39)$$

$$(\bar{w}^2)^{kj} = (\bar{w}^j)^{2k} = (-\bar{w}^{j'})^{2k} = (\bar{w}^2)^{kj'}. \quad (3.40)$$

Therefore, the following recursive formula follows immediately for $j = 0, \dots, M - 1$

$$F_j = F'_j + \bar{w}^j F''_j \quad (3.41)$$

$$F_{j+M} = F'_j - \bar{w}^j F''_j. \quad (3.42)$$

Observe that with repeated use of this recursive relationship the F'_j and F''_j can be computed with half the number of steps on each application. This is the essential step which makes the fast Fourier transform so efficient.

The improvement of the fast Fourier transform over the original brute force approach can be seen by computing the new order of complexity. First assume that all the w^j are calculated and let $W(N)$ be the required number of operations for the FFT with N points. Then we have that

$$W(2M) = 2W(M) + 4M \quad (3.43)$$

since we split the problem in half arriving at the term $2W(M)$ and then have the multiplications $\bar{w}^j F_j''$ followed by the addition $+\bar{w}^j F_j''$ and subtraction $-\bar{w}^j F_j''$ and recombination $F_j + F_{j+M}$ yielding the term $4M$. Note also that $W(1) = 0$ since there is nothing to compute on just a single point.

Now if the number of points being considered in the FFT are a power of 2 then $N = 2^n$. Let $\omega_j = W(2^j)$ so that

$$\omega_j = 2\omega_{j-1} + 2(2^j), \quad (3.44)$$

with $\omega_0 = 0$.

Multiplying through (3.44) by a factor of 2^{n-j} and summing we arrive at

$$\sum_{j=1}^n 2^{n-j} \omega_j = 2 \sum_{j=1}^n 2^{n-j} (\omega_{j-1} + 2^j) \quad (3.45)$$

$$= 2n2^n + \sum_{j=1}^n 2^{n-j+1} \omega_{j-1} \quad (3.46)$$

$$= \{\text{substituting } k = j - 1 \text{ into the summation}\}$$

$$= 2n2^n + \sum_{k=0}^{n-1} 2^{n-k} \omega_k \quad (3.47)$$

$$= 2n2^n + \sum_{k=1}^{n-1} 2^{n-k} \omega_k, \quad (3.48)$$

where the fact that $\omega_0 = 0$ has been used in the last step to discard the first term of the sum. Observe that final summation on the right hand side of the equal sign and the one on the left run over the same indices, except the n th term, thus giving the final equation

$$\omega_n = 2n2^n. \quad (3.49)$$

Finally,

$$W(N) = \omega_n = 2n2^n = 2N \log_2 N. \quad (3.50)$$

Therefore, the order of complexity of the FFT is $\mathcal{O}(N \log_2 N)$ which is a vast improvement over the original brute force approach with order $\mathcal{O}(N^2)$.

A few final observations about the FFT make the algorithm output easier to follow in numerical applications. Firstly observe that

$$F_0 = \frac{1}{N} \sum_{k=0}^{N-1} f_k. \quad (3.51)$$

Therefore, F_0 is simply an average of the input values. As such it doesn't provide much additional detail and is usually not included in graphs of FFT outputs.

Another remark about the FFT is that for real valued inputs

$$\overline{F_j} = F_{N-j}. \quad (3.52)$$

This can be proven by observing that

$$F_{N-j} = \frac{1}{N} \sum_{k=0}^{N-1} f_k e^{-ihk(N-j)} \quad (3.53)$$

$$= \frac{1}{N} \sum_{k=0}^{N-1} f_k e^{-i\frac{2\pi}{N}k(N-j)} \quad (3.54)$$

$$= \frac{1}{N} \sum_{k=0}^{N-1} f_k e^{-i2\pi k} e^{i\frac{2\pi}{N}kj} \quad (3.55)$$

$$= \frac{1}{N} \sum_{k=0}^{N-1} f_k e^{i\frac{2\pi}{N}kj} \quad (3.56)$$

$$= \overline{F_j} \quad (3.57)$$

where the assumption that the f_k are real valued has been used in the final step.

The reason the symmetrical relationship $\overline{F_j} = F_{N-j}$ is important is that when applying the FFT numerically it is only required to retrieve half of the output. This simplifies the analysis of FFT applications.

A very fast implementation of the FFT, and the one used in this chapter, is contained in the C++ package FFTW3 (Frigo and Johnson, 2005).

3.3.5 Fourier Analysis Application

Recall that the aim of using the Fourier series is to quantify the periodicities in the monthly rainfall and AS event data. In particular, Fourier series will be fitted to the monthly counts of each of these datasets using the FFT. This will be repeated for each of the phase 3 and 4 datasets within the BT and VAC categories. Since the aim of using the FFT is to compare the periods in the AS events and the rainfall data

over the same time period it is necessary to ensure that the datasets analysed begin and end on the same month.

Let N be the number of bins and let each bin be represented by $i \in \{0, \dots, N-1\}$. We use the left edge of each bin and corresponding bin height as input to the FFT.

We have from the symmetry properties of the FFT output in (3.52) that the discrete Fourier series is given by

$$f(t) = \frac{F_0}{N} + \frac{2}{N} \sum_{k=1,2,\dots,\text{ceil}(\frac{N}{2})} a_k \cos\left(\frac{2\pi tk}{N}\right) - b_k \sin\left(\frac{2\pi tk}{N}\right), \quad (3.58)$$

where the ceil indicates that the number should be rounded upwards and the complex numbers $F_k = a_k + ib_k$ for real valued a_k and b_k . The $\frac{1}{N}$ scaling factor of coefficients results from the computation method used by the FFTW3 package (Frigo and Johnson, 2005).

In cases where a linear detrend has been applied to the data prior to using the FFT the Fourier series takes the form

$$f(t) = \alpha t + \beta + \frac{F_0}{N} + \frac{2}{N} \sum_{k=1,2,\dots,\text{ceil}(\frac{N}{2})} a_k \cos\left(\frac{2\pi tk}{N}\right) - b_k \sin\left(\frac{2\pi tk}{N}\right). \quad (3.59)$$

On the other hand, when using the mean detrending method the Fourier series is computed as

$$f(t) = m + \frac{F_0}{N} + \frac{2}{N} \sum_{k=1,2,\dots,\text{ceil}(\frac{N}{2})} a_k \cos\left(\frac{2\pi tk}{N}\right) - b_k \sin\left(\frac{2\pi tk}{N}\right), \quad (3.60)$$

where m is the mean number of events.

Note that the rainfall data is already provided on a monthly scale. However, the AS event data is provided on a daily timeframe. Therefore, to convert each AS datapoint, denoted by t , to a monthly scale the transformation $\frac{t}{30}$ must be applied, where, an approximation of one month equaling 30 days is used.

In this case the equations in (3.58)-(3.60) become

$$f(t) = \frac{F_0}{N} + \frac{2}{N} \sum_{k=1,2,\dots,\text{ceil}(\frac{N}{2})} a_k \cos\left(\frac{2\pi tk}{30N}\right) - b_k \sin\left(\frac{2\pi tk}{30N}\right) \quad (3.61)$$

$$f(t) = \alpha t + \beta + \frac{F_0}{N} + \frac{2}{N} \sum_{k=1,2,\dots,\text{ceil}(\frac{N}{2})} a_k \cos\left(\frac{2\pi tk}{30N}\right) - b_k \sin\left(\frac{2\pi tk}{30N}\right) \quad (3.62)$$

$$f(t) = m + \frac{F_0}{N} + \frac{2}{N} \sum_{k=1,2,\dots,\text{ceil}(\frac{N}{2})} a_k \cos\left(\frac{2\pi tk}{30N}\right) - b_k \sin\left(\frac{2\pi tk}{30N}\right). \quad (3.63)$$

When there is an even number of input data points, so that, $\text{ceil}(\frac{N}{2}) = \frac{N}{2}$, there is no need to multiply the $\frac{N}{2}$ term inside the summations in (3.58)-(3.63) by 2. The frequency at this term is known as the Nyquist frequency (Grenander, 1959) which only occurs once in the spectrum marking the point before the complex symmetry described in (3.52) occurs.

In addition, it should be noted that with the definition of the Fourier series in (3.58)-(3.63) the frequency of each term is given by

$$\frac{k}{N} \quad (3.64)$$

and, therefore, the periodicity associated to each term is found via the inverse

$$\frac{N}{k}. \quad (3.65)$$

The formula in (3.65) will be used to compare the lengths of the periods in the event and rainfall data.

Observe that the equations in (3.58)-(3.63) use all possible Fourier series terms derived from the FFT. As more terms are used the series becomes more sensitive to noise as opposed to the important long term periodicities. An approach to minimise this is to concentrate only on those terms which have the most significant impact on the Fourier series function. To determine significance the absolute value of the coefficients of each term can be computed with larger values indicating more importance. In the complex plane a number of the form $a + ib$ can be considered as a vector with absolute value $|a + ib| = \sqrt{a^2 + b^2}$. Therefore, the values $|a_k + ib_k|$, for coefficients a_k, b_k as defined in (3.58)-(3.63), can be computed to extract the most important terms. This information is usually displayed in the form a power spectrum displaying each term's coefficient magnitude against frequency. The Fourier series of (3.58)-(3.63) can then be reconstructed keeping only those terms with the largest relative absolute value. For the analyses conducted in this chapter only the two highest powered frequencies will be included in the final Fourier series approximations.

Since the focus of using the Fourier series is to find the time varying dynamics via periodicities in the rainfall and AS event databases the constant leading terms $\frac{E_0}{N}$, β and m will be dropped from the Fourier series.

Another issue of concern when utilising the FFT is that the window of time

over which the data is observed should contain full periods. As detailed by Staudt (1998) when this assumption is not met a problem known as spectral leakage may be introduced into the FFT output. When spectral leakage occurs the power spectrum of the FFT will be distorted with insignificant frequencies having high powers. One approach to overcome this issue is to introduce a window function prior to performing FFT calculations. Simply applying the FFT is the equivalent of using a rectangular window function which takes the value 1 inside the timeframe under consideration and 0 elsewhere. A second common type of window function is known as the Hanning window (Harris, 1978; Staudt, 1998). The Hanning window acts to smooth the end points of the dataset so that the data “wraps” around ensuring full periodicities in the FFT. However, since this chapter uses finite length datasets the issue of extrapolation beyond the final timestamps is not of concern. Moreover, only the two frequencies with highest power will be used in the Fourier series thus providing a filter for noisy terms. Therefore, the rectangular window will be used henceforth.

3.3.6 Results of Fourier Series Analysis

In this section the foundations of Fourier series analysis are used to extract the two most important periods for each of the AS and weather datasets. From this analysis the periods can be compared to determine the relationship between rainfall in Somalia and AS terrorism.

Tables 3.7 - 3.10 provide the most significant periods found using the FFT. In Figures 3.4 - 3.5 the histograms, detrended histograms, FFT power spectrums and Fourier series functions are presented for the rainfall data in phases 3 and 4 cor-

responding to the BT dataset. Similar results are obtained for the rainfall data corresponding to the VAC datasets and, thus, are omitted. The results of the same analyses for BT and VAC AS events are displayed in Figures 3.6 - 3.9.

Table 3.7: Battle Territory Phase 3 Periods

Power Spectrum Rank	Event Period (Months)	Rainfall Period (Months)
1	2.9	5.6
2	2.7	6.5

Table 3.8: Battle Territory Phase 4 Periods

Power Spectrum Rank	Event Period (Months)	Rainfall Period (Months)
1	54	5.9
2	27	2.9

Table 3.9: Violence Against Civilians Phase 3 Periods

Power Spectrum Rank	Event Period (Months)	Rainfall Period (Months)
1	2.2	6.3
2	2.3	5.4

Table 3.10: Violence Against Civilians Phase 4 Periods

Power Spectrum Rank	Event Period (Months)	Rainfall Period (Months)
1	54	5.9
2	27	2.9

BT Rainfall Phase 3

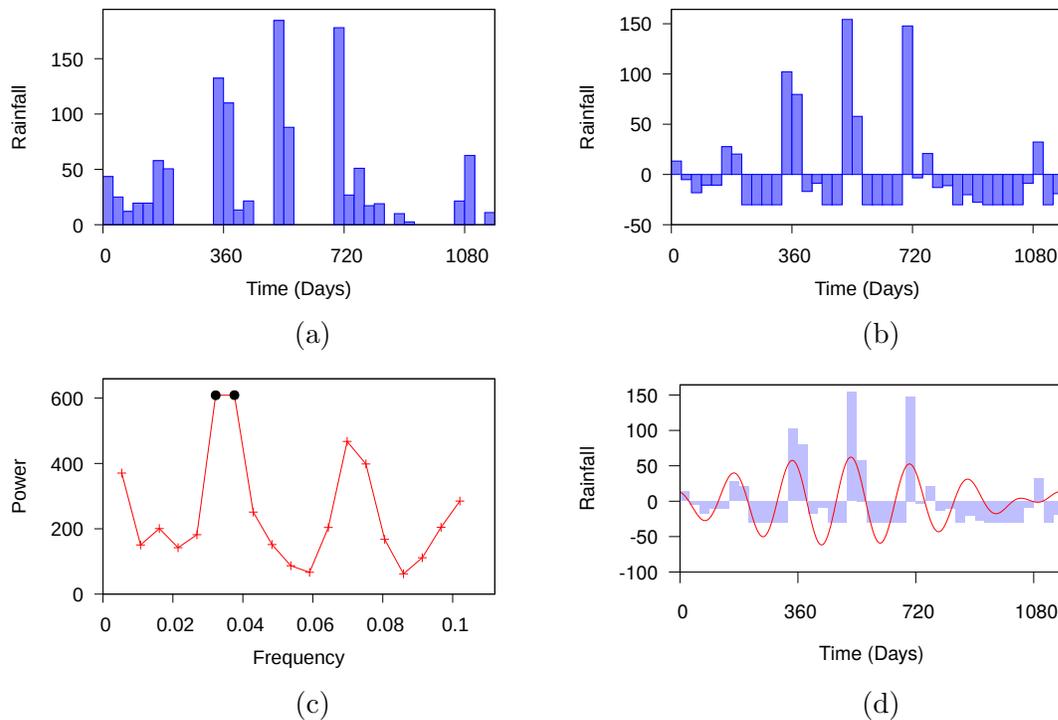


Figure 3.4: (a) Battle territory phase 3 weather data histogram. (b) Battle territory phase 3 weather data mean detrended histogram. (c) Battle territory phase 3 monthly weather data FFT frequencies power spectrum. The higher the power the greater the significance of the frequency. The highlighted points (black circles) correspond to the two frequencies with highest power which are used to construct a Fourier series approximation for the data. (d) Battle territory phase 3 weather data mean detrended histogram with Fourier series constructed from two frequencies with highest power.

BT Rainfall Phase 4

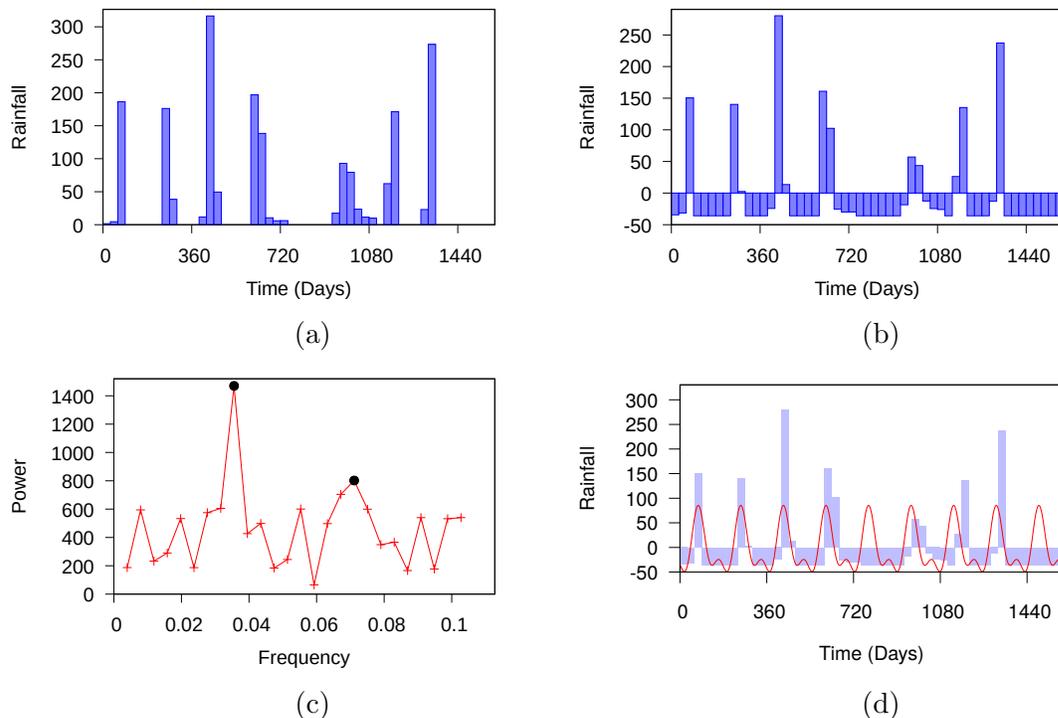


Figure 3.5: (a) Battle territory phase 4 weather data histogram. (b) Battle territory phase 4 weather data mean detrended histogram. (c) Battle territory phase 4 monthly weather data FFT frequencies power spectrum. The higher the power the greater the significance of the frequency. The highlighted points (black circles) correspond to the two frequencies with highest power which are used to construct a Fourier series approximation for the data. (d) Battle territory phase 4 weather data mean detrended histogram with Fourier series constructed from two frequencies with highest power.

Battle Territory Phase 3

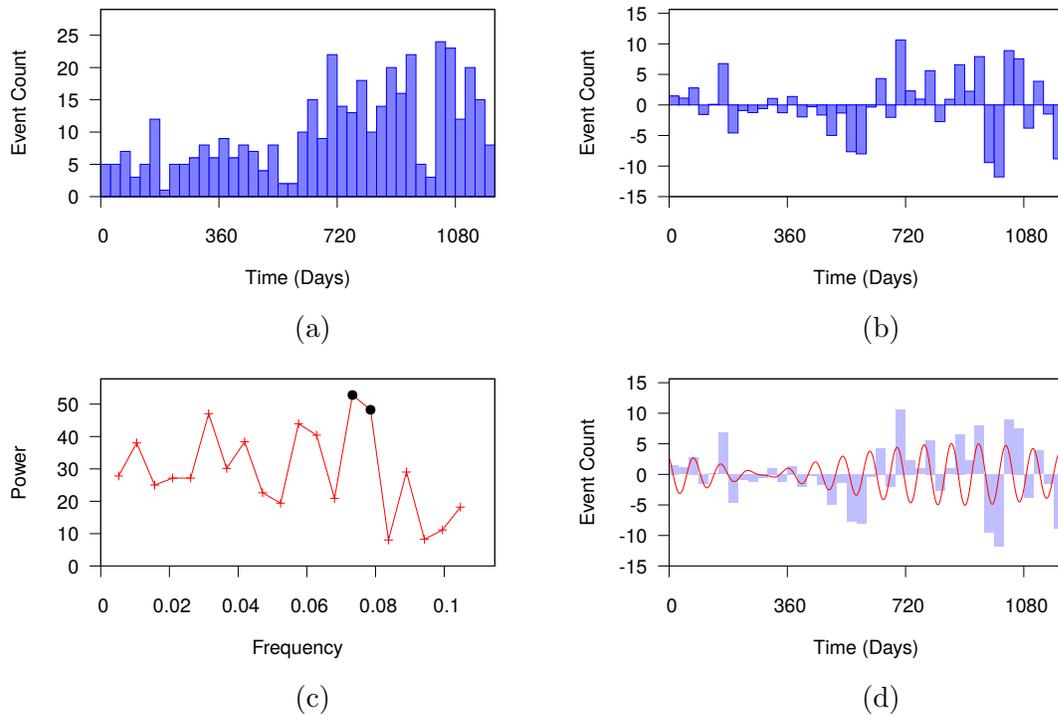


Figure 3.6: (a) Battle territory phase 3 event data histogram. (b) Battle territory phase 3 event data linearly detrended histogram. (c) Battle territory phase 3 monthly event data FFT frequencies power spectrum. The higher the power the greater the significance of the frequency. The highlighted points (black circles) correspond to the two frequencies with highest power which are used to construct a Fourier series approximation for the data. (d) Battle territory phase 3 event data linearly detrended histogram with Fourier series constructed from two frequencies with highest power.

Battle Territory Phase 4

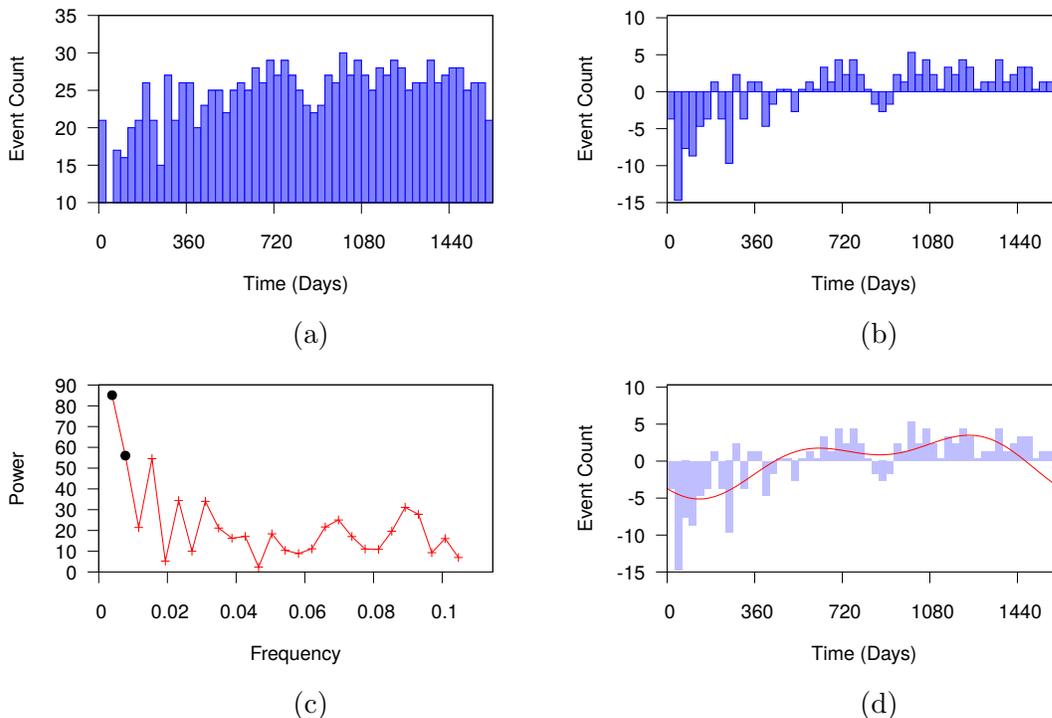


Figure 3.7: (a) Battle territory phase 4 event data histogram. (b) Battle territory phase 4 event data mean detrended histogram. (c) Battle territory phase 4 monthly event data FFT frequencies power spectrum. The higher the power the greater the significance of the frequency. The highlighted points (black circles) correspond to the two frequencies with highest power which are used to construct a Fourier series approximation for the data. (d) Battle territory phase 4 event data mean detrended histogram with Fourier series constructed from two frequencies with highest power.

Violence Against Civilians Phase 3

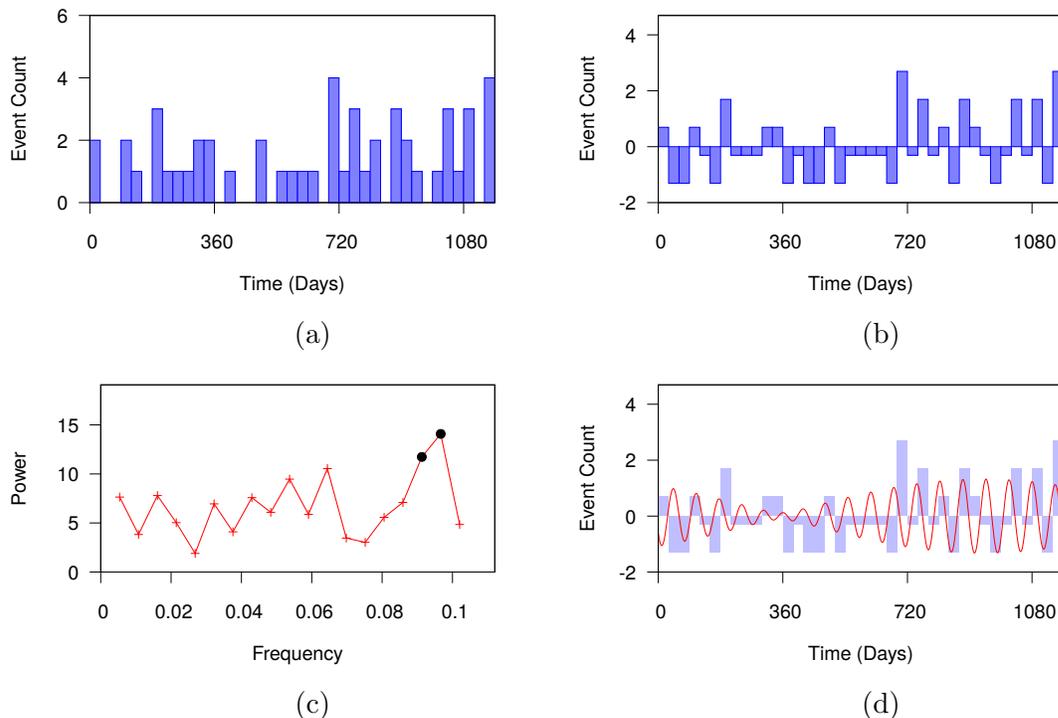


Figure 3.8: (a) Violence against civilians phase 3 event data histogram. (b) Violence against civilians phase 3 event data mean detrended histogram. (c) Violence against civilians phase 3 monthly event data FFT frequencies power spectrum. The higher the power the greater the significance of the frequency. The highlighted points (black circles) correspond to the two frequencies with highest power which are used to construct a Fourier series approximation for the data. (d) Violence against civilians phase 3 event data mean detrended histogram with Fourier series constructed from two frequencies with highest power.

Violence Against Civilians Phase 4

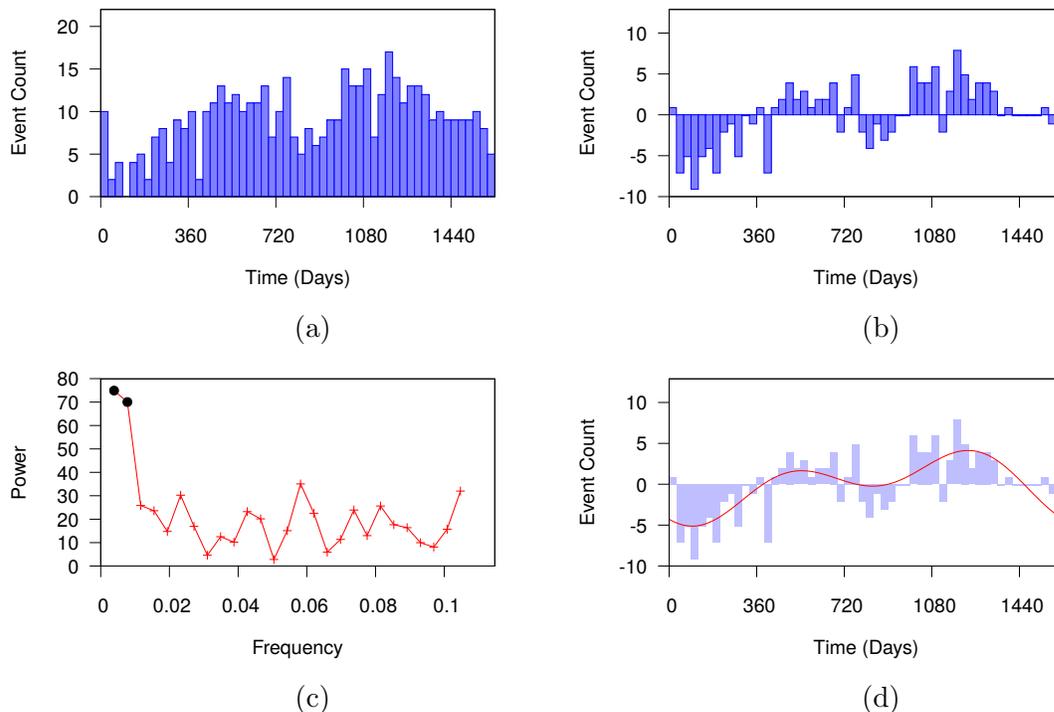


Figure 3.9: (a) Violence against civilians phase 4 event data histogram. (b) Violence against civilians phase 4 event data mean detrended histogram. (c) Violence against civilians phase 4 monthly event data FFT frequencies power spectrum. The higher the power the greater the significance of the frequency. The highlighted points (black circles) correspond to the two frequencies with highest power which are used to construct a Fourier series approximation for the data. (d) Violence against civilians phase 4 event data mean detrended histogram with Fourier series constructed from two frequencies with highest power.

Observe that the BT and VAC phase 3 periods are between 2-3 months whilst the weather data for these phases have periods in the range 5-6. Significantly the 3 month periods suggest that the four cycles of rainfall in Somalia from April-June and October-December appear as seasonal dynamics in AS attacks.

On the other hand, in phase 4 the AS data in both the BT and VAC categories have periods of 27 and 54 months. The rainfall datasets for phase 4 have periods of 2.9 and 5.9 months. Therefore, in phase 4 the AS attack pattern does not appear to show a yearly cycle as seen in phase 3.

One possible explanation for the difference between the periods of the rainfall and attack data in phase 4 is the changing strategies of AS. Phase 4 covers the timeline of renewed resistance against AS from international forces. In particular, in 2011 Kenyan troops crossed into South Somalia to resist the growing security threat of AS at Kenya's border. Simultaneously, the African Union Mission to Somalia (AMISOM) consisting of international forces were also fighting against AS. A result of all these pressures caused AS to retreat. In response AS shifted strategically to rely more on the use of IED attacks as well as engaging in high profile events, such as, an attack on the Westgate Mall in Nairobi in 2013. In contrast to military type assaults against government forces which require mobilisation of AS members and planning for weather related events, IED attacks can be organised in a more spontaneous fashion. Therefore, the cycles observed in the Fourier analysis of the phase 4 datasets may be demonstrating this difference in the strategies employed by AS.

3.4 Hawkes Process Modelling

3.4.1 Hawkes Process with Time Independent Background Rate

In Chapter 2 of this thesis the Hawkes self-exciting point process model was discussed in detail. The important components necessary for applying the Hawkes model will be presented here in summary.

The premise of the Hawkes model is that given a series of timestamps $\{t_i\}_{i=1}^N$ an intensity function of the form (Hawkes, 1971)

$$\lambda(t) = \mu + k_0 \sum_{t > t_i} \omega e^{-\omega(t-t_i)} \quad (3.66)$$

can be fitted using the method of maximum likelihood estimation (MLE) (Ozaki, 1979). In this case μ is a time independent baseline rate at which events occur, k_0 describes the rise in the rate of events following an initial event and ω is determined by the length of time a series of self-exciting events propagates.

It should be noted that in the simple case $\lambda(t) = \mu$ the model becomes a Poisson process (Ross, 2010). In this case the occurrence of events is assumed to have a background rate $\mu = \frac{N}{T}$ being the average number of events per unit time. This Poisson model can be used as a baseline for comparison with the Hawkes process model.

The parameter values for the model in (3.66) are computed by maximising the

formula

$$\log L = \sum_{i=1}^N \left[\log \left(\mu + k_0 \sum_{t_i > t_j} \omega e^{-\omega(t_i - t_j)} \right) + k_0 (e^{-\omega(T - t_i)} - 1) \right] - \mu T. \quad (3.67)$$

The overall fit of the Hawkes process model can be determined using the Kolmogorov-Smirnov test (Brown et al., 2002; Massey, 1951). For this test the following statistic is calculated

$$D_n = \max_k \left(\left| U_k - \frac{k-1}{N} \right|, \left| \frac{k}{N} - U_k \right| \right). \quad (3.68)$$

If $D_n < D_\alpha$ for some critical value D_α then there is evidence that the Hawkes process is capturing the dynamics present in the dataset. The values for U_i are evaluated using

$$U_i = 1 - \exp \left(- \left(\mu(t_i - t_{i-1}) + k_0 \sum_{j=1}^{i-1} [e^{-\omega(t_{i-1} - t_j)} - e^{-\omega(t_i - t_j)}] \right) \right). \quad (3.69)$$

A comparative fit between Hawkes process models, applied to the same dataset, can be determined using the Akaike Information Criterion (AIC) (Akaike, 1974)

$$\text{AIC} = 2k - 2 \log L, \quad (3.70)$$

where, k is the number of model parameters. A lower value of AIC implies a better fitting model.

The results of applying the simple Poisson process model to the datasets studied in this chapter are provided in Tables 3.11 - 3.12. None of the Poisson models prove

to be a good fit for the data.

Table 3.11: Battle Territory Poisson Process Results

Dataset	μ	AIC	KS Test	KS Critical 95%	KS Critical 99%
Phase 3	0.3449	1682	1.0000	0.0674	0.0808
Phase 4	0.8269	3175	1.0000	0.0372	0.0446

Table 3.12: Violence Against Civilians Poisson Process Results

Dataset	μ	AIC	KS Test	KS Critical 95%	KS Critical 99%
Phase 3	0.0444	422	1.0000	0.1904	0.2282
Phase 4	0.3075	2146	1.0000	0.0613	0.0735

Applying the Hawkes process model in the form of (3.66) to the “Battle Territory” and “Violence Against Civilians” datasets yields the results in Tables 3.13 - 3.14. From these results it can be seen that KS test statistics exceed the critical values except the VAC phase 3 dataset at the 99% critical value. However, the k_0 value in VAC phase 3 does not satisfy the requirements for a non-explosive Hawkes process as detailed in Chapter 2. Therefore, these results suggest that a Hawkes process as in (3.66) does not provide a sufficiently good fit to the Al Shabaab data.

Table 3.13: Battle Territory Time Invariant Background Hawkes Process Results

Dataset	μ	k_0	ω	AIC	KS Test	KS Critical 95%	KS Critical 99%
Phase 3	0.0962	0.7430	0.0562	1612	0.1970	0.0674	0.0808
Phase 4	0.5715	0.3845	0.0034	3161	0.4442	0.0372	0.0446

Table 3.14: Violence Against Civilians Time Invariant Background Hawkes Process Results

Dataset	μ	k_0	ω	AIC	KS Test	KS Critical 95%	KS Critical 99%
Phase 3	0.0321	1.0000	0.0007	424	0.2175	0.1904	0.2282
Phase 4	0.1288	0.6492	0.0064	2121	0.2136	0.0613	0.0735

3.5 Hawkes Process with Seasonal Background Rate

It has been shown that the AS events have a cyclical dynamic as shown in Tables 3.7 - 3.10. A consequence of this finding is that the time invariant background rate of the Hawkes process model in (3.66) may not be adequately well suited to capture the variation in the underlying rate of AS events. Therefore, an extension of this model is required which incorporates the FFT periodicities discovered previously.

A similar problem has been considered in the field of seismology. In the work of Ogata (1999) a review is made of modelling seasonal variations in earthquake activity. One of the pieces of research covered is by Ogata (1983) which demonstrates that an appropriate extension of the Hawkes process to incorporate time trends in the data being studied is

$$\lambda(t) = \mu + P(t) + C(t) + k_0 \sum_{t > t_i} \omega e^{-\omega(t-t_i)}. \quad (3.71)$$

Here $P(t)$ is a polynomial equation. A polynomial can be used to capture trends in the data such as linear or non-linear trends over time. On the other hand, the term $C(t)$ is a cyclic term consisting of trigonometric functions which can capture patterns, such as, seasonal variations in datasets.

Fitting a Hawkes process model without $P(t)$ and $C(t)$ to a set of timestamps

requires the use of the method of maximum likelihood estimation (MLE) (Ozaki, 1979). Discussed in detail in Chapter 2 of this thesis the MLE for a Hawkes process model with timestamps $\{t_i\}_{i=1}^N$, final time $t_N = T$ and intensity function $\lambda(t)$ is given by (Rubin, 1972)

$$\log L(\{t_i\}; \mu, k_0, \omega) = \sum_{i=1}^N \log(\lambda(t_i)) - \int_0^T \lambda(t) dt. \quad (3.72)$$

Here the aim is to find the parameter set $\{\mu, k_0, \omega\}$ which maximises this equation. A major issue that must be overcome when using the MLE is that the function $\log L$ is highly multimodal (Ogata and Akaike, 1982). Therefore, the starting conditions of any optimisation algorithm applied to (3.72) must be chosen carefully to obtain a meaningful parameter set. One way to find the Hawkes process parameters is to initialise the optimisation routine at multiple starting points and taking the resulting parameter set yielding the highest value of (3.72) (Egesdal et al., 2010).

When extending the Hawkes process to incorporate a time varying background rate as in (3.71) it is also necessary to determine the form of the functions $P(t)$ and $C(t)$. A possible approach to find these functions would be to determine them inside the MLE routine. For example, expressing the unknown functions as

$$P(t) = \sum_{n=0}^N a_n t^n \quad (3.73)$$

$$C(t) = b_0 + \sum_{n=1}^N b_n \cos\left(\frac{2\pi n t}{N}\right) + c_n \sin\left(\frac{2\pi n t}{N}\right) \quad (3.74)$$

we would then need to consider an MLE problem of the form

$$\log L(\{t_i\}; \{a_n\}, \{b_n\}, \{c_n\}, \mu, k_0, \omega). \quad (3.75)$$

In this chapter a better approach to studying the functions $P(t)$ and $C(t)$ has been presented. In particular, we determine their form prior to performing the MLE method. This provides both an additional layer of analysis from which to understand the data and takes advantage of the fast computational time of the FFT.

However, since the computations of $P(t)$ and $C(t)$ are based on monthly event counts it is unreasonable to assume that the coefficients found using FFT will be directly usable inside the Hawkes process model (3.71). Therefore, it is necessary to introduce a scaling factor using the following approach

1. Fit a background of the form $B(t) = P(t) + C(t)$ to the dataset using the FFT.
2. Fit a model of the form $\lambda(t) = A \times B(t) + \mu + k_0 \sum_{t > t_i} \omega e^{-\omega(t-t_i)}$ to the dataset using MLE as in (3.72), where, the loglikelihood now has parameters $\log L(\{t_i\}; A, \mu, k_0, \omega)$. Here A is a real number to be determined which scales the function B .

For the MLE calculations in this chapter the “Nelder-Mead” (see Section 2.3) routine from the C++ library “nlopt” (Johnson, 2017) is used to find the model parameters. Note that “Nelder-Mead” is a minimisation algorithm and hence the equivalent problem of finding parameters which minimised $-\log L$ was considered.

To overcome the issue of trying to find the minimum of the multimodal loglikelihood function the minimisation algorithm was started at 100 different initial vectors.

There are four parameters to find $\{A, \mu, k_0, \omega\}$. Since A is used to scale the Fourier series to daily timestamps it is reasonable to bound the search space to $[0, 1]$. Similarly the search space for parameter μ was $[0, 1]$. A necessary condition on k_0 to ensure non-explosion of the Hawkes process is that $k_0 < 1$ (Hawkes and Oakes, 1974) so the search space for this parameter was $[0, 1]$. Finally since the inverse parameter $\frac{1}{\omega}$ can be interpreted as the average number of days over which a series of self exciting events persists (Lewis et al., 2012) then for daily timestamps it is expected that $\omega < 1$. Hence the bounds on the decay parameter were also $[0, 1]$. Therefore, on each run of the minimisation algorithm a vector with four random numbers each in the range $[0, 1]$ was used as a starting value. After finding the Hawkes process parameters the KS test and AIC value of the models can be computed to judge the fit of the model.

3.6 Results

In this section the theoretical details discussed previously will be used in application to study Al Shabaab data using a Hawkes process model with Fourier series background rate. Each of the datasets “Battle for Territory” (BT) and “Violence Against Civilians” (VAC) are studied across phases 3 and 4 of the AS group. Tables 3.15 - 3.18 present the results of analysing a Poisson process, Hawkes process with constant background rate, Hawkes process with a Fourier series background rate consisting of the highest powered FFT frequency and a Hawkes process with a Fourier series background rate consisting of the two highest powered FFT frequencies.

Table 3.15: Battle Territory Phase 3 Hawkes Process Results

Model	A	μ	k_0	ω	AIC	KS	KS 95%	KS 99%
Poisson Process	–	0.3449	–	–	1682	1.0000	0.0674	0.0808
No Seasonality	–	0.0962	0.7430	0.0562	1612	0.1970	0.0674	0.0808
1 Fourier Term	0.0152	0.0832	0.4768	0.0748	1604	0.2067	0.0674	0.0808
2 Fourier Terms	0.0163	0.0802	0.4656	0.0729	1600	0.2026	0.0674	0.0808

Table 3.16: Battle Territory Phase 4 Hawkes Process Results

Model	A	μ	k_0	ω	AIC	KS	KS 95%	KS 99%
Poisson Process	–	0.8269	–	–	3175	1.0000	0.0372	0.0446
No Seasonality	–	0.5715	0.3845	0.0034	3161	0.4442	0.0372	0.0446
1 Fourier Term	0.0032	0.5854	0.3634	0.0034	3163	0.4436	0.0372	0.0446
2 Fourier Terms	0.0109	0.6428	0.2849	0.0030	3163	0.4522	0.0372	0.0446

Table 3.17: Violence Against Civilians Phase 3 Hawkes Process Results

Model	A	μ	k_0	ω	AIC	KS	KS 95%	KS 99%
Poisson Process	–	0.0444	–	–	422	1.0000	0.1904	0.2282
No Seasonality	–	0.0321	1.0000	0.0007	424	0.2175	0.1904	0.2282
1 Fourier Term	0.0089	0.0331	1.0000	0.0006	425	0.1018	0.1904	0.2282
2 Fourier Terms	0.0075	0.0332	1.0000	0.0006	425	0.0894	0.1904	0.2282

Table 3.18: Violence Against Civilians Phase 4 Hawkes Process Results

Model	A	μ	k_0	ω	AIC	KS	KS 95%	KS 99%
Poisson Process	–	0.3075	–	–	2146	1.0000	0.0613	0.0735
No Seasonality	–	0.1288	0.6492	0.0064	2121	0.2136	0.0613	0.0735
1 Fourier Term	0.0138	0.1729	0.4794	0.0075	2122	0.2094	0.0613	0.0735
2 Fourier Terms	0.0289	0.2779	1.0000	0.0001	2109	0.2080	0.0613	0.0735

3.7 Discussion

In this chapter a novel modelling framework has been presented which can be used to capture cyclical and history dependent dynamics in terrorist attacks. In particular, by utilising the FFT significant periodicities were identified and used to build a Fourier series. This Fourier series was then used as a background rate for the past dependent Hawkes point process model.

The chapter opened with an examination of the historical roots of the AS group. In particular, AS was found to have operated in four phases according to strategic developments (Anderson and McKnight, 2015a; Hansen, 2013; Marchal, 2009; Maszka, 2017; Wise, 2011). These phases spanned the period from the groups inception through to December 2007, January 2008 to April 2008, May 2008 to July 2011 and August 2011 to the present time.

During the AS group's history there were three leaders as follows

- **2006-2008:** Aden Hashi Ayro (killed by a US airstrike).
- **2008-2014:** Ahmed Abdi Godane (killed by a US drone strike).
- **2014 - Present:** Abu Ubeyda.

Initially AS had success in fighting against government forces. To reverse the gains achieved by AS a force consisting of international troops known as the African Union Mission in Somalia (AMISOM) was sent into Somalia. In addition, to provide a bulwark against AS on the border Kenyan troops also moved into South Somalia. This combination of international forces had the effect of AS losing ground.

As a consequence of the pressure placed on AS the group evolved strategically. Specifically, the group relied more on attacks involving IEDs. Accompanying this shift in tactics AS also began launching high profile international events including an attack on the Westgate Mall in Nairobi in 2013 (Williams, 2014) and another on the Garissa University in 2015 (Lyons et al., 2015). The incorporation of less discriminant attacks by AS lead to numerous civilian deaths.

From these observations concerning AS the datasets studied in this chapter were analysed in four sets - BT (phases 3 and 4) and VAC (phases 3 and 4). Phases 1 and 2 were excluded from analysis due to a limited number of datapoints.

Section 3.3 presents the results of using the FFT to find periodicities in the AS event data in monthly counts and Somalia rainfall data also in monthly counts. It was shown that the BT and VAC phase 3 datasets followed a four cycle period each year. On the other hand, the BT and VAC phase 4 datasets followed a much longer 27 and 54 months cycles.

These observations concerning the periodicities in the AS data can be interpreted in terms of the real-world changes the group underwent.

In phase 3 the group was more involved in fighting government forces. Since these type of events involved more traditional military tactics organising AS members and accounting for weather patterns influenced the trends seen in the groups attack dynamics.

In contrast during phase 4 AS had been forced to retreat by the arrival of international forces in Somalia. Hence the group relied more on IED attacks which could be organised in a more spontaneous manner.

Having studied the cyclical nature of AS events this chapter focused on generating a model capable of capturing any past dependent behaviour in the attacks. This was achieved via the use of seasonal background rate in a Hawkes self-exciting point process model. The background rate took the form of a Fourier series constructed from the two frequencies with highest power identified with the FFT.

For the BT phase 3 dataset the best performing model is the seasonal Hawkes process model with a two term Fourier series background rate. Also for this dataset the simple stationary background rate Hawkes process outperforms a Poisson model. Therefore, despite no overall goodness of fit there does appear to be some evidence for historical dependence between the datapoints.

For the dataset of BT phase 4 the Hawkes process with constant background rate has the lowest AIC value. Hence there appears to be historical dependence in this dataset but the inclusion of seasonal trends does not provide any improvement.

The best performing model, according to the AIC, for the VAC phase 3 dataset was a simple Poisson process. However, it can be observed from the KS test statistics that the seasonal Hawkes process models have a significant fit. Moreover, the AIC values of the Poisson and Hawkes models are similar. This provides some evidence for both cyclical and historical dependence. However, it should be noted that the values of k_0 for the VAC phase 3 models are not in the allowed range for a Hawkes process.

In the results for the dataset VAC phase 4 the seasonal Hawkes process with two Fourier terms has the lowest AIC value, and, is thus the best performing model. Hence evidence exists for both a cyclical dynamic and historical dependence in the

dataset. As was noted for the phase 3 VAC models the value of $k_0 = 1$ for the VAC phase 4 Hawkes process model with a two termed Fourier series background is not a valid value for the Hawkes model.

In addition to studying the quantitative results from the analyses in this chapter it is also insightful to relate the results to real-world events.

For the scaling parameter only small changes are observed for the seasonal Hawkes process models across the datasets. This implies that the magnitude of the seasonal effects are consistent. From this observation it can be deduced that seasonal trends are a stable dynamic in each AS dataset.

The background rate of the seasonal Hawkes process models show a greater shift across the data phases. For the BT and VAC phase 3 datasets the parameter μ has smaller values than for the corresponding phase 4 datasets. A possible reason for this observation was that phase 4 was associated with a greater usage of IED attacks which may be easier to launch leading to a higher background rate of attacks.

After an event the Hawkes process model aims to capture a boost in successive events via the parameter k_0 . There is a significant change in this parameter value between the phase 3 and 4 datasets. In particular, after attacks in phase 3 the boosts were much greater than those observed in phase 4. Phase 3 was associated with the arrival of forces made up of Ugandan and Burundian troops aimed at countering AS. Therefore, the boost parameter in the seasonal Hawkes model may be reflecting this increased domestic focus of AS. On the other hand, in phase 4 AS was forced to retreat from territory and adopted a wider use of IED and high profile international attacks. These types of attacks may have been more standalone in nature meaning

follow up events were less likely.

Since a boosted level of attacks is unsustainable the parameter ω controls the decay rate of related events. Across phases the decay parameters do not change significantly for either the BT or VAC datasets. This implies the length of the memory of AS attacks was consistent under each category. This implies that in the longer term tactics used by Godane, the second leader of AS, were continued under the group's third leader, Ubeyda. This observation is similar to that of (Maszka, 2017) that Ubeyda has not yet altered the tactical approach of AS.

In addition to analysing the results across the BT and VAC datasets it is also interesting to discuss the differences between corresponding phases. For the BT and VAC phase 3 datasets the background rates μ are slightly higher for BT events whilst the boost is greater for the VAC events. The decay factors are greater for the BT events. This observations suggest that the BT events in phase 3 occurred at a slightly higher average rate and decade much faster than VAC events. However, after an initial attack the VAC phase 3 events had a greater increase in rate. These results may be explained by AS attacks being focused on obtaining territory from governmental forces in phase 3. Once territory was obtained further attacks would be unnecessary. On the other hand violence against civilians could have been used for population control which would need to be maintained during AS occupation. This observation resembles a similar discussion by Maszka (2017) who notes that Godane aimed to impose strict laws on the people in AS held territory.

Between BT and VAC phase 4 the values of μ move from higher to lower. On the other hand the boost values are higher in the VAC category. In contrast, the decay

parameters for VAC phase 4 are similar to those for BT except for the seasonal Hawkes process with a two termed Fourier series for the VAC dataset where the decay parameter is smaller than for BT. The results for the value of the background rate μ and boost k_0 are similar for the AS events studied in phase 3. However, the similarity of the decay rates in phase 4 demonstrate a shift in tactics by the AS group. Specifically, the group was losing territory in phase 4 and had begun relying on IED attacks. The usage of more indiscriminate methods could be leading to convergence in BT and VAC attack profiles.

Although this chapter has provided little evidence for statistical significance of the seasonal Hawkes process model for AS attacks some interesting observations have still be obtained. Firstly, by utilising the FFT prior to the Hawkes model fitting it was possible to confirm insights about seasonal trends in the event datasets. In particular, some evidence was found for Hypothesis 3 that AS attacks have a 3 month cycle relating to rainy seasons in Somalia. Having obtained the most significant periods present in the AS event data a seasonal Hawkes process model was derived. Specifically, the two FFT periods of highest power were used to construct a Fourier series background rate for the Hawkes model. The analyses of these seasonal models did not provide much evidence for goodness of fit. Therefore, there is little evidence to support Hypothesis 1 in the case of AS that there is historical dependence in the data similar to earthquakes and aftershocks. However, after fitting the full seasonal Hawkes process the model parameters provided a link between the models and the real world event occurrences. Hence, some evidence exists for Hypothesis 4 that the Hawkes model can be used as a real-time analysis tool. Moreover, the use of the

extensions of the Hawkes process in this chapter demonstrate how the model can be used to study terrorism internationally. In particular, the research conducted in this chapter illustrates the usage of the Hawkes model to study terrorism in a non-western setting. Developing universal models has implications for counter-terrorism practitioners and researchers by providing the foundations of versatile modelling approaches adaptable to a variety of environmental conditions.

There are several possible approaches to extending the work presented in this chapter. A simple change that could be taken within the framework already established would be to consider different kernels for the Hawkes process. Here an exponential self-exciting kernel was studied but as discussed by Ogata (1999) there are numerous other approaches that can be studied. Another direction that could be taken to improve the model developed in this chapter is to consider a marked point process (Daley and Vere-Jones, 2003). A marked point process could be adapted to explicitly model weather conditions inside the point process model. Finally, it could be possible to incorporate the change point detection methods from Chapter 2 of this thesis to confirm whether the points used to delineate Al Shabaab's phases were mathematically correct.

CHAPTER 4

GEOGRAPHIC PROFILING USING THE DIRICHLET PROCESS MIXTURE MODEL

A fundamental principle in the field of crime science is that environmental factors have a significant influence on the occurrence of crime (Brantingham and Brantingham, 1981; Brantingham and Brantingham, 1993a; Eck, 2010). These factors can be disaggregated into three classes - spatial, temporal and spatio-temporal. In previous chapters covered in this thesis the main focus has been on modelling temporal dynamics of terrorist attacks. On the other hand, spatial components of attacks have been treated implicitly via data segmentation according to geographic locations of interest. Hence there is potential that some important insights may have been missed concerning the underlying factors driving terrorist incidents. Therefore, in this chapter the temporal studies of IED attacks by the Provisional Irish Republican Army will be further enhanced through consideration of a spatial model.

Motivation to study spatial patterns in the field of Crime Science stems from observations made concerning urban crime. Currently the predominant theory in this setting, as discussed in Cornish and Clarke (1986), is that an offenders spatial targeting of where to commit a crime is based on a rational thought process. This theory is also padded with the caveat that this rational approach is based on a utility maximisation based on multiple factors, such as, risks vs rewards, knowledge, time constraints and possible resource requirements needed for the offence. Building on the rational offender view is the crime pattern theory of Brantingham and Brantingham (1993a). Here spatial patterns are modelled based on the assumption that offences occur at locations coinciding with an offenders awareness space developed as a result of the offenders routine activities.

Combining the approaches of these theories (Cohen and Felson, 1979) crimes

are more likely to occur when there exists an overlap between the awareness space of an offender and targets which are lacking guardianship. These hypotheses have been demonstrated to be correct in the urban crime studies literature. For example, crimes, such as, burglary and robbery have higher probability to be located on road networks that are widely known (Beavon, Brantingham and Brantingham, 1994; Davies and Johnson, 2015; Johnson and Bowers, 2010) as well as potential nodes, such as, shops and bars, within an offenders awareness space (Bowers, 2014; Groff and Lockwood, 2014).

The spatial targeting methods of offenders have also been shown to be similar in nature to ways animals are observed to pursue foraging strategies (Johnson, Summers and Pease, 2009). In particular, evidence has been found that offenders choose to commit crimes at locations which have proven successful in the past (Bernasco, 2008; Bernasco, Johnson and Ruiter, 2015; Johnson, 2008; Johnson, Summers and Pease, 2009; Pitcher and Johnson, 2011). The theory underpinning this foraging behaviour is that offenders are limited in their journey-to-crime distance. Therefore, repeatedly targeting a geographic location can ensure that good knowledge of an area is obtained which can be used to ensure future successful offences. On the other hand, this type of offending pattern can both deplete the available resources in an area and attract the attention of authorities. As a result of these two factors there is a shift in the offenders risk-reward analysis whereby after some period of criminal activity in an area the offender will move to a new geographical target.

Parallels may be drawn between the spatial patterns of criminal offenders and those of terrorist attacks. Specifically terrorist groups although existing under the

name of their organisations can actually carry out attacks at a more refined smaller group level. An example of this occurred within the Provisional IRA which, after suffering heavy infiltration by British Security Forces, splintered into a cellular organisation in 1977 (Asal et al., 2013). The terrorist actors versus state actors at this group scale can now be treated somewhat analogous to the urban criminal versus a police force. Both situations present an asymmetry between the resources accessible to the actors on each side. It has been argued in the literature (Johnson and Braithwaite, 2009; Townsley, Johnson and Ratcliffe, 2008) that this observation makes it suitable to apply Crime Science approaches to study spatial dynamics of terrorist attacks.

One area of overlap between cell based terrorist group attacks and urban crimes is the focusing of spatial targets as described previously. It was a specified aim of the PIRA, presented in the groups “Green Book” (O’Brien, 1999), to wage a “war of attrition against enemy personnel which is aimed at causing as many casualties and deaths as possible so as to create a demand from their people at home for their withdrawal”. This method of trying to exhaust the enemy centered on a widescale campaign of improvised explosive device (IED) attacks with Coogan (2002) explaining that PIRA had the ability to “turn ... bombing[s] on and off like a tap”. In this respect the cells within PIRA were operating in a similar way to an urban criminal targeting an area before moving on when the risk of detection by authorities outweighs any rewards. Investigations of the relationship between insurgent and counterinsurgent strategies (Braithwaite and Johnson, 2012; Johnson et al., 2011) have revealed that space-time patterns hold similar to those of urban offender theo-

ries described earlier.

In this chapter research into spatial clustering of terrorist attacks will be expanded through research into the relationship between PIRA IED attacks and the closure of IED factories by British Security Forces. The methodology used to study the spatial patterns of attacks in this chapter are based on research into geographic profiling (GP). In essence GP models incorporate the theoretical aspects of environmental influences from crime science and mathematical techniques that analyse groupings of objects to determine likely geographical locations for the source of criminal activities (Rossmo, 2000). Although such approaches have been proven effective to understand the spatial nature of offending (Canter et al., 2000; Levine, 2006; Miller, 2003; Verity et al., 2014) there is a gap for their extension into terrorism literature.

For the research presented here a GP model known as the Dirichlet Process Mixture Model (DPMM) will be introduced and used to study spatial components of terrorist attacks. The DPMM is a method based on Bayesian analysis which uses the relationship between prior and posterior distributions to convert the locations of events into information pertaining to the most likely sources of the events (Verity et al., 2014). The DPMM will be used here in a novel approach which studies how the grouping of IED attacks is influenced by the closure of IED factories. This research can be directly employed by those developing counter-terrorism policies to inform best practise approaches.

4.1 Datasets

There were two separate datasets used in the analyses undertaken in this chapter corresponding to PIRA IED attacks and IED factory discoveries by British Security Forces. Both of these datasets were collated from open sources as part of a research project by Asal et al. (2013).

The IED dataset was derived from open source accounts of attacks in Northern Ireland using LexisNexis and the Irish Times reporting archive. In total the database consists of 5461 entries. For the research conducted in this chapter the main fields extracted from the database were times and locations of the events. Timing data for the attacks was measured on a daily scale in the form of year/month/day. Location data was recorded on a county level scale covering the six counties of Northern Ireland (Antrim, Armagh, Derry, Down, Fermanagh, Tyrone) and the capital city (Belfast). Further data also exists for the border counties of Cavan, Donegal, Leitrim, Louth, and Monaghan. However, the focus in this chapter was restricted to Belfast, which, as the capital city witnessed a high volume of terrorist activity (Bell, 2000; Coogan, 2002; Fay, Morrissey and Smyth, 1999; McKittrick, 2007). Alongside county information the IED dataset also contains some limited street level data, which, is needed to perform a more refined spatial analysis using the DPMM. From this reduced dataset unique events are separated where unique is defined by the condition that no two entries have the same time and location. The use of “and” in this context is as a logical operator allowing events to have overlaps of time or location but excluding any cases where these features are both identical. In particular, there is

data covering 192 street level events in Belfast. Although this disaggregation reduces the number of available data points used from the original dataset it is also necessary to obtain more detailed results from applying the DPMM. It is also of importance to state that the dataset constructed is ordered temporally ranging from earliest to latest event. This arrangement of the data by increasing time is important for the novel implementation of the DPMM that will be introduced later in Section 4.4.

For data pertaining to the IED factory closures Asal et al. (2013) explain that “such data were difficult to collect in a systematic way”. However, Asal et al. did compile an unstructured dataset of news reports covering some of the IED factory discoveries that occurred in Northern Ireland. Focusing on Belfast it was possible to extract from the original unstructured dataset a series of 19 events involving IED factory identification. This subset of data had dates of identification on a daily scale recorded in the form year/month/day and street level addresses.

A further step that was taken before using the IED and factory datasets was to code their geographic information into latitude (lat)/longitude (long) pairs. To achieve this coding the Python module requests (Chandra and Varanasi, 2015) was used to make calls to OpenStreetMap (OSM) search engine Nominatim (www.nominatim.openstreetmap.org/) and extracted the returned lat/long coordinates. When the OSM search failed to return coordinates the same script would instead use the Google Maps Api (<https://maps.googleapis.com>).

In Tables 4.1 - 4.2 summaries of the datasets used in this chapter are provided

Table 4.1: PIRA IED Dataset

Field	Data
Date	Time of the recorded event in the form year(1970-1998)-month(1-12)-day(1-31)
Lat/Long	Latitude and longitude coordinates of street level events in the Belfast area

Table 4.2: Factory Closure Dataset

Field	Data
Date	Time of the recorded event in the form year(1970-1994)-month(1-12)-day(1-31)
Lat/Long	Latitude and longitude coordinates of street level closures in the Belfast area

4.2 Geographic Profiling

In the following sections the historical and mathematical development of the Dirichlet Process Mixture Model in the field of geographic profiling will be covered. This section also serves as a way to motivate the use of the DPMM as a method for studying the relationship between terrorist IED factories and counter-terrorism approaches aimed at closing them.

4.2.1 Criminal Geographic Profiling

The research field of geographic profiling and its application to studying criminal geographic targeting (CGT) stems from the idea that criminal actors are bound by

environmental factors. In particular, López (2005) describes four parameters which determine an offenders search space: operational range (furthest distance traveled by an offender), distance decay (the relationship between offender travel distance and number of crimes committed), a buffer region (offenders tend not to commit crimes too close to home so as to avoid incriminating themselves) and direction of travel (potential for an offender to prioritise a specific direction e.g. due to better area knowledge). These criminal considerations are then coupled with non-criminal nodes, such as, pathways and edges associated with daily activities (Brantingham and Brantingham, 1993b), for example, traveling to socialise or shopping.

A mathematical description of the theory underpinning criminal geographic profiling is given in the work of Rossmo (2000). In studying CGT Rossmo derives the formula shown in (4.1) to determine the likelihood, p_{ij} , that a point, (x_i, y_j) , inside a specific 2 dimensional region is the source location of observed criminal event sites $\{x_n, y_n\}$. To make the function more applicable to the study of criminal activity in the research of Rossmo the Manhattan metric is used which models the distance between the source of crimes and the criminal events on a grid network which the offender must traverse.

$$p_{ij} = k \sum_{n=1}^N \left[\frac{\phi}{(|x_i - x_n| + |y_j - y_n|)^f} + \frac{(1 - \phi)(B^{g-f})}{2B - (|x_i - x_n| + |y_j - y_n|)^g} \right] \quad (4.1)$$

with the conditions that

$$|x_i - x_n| + |y_j - y_n| > B \supset \phi = 1, \quad (4.2)$$

and

$$|x_i - x_n| + |y_j - y_n| \leq B \supset \phi = 0. \quad (4.3)$$

Here the value p_{ij} is known as the score function indicating the likelihood that a geographic location is the source of the observed events. For each possible source (x_i, y_j) the score function is calculated by the formula in (4.1). The function ϕ is used to define a buffer zone or region where the crimes are most likely located based on the ideas of distance to crime researched discussed earlier. The buffer zone is defined to have radius B . Inside the buffer zone the function $\phi = 0$ and outside $\phi = 1$. The constant value k is determined empirically. When summing over the crime site locations the total number of crimes is N . Finally the empirically determined parameters f and g control the decay rate around a potential source location. After applying the formula in (4.1) a 3 dimensional surface plot can be made illustrating the likelihood of source locations of the observed events.

4.2.2 Bayesian Framework of CGT

One of the major drawbacks of the CGT method proposed by Rossmo (2000) is that the model assumes a fixed formulation. In particular, the model utilises a specific decay function and distance metric. As a result of this fixed nature the equation in (4.1) has somewhat limited usage.

A method widely used in the fields of probability and statistics which allows for models to evolve based on other pieces of information is known as Bayes' rule. This rule is derived as follows (Ross, 2010).

First define the probability of event A given another event B as

$$P(A|B) = \frac{P(A \cap B)}{P(B)}. \quad (4.4)$$

Similarly

$$P(B|A) = \frac{P(A \cap B)}{P(A)}. \quad (4.5)$$

From the definitions in (4.4) and (4.5) it follows immediately that

$$P(A|B)P(B) = P(B|A)P(A), \quad (4.6)$$

which leads to Bayes' rule

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}. \quad (4.7)$$

The use of Bayes' rule in CGT began with the research of O'Leary (2009). Specifically, O'Leary formulated that the posterior distribution for a source location \mathbf{z} and an average distance to crime α given the observed location \mathbf{x} of a single criminal event is

$$P(\mathbf{z}, \alpha | \mathbf{x}) = \frac{P(\mathbf{x} | \mathbf{z}, \alpha) \pi(\mathbf{z}, \alpha)}{P(\mathbf{x})}. \quad (4.8)$$

Note that in (4.8) the source $\mathbf{z} = (z^{(1)}, z^{(2)})$ and crime site $\mathbf{x} = (x^{(1)}, x^{(2)})$ are two-dimensional vectors. In this chapter these vectors are lat/long coordinates.

Here the term $\pi(\mathbf{z}, \alpha)$ is called the prior distribution and represents the distribution of the potential source location \mathbf{z} and distance to crime α before information

about the crime location is known. O'Leary also assumes a situation in which the source location and distance parameter are independent thus factorising

$$\pi(\mathbf{z}, \alpha) = H(\mathbf{z})\pi(\alpha). \quad (4.9)$$

Now the H is the prior for the source location and π is the prior for the average distance traveled by the offender.

Since the marginal distribution $P(\mathbf{x})$ is independent of the \mathbf{z} and α this term can be ignored and hence the final formulation becomes

$$P(\mathbf{z}, \alpha | \mathbf{x}) \propto P(\mathbf{x} | \mathbf{z}, \alpha) H(\mathbf{z}) \pi(\alpha). \quad (4.10)$$

O'Leary (2009) also further extends the Bayesian CGT model to incorporate the observation of several crime site events given by the series $\{\mathbf{x}_i\}_{i=1}^N$. Following similar steps to the case of observing only a single crime site we obtain the formula

$$P(\mathbf{z}, \alpha | \mathbf{x}_1, \dots, \mathbf{x}_N) \propto P(\mathbf{x}_1, \dots, \mathbf{x}_N | \mathbf{z}, \alpha) H(\mathbf{z}) \pi(\alpha). \quad (4.11)$$

As before the term $P(\mathbf{z}, \alpha | \mathbf{x}_1, \dots, \mathbf{x}_N)$ is the posterior distribution of source, \mathbf{z} , and average distance, α , whilst $P(\mathbf{x}_1, \dots, \mathbf{x}_N | \mathbf{z}, \alpha)$ gives the joint probability of observing the locations $\{\mathbf{x}_i\}_{i=1}^N$ given the source and average distance.

Assuming that the criminal event sites are independent yields

$$P(\mathbf{x}_1, \dots, \mathbf{x}_N | \mathbf{z}, \alpha) = P(\mathbf{x}_1 | \mathbf{z}, \alpha) P(\mathbf{x}_2 | \mathbf{z}, \alpha) \dots P(\mathbf{x}_N | \mathbf{z}, \alpha). \quad (4.12)$$

Therefore we can write

$$P(\mathbf{z}, \alpha | \mathbf{x}_1, \dots, \mathbf{x}_N) \propto \left(\prod_{i=1}^N P(\mathbf{x}_i | \mathbf{z}, \alpha) \right) H(\mathbf{z}) \pi(\alpha), \quad (4.13)$$

where \prod is used to denote the product of the terms.

The final step in O’Leary’s formulation is to observe that since we are only interested in the source location, \mathbf{z} , for the crimes it is possible to take the conditional distribution by integrating as follows

$$P(\mathbf{z} | \mathbf{x}_1, \dots, \mathbf{x}_N) \propto \int \left(\prod_{i=1}^N P(\mathbf{x}_i | \mathbf{z}, \alpha) \right) H(\mathbf{z}) \pi(\alpha) d\alpha. \quad (4.14)$$

Hence to search for a likely source location it is necessary to find a geographic point with a high value of $P(\mathbf{z} | \mathbf{x}_1, \dots, \mathbf{x}_N)$ computed via (4.14).

At this point it is worth noting a major drawback that also occurs with this Bayesian approach. When introducing the method of Rossmo (2000) no assumptions were made concerning the number of potential sources. However with the Bayes approach described by O’Leary (2009) only one source location is considered. Hence, although allowing more flexibility to include new information into the model, the Bayes method is not suitable in this form for studies with potentially multiple sources. A goal, therefore, is to derive a model capable of combining more than one source location as in Rossmo (2000) with the adaptability of adding new information with the use of a Bayesian approach as in O’Leary (2009). This problem is solved using the Dirichlet Process Mixture Model.

4.3 Dirichlet Process Mixture Model

4.3.1 Model Theory

The problem of assigning event locations to more than one source location can be solved by appealing to Dirichlet Process Mixture Models (DPMM). In this section the theoretical foundations of the DPMM will be presented alongside practical aspects of how the model can be used in application.

Extensive research covering the mathematical foundation and properties of the DPMM already exist (Escobar and West, 1995; Ferguson, 1983; Green and Richardson, 2001). One property in particular which makes the DPMM such a useful model is the flexibility it affords concerning the possible number of source locations that can be found. Specifically the DPMM has the property of allowing a potentially infinite number of sources to exist. Therefore, when applying the DPMM to a problem where N events are observed all possible groupings of N events in up to N groups are considered with the most likely model being chosen.

The basic design of the DPMM and details of model inference are provided in the research of Neal (2000). Here an overview of the work of Neal will be provided. From this point the model will be established with the aim of applying the DPMM in a two dimensional discrete setting. Therefore, all crime event locations and source locations will be assumed to be points of the form $(\omega^{(1)}, \omega^{(2)}) \in \Omega \subset \mathbb{R}^2$.

Assume that the total number of observed crime sites is given by N and that these sites are denoted by the set $\{\mathbf{x}_i\}_{i=1}^N$. Moreover, assume that the infinite possible

source locations for the crimes are denoted by $\{\mathbf{z}_i\}_{i=1}^{\infty}$ or $\mathbf{z}_{1,\dots,\infty}$. It should be noted that in application the actual number of sources will be bounded by the number of events. However, to demonstrate the potential of the DPMM it was decided to present arguments with a potentially infinite number of sources. Prior to adding information concerning the observed event locations \mathbf{x}_i it is also assumed that the \mathbf{z}_i are drawn from some prior distribution denoted $G_0(\mathbf{z}_i)$. In the language of DPMM G_0 is known as the base distribution of the Dirichlet Process. Then each crime event location \mathbf{x}_i can be associated with a source via an index denoted c_i . In particular, it is said that the crime at \mathbf{x}_i originated from the source at \mathbf{z}_{c_i} . The point \mathbf{x}_i can be thought of as a draw from some distribution $F(\mathbf{z}_{c_i})$. Here the distribution F will be taken to be a 2 dimensional Gaussian centered on the location of \mathbf{z}_{c_i} . This assumption on the distribution is justified based on the theory of criminal geographic profiling where it is assumed that an offenders travel distance decays moving away from their home location (Brantingham and Brantingham, 1981; Liu and Eck, 2008; Rhodes and Conly, 1981).

Subsequently the c_i are considered to be drawn from a method known as the Chinese Restaurant Process (CRP). The CRP is associated with a concentration parameter α . If the model has determined that there exists a group A from which a total of N_A crime events are grouped then the probability that the next observed crime event also belongs to group A is given by

$$\frac{N_A}{N + \alpha} \tag{4.15}$$

where N_A is the number of elements already in A and N is the total number of

observed crime events observed thus far. Conversely the probability that the next observation stems from a group with a different source is given by

$$\frac{\alpha}{N + \alpha}. \quad (4.16)$$

Observe that under the CRP probabilities in (4.15) and (4.16) there is preference for groups with a higher number of events to grow larger than the emergence of a new group.

Considering all of the above assumptions the DP model can expressed as

$$\mathbf{x}_i | \mathbf{z}_{c_i} \sim \mathcal{F}(\mathbf{z}_{c_i}) \quad (4.17)$$

$$\mathbf{z}_{1, \dots, \infty} \sim \mathcal{G}_0 \quad (4.18)$$

$$c_i \sim \text{CRP}(\alpha). \quad (4.19)$$

With the DPMM as stated above the concentration parameter of the CRP must be estimated from the dataset being studied. However, this estimation procedure can be avoided using an extension of the DPMM made by Verity et al. (2014). In particular, a hyper-prior distribution can be placed on α denoted by H . In this case hyperprior simply means a prior distribution placed on α but it is termed a hyperprior in order to distinguish that the prior is over the parameter α of the CRP model. Therefore, the full DPMM can be described as follows

$$\mathbf{x}_i | \mathbf{z}_{c_i} \sim \mathcal{F}(\mathbf{z}_{c_i}) \quad (4.20)$$

$$\mathbf{z}_{1, \dots, \infty} \sim \mathcal{G}_0 \quad (4.21)$$

$$c_i \sim \text{CRP}(\alpha) \quad (4.22)$$

$$\alpha \sim \mathcal{H}. \quad (4.23)$$

The \mathcal{F} , \mathcal{G}_0 and \mathcal{H} each denote a type of distribution, such as, normal or exponential. On the other hand $F(\mathbf{x}_i|\mathbf{z}_{c_i})$, $G_0(\mathbf{z}_i)$ and $H(\alpha)$ represent probability mass/density functions.

4.3.2 DPMM Formulae

Having established a description of the DPMM in terms of distributions it is now possible to derive relevant formulae that can be used to compute useful information from the DPMM framework. The main application of the DPMM in this chapter is to cluster data. Therefore, consider a specific partition defined by the indices c_i and assign the indices to the vector $\mathbf{c} := (c_1, \dots, c_N)$. With this partitioning there is a corresponding value $N_j = |\{c_i : c_i = j\}|$ which denotes the number of elements in group j . Moreover, for a total number of groups $u \leq N$ we have that $\sum_{j=1}^u N_j = N$. In the work of Antoniak (1974) it was shown that the conditional probability

$$P(\mathbf{c}|\alpha) = \frac{\alpha^u \Gamma(\alpha)}{\Gamma(N + \alpha)} \prod_{j=1}^u \Gamma(N_j). \quad (4.24)$$

Here the function Γ is defined by (Artin and Butler, 2015)

$$\Gamma(\alpha) = \int_0^\infty s^{\alpha-1} e^{-s} ds. \quad (4.25)$$

As is common in Bayesian settings (Gelman et al., 2014) we can integrate over the α to obtain

$$P(\mathbf{c}) = \int P(\mathbf{c}|\alpha)P(\alpha)d\alpha = \int P(\mathbf{c}|\alpha)H(\alpha)d\alpha \quad (4.26)$$

where $P(\alpha) = H(\alpha)$ follows from the definition of the DPMM.

In the research of Verity et al. (2014) the authors define the function

$$t(u) = \int_0^\infty \frac{\alpha^u \Gamma(\alpha)}{\Gamma(N + \alpha)} H(\alpha) d\alpha. \quad (4.27)$$

Therefore, we obtain

$$P(\mathbf{c}) = t(u) \prod_{j=1}^u \Gamma(N_j). \quad (4.28)$$

Another quantity of interest that is easily obtained from the definition of the DPMM, and using Bayesian techniques, is the marginal probability for the vector of data points $\mathbf{x} := (\mathbf{x}_1, \dots, \mathbf{x}_N)$

$$P(\mathbf{x}|\mathbf{c}) = \prod_{j=1}^\infty \sum_{\boldsymbol{\omega} \in \Omega} G_0(\boldsymbol{\omega}) \prod_{i:c_i=j} F(\mathbf{x}_i|\boldsymbol{\omega}). \quad (4.29)$$

Observe that the first product in this formula considers the possibility of an infinite number of source locations. However, in application there will always be a finite number of sources. In this case only a subset of the indices j are associated with sources whilst other indices in the summation will correspond to a sum over the prior which is equal to 1.

The probabilities found in (4.28) and (4.29) can be combined using Bayes' rule

to obtain

$$P(\mathbf{c}|\mathbf{x}) = \frac{P(\mathbf{x}|\mathbf{c})P(\mathbf{c})}{\sum_{\mathbf{c} \in \mathcal{P}} P(\mathbf{x}|\mathbf{c})P(\mathbf{c})}, \quad (4.30)$$

where the set \mathcal{P} denotes all possible ways to partition N data points in at most N groups.

In the previous derivations formulae have been obtained relating to a particular partition defined by the vector \mathbf{c} . However, when performing inference techniques for the DPMM discussed later in this chapter it is also necessary to have a formula for the distribution of a single source location \mathbf{z}_j . Again using the distributions defined in the DPMM it holds that

$$P(\mathbf{z}_j|\mathbf{c}, \mathbf{x}) = \frac{G_0(\mathbf{z}_j) \prod_{i:c_i=j} F(\mathbf{x}_i|\mathbf{z}_j)}{\sum_{\boldsymbol{\omega} \in \Omega} G_0(\boldsymbol{\omega}) \prod_{i:c_i=j} F(\mathbf{x}_i|\boldsymbol{\omega})}. \quad (4.31)$$

So far in this exposition of the DPMM the necessary equations have been provided which allow us to compute the required probabilities for the model. These numerical values can contain interesting information but are harder to interpret raw than a visual representation. Moreover, if models like the DPMM are to be useful in real world settings by practitioners unfamiliar with the mathematical theory a simple to interpret representation of the model outcomes is essential. To this end the formulae derived above can now be combined in order to provide a method for calculating a jeopardy surface visually demonstrating the most likely sources for crime or terrorist attacks.

As the research of Verity et al. (2014) explains the problem of interest is to find the location of any source location \mathbf{z} . In probabilistic terms this problem can be

reduced approximately to finding the mean of the probabilities $P(\mathbf{z}_j|\mathbf{c}, \mathbf{x})$ (exactly in the case of a continuous setting with the use of probability density functions instead of probability mass functions). This mean is computed as

$$\mathcal{J}(\mathbf{z}|\mathbf{c}, \mathbf{x}) = \frac{1}{u} \sum_{j=1}^u P(\mathbf{z}_j|\mathbf{c}, \mathbf{x}), \quad (4.32)$$

where \mathcal{J} represents the jeopardy surface of source locations conditioned on a particular partition \mathbf{c} and event locations \mathbf{x} .

However, we have from (4.30) a distribution for partitions given events. Therefore, a full jeopardy surface for source locations given event locations can be obtained via the formula

$$\mathcal{J}(\mathbf{z}|\mathbf{x}) = \sum_{\mathbf{c} \in \mathcal{P}} \mathcal{J}(\mathbf{z}|\mathbf{c}, \mathbf{x}) P(\mathbf{c}|\mathbf{x}). \quad (4.33)$$

4.3.3 Markov Chain Monte Carlo Method of Inference

The fundamental idea underpinning the DPMM is to consider each point of a dataset in turn and compute the likelihood that the point should be included in a previously found cluster or added to a new cluster. This process is then repeated iteratively through the dataset until the most likely combination of the datapoints into clusters is discovered. As demonstrated in Section 4.3.2 for a dataset consisting of N entries the problem requires a summation over N points into at most N clusters. This type of problem has been considered in the field of combinatorial mathematics where Bell numbers are used to describe the total possible partitions of a set (Bell, 1934; Bell, 1938). As pointed to by Verity et al. (2014) even for very small datasets with 10

data entries the Bell number is large, $B_{10} = 115975$. Therefore, adopting a brute force approach by directly trying to apply the DPMM on even small datasets is computationally expensive. Hence it is necessary to use a probabilistic sampling method to infer the best clustering of a given set of data.

When performing statistical inference a well studied and highly effective method to minimise problems of expensive computational time is to appeal to a class of algorithms known as Markov Chain Monte Carlo (MCMC) methods (Berg, 2004; Gamerman, 1997; Gilks, Richardson and Spiegelhalter, 1996). MCMC algorithms are used to reconstruct a probability distribution. The algorithm starts from some initial value and constructs a new value, for example by adding a random number, and then compares the current and new value using some chosen criteria in order to determine whether to keep the current value or move to the new value. Since this step only requires the current and new values, i.e. no past dependence, a Markov chain is being constructed. This method must be repeated until a required level of convergence is achieved. The latter step is the Monte Carlo method. By carefully choosing the comparison technique at each stage of the algorithm it can be proven that this method will converge on the original probability distribution.

4.3.4 Gibbs Sampling for DPMM

A specific type of MCMC algorithm which has been used for DPMM inference is known as a Gibbs sampler. In the research of Neal (2000) a full account of the application of Gibbs sampling to DPMM is provided. Here an overview of the work of Neal is presented.

Recall that the index c_i links the location of event \mathbf{x}_i to a source \mathbf{z}_{c_i} . Assume that on the current step of the Gibbs sampling method a draw of value k is made for c_i . Then let $c_{-i} := \{c_j | j \neq i\}$ be the subset of indices excluding c_i . From this subset make a further definition of the subset $c_{-i,k} := \{c_j | j \neq i, j = k\}$ of index c_k with $k \neq i$. Corresponding to $c_{-i,k}$ we have the value $N_{-i,k} = |c_{-i,k}|$ which represents the number of elements of the subset $c_{-i,k}$. In this case formula 3.6 from Neal (2000) implies that given the most recent source locations we have

if $k \in c_{-i}$

$$P(c_i = k | c_{-i}, \mathbf{x}_i, \mathbf{z}_{1,\dots,\infty}, \alpha) = b \frac{N_{-i,k}}{N - 1 + \alpha} F(\mathbf{x}_i | \mathbf{z}_k) \quad (4.34)$$

else if $k \notin c_{-i}$

$$P(c_i = k | c_{-i}, \mathbf{x}_i, \mathbf{z}_{1,\dots,\infty}, \alpha) = b \frac{\alpha}{N - 1 + \alpha} \sum_{\boldsymbol{\omega} \in \Omega} F(\mathbf{x}_i | \boldsymbol{\omega}) G_0(\boldsymbol{\omega}) \quad (4.35)$$

for some normalising constant b . These formulas hold for $i = 1, \dots, N$.

In terms of the CRP (4.34) states that the probability of placing \mathbf{x}_i in an already occupied cluster is proportional to the number $N_{-i,k}$ of other data points associated to that cluster. On the other hand (4.35) states that there exists a probability proportional to α of \mathbf{x}_i being assigned to a new cluster. After finding the most likely of these two situations for each i it is then possible to draw a new set of \mathbf{z}_{c_i} based on the formula derived in (4.31).

As discussed by Verity et al. (2014) the formulations in (4.34) and (4.35) can be

simplified by removing the dependence on α via integration. In particular, using the function t defined in (4.27) Verity et al. (2014) derive the formula

if $k \in c_{-i}$

$$P(c_i = k | c_{-i}, \mathbf{x}_i, \mathbf{z}_1, \dots, \infty) = b' t(u_{-i}) N_{-i,k} F(\mathbf{x}_i | \mathbf{z}_k) \quad (4.36)$$

else if $k \notin c_{-i}$

$$P(c_i = k | c_{-i}, \mathbf{x}_i, \mathbf{z}_1, \dots, \infty) = b' t(u_{-i} + 1) \sum_{\boldsymbol{\omega} \in \Omega} F(\mathbf{x}_i | \boldsymbol{\omega}) G_0(\boldsymbol{\omega}) \quad (4.37)$$

with u_{-i} denoting the number of unique elements of the set c_{-i} . The value b' is used to indicate a new constant of proportionality.

The last step in deriving the formulae for the Gibbs sampling method is to move from the discrete setting into the case of continuous space. In this case the functions $F(\cdot)$ and $G_0(\cdot)$ are probability density functions (pdfs). To simplify the calculations needed for the DPMM these pdfs are assumed to be conjugate. Conjugate means that the density functions $F(\cdot)$ and $G_0(\cdot)$ belong to the same family (Gelman et al., 2014). A common example of conjugation is when the pdfs take the form of an exponential. In this case the relationship between posterior and prior density functions corresponds, according to Bayes' theorem, to exponential multiplication which results in arithmetical changes to the exponent. A well studied example of an exponential family are Gaussian distributions. As discussed earlier in this chapter the 2 dimensional multivariate Gaussian is employed.

We can now describe the analogous probabilities to (4.36) and (4.37), under the assumption of continuous space and conjugate pdfs

if $k \in c_{-i}$

$$P(c_i = k | c_{-i}, \mathbf{x}_i) = b'' t(u_{-i}) N_{-i,k} \int_{\Omega} F(\mathbf{x}_i | \boldsymbol{\omega}) dy_{-i,k}(\boldsymbol{\omega}) \quad (4.38)$$

else if $k \notin c_{-i}$

$$P(c_i = k | c_{-i}, \mathbf{x}_i) = b'' t(u_{-i} + 1) \int_{\Omega} F(\mathbf{x}_i | \boldsymbol{\omega}) dG_0(\boldsymbol{\omega}). \quad (4.39)$$

The term $y_{-i,k}(\boldsymbol{\omega})$ represents a posterior distribution over the $\boldsymbol{\omega}$ which stems from the prior G_0 and the observed data points \mathbf{x}_j where $j \neq i$ and $c_j = k$. As in the previous derivations a new constant of proportionality b'' has been introduced. It should be noted however that none of the proportionality constants actually require numerical evaluation since we only need to sample from the distributions and compare their relative likelihoods when using the DPMM for cluster analysis.

4.3.5 DPMM Parameter Fitting

A major assumption that has been made throughout the current exposition of the DPMM is that at each step of the MCMC algorithm only the clustering of data points needs to be computed. This neglects the fact that the probability distributions being used to determine the clusters have parameter values which require initialisation. Specifically, under the model used here where the risk surface consists of multivariate

distributions over the source locations there are two parameters - the mean and the standard deviation. In the description of the group allocation method described in Section 4.3.4 it can be seen that the mean is actually integrated out of the model to improve efficiency. However, the standard deviation has to be user-defined. One method described by Stevenson (2013) is to set the standard deviation manually by consulting a histogram of pairwise distances between the event data points. But this clearly introduces subjectivity into the problem. Therefore, Stevenson (2013) proposes incorporating the method now described into the MCMC algorithm which results in an adaptive model-derived estimation of the standard deviation parameter.

On the level of distributions the adaptive estimation method proposed by Stevenson (2013) is manifested in the following way

$$\mathbf{X}|\boldsymbol{\mu}, \sigma \sim \mathcal{N}(\boldsymbol{\mu}, \sigma), \quad (4.40)$$

$$\boldsymbol{\mu} \sim \mathcal{N}(\tilde{\boldsymbol{\mu}}, \sigma), \quad (4.41)$$

$$\sigma \sim \text{IG}(\delta, \beta), \quad (4.42)$$

where $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_N)$ is the vector of observed event data, \mathcal{N} denotes a normal distribution and IG is an inverse gamma distribution.

In the case of a one dimensional example a random variable x_i which has an $\mathcal{N}(\mu, \sigma^2)$ distribution has a probability density function given by (Ross, 2010)

$$f(x_i) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left\{-\frac{(x_i - \mu)^2}{2\sigma^2}\right\}. \quad (4.43)$$

When analysing multidimensional data, as in this chapter, this one dimensional

density can be generalised. For a random vector of N dimensions, denoted \mathbf{X} , with a multivariate normal distribution having parameters $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ the probability density function is given by the formula (Ross, 2010)

$$f(\mathbf{X}) = \frac{1}{(2\pi)^{\frac{N}{2}} |\boldsymbol{\Sigma}|^{\frac{1}{2}}} \exp \left\{ -\frac{1}{2} (\mathbf{X} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{X} - \boldsymbol{\mu}) \right\}, \quad (4.44)$$

where, $(\mathbf{X} - \boldsymbol{\mu})^T$ denotes the usual vector transpose, which, takes a column vector to a row vector.

In application to the DPMM the source locations are assumed to correspond to the same standard deviation σ with this value assumed to hold in both the x- and y-axes dimensions for each data point. Here the value of σ represents knowledge about the dispersal distance expected around the sources. A justification for this assumption in the context of PIRA IED studies can be found in the research of Horgan et al. (2013). In particular, the research group of Horgan et al found that 63.4% of PIRA members travelled less than 4 miles to commit their attacks.

A final assumption needed to simplify the model is that the dispersal patterns along the x- and y-axes are independent. This simplifying assumption removes any matrix multiplications arising from covariance and minimises the number of parameters to be fitted. Furthermore, under this independence assumption it is only necessary to derive formulae in one dimension with the formulae in the second dimension being the same with a simple change in variable name.

Here we will refer to the spatial x-components via \mathbf{X}_x and the y-components as \mathbf{X}_y . The formulae derived using \mathbf{X}_x will be the same for \mathbf{X}_y except for a change of

variable.

In each group update step of the MCMC algorithm the density function modelling the event data is dependent on an $\mathcal{N}(\mu_x, \sigma^2)$ distribution. Incorporating this information into (4.44) yields the following density

$$f(\mathbf{X}_x | \mu_x, \sigma^2) = \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \exp \left\{ -\frac{1}{2\sigma^2} \sum_{i=1}^N (x_i - \mu_x)^2 \right\}. \quad (4.45)$$

For the density in (4.45) to be useful it is also necessary to draw values for the mean and variance at each update step. To achieve this objective it is assumed that each source location is drawn from a prior distribution modelled as a bivariate normal. For this source distribution we follow the same labelling notation as in Stevenson (2013) denoting the means along each spatial dimension via the pair (d_x, d_y) and the corresponding variances (T_x^2, T_y^2) . With these definitions for the prior on source locations it is now possible to obtain the following likelihood which models the joint conditional distribution of event data \mathbf{X}_x and mean μ_x along the x-axis spatial direction

$$f(\mathbf{X}_x, \mu_x | d_x, \sigma^2, T_x^2) = \frac{1}{(2\pi\sigma^2)^{\frac{N}{2}}} \frac{1}{\sqrt{(2\pi T_x^2)}} \exp \left\{ -\frac{1}{2} \sum_{i=1}^N \frac{(x_i - \mu_x)^2}{\sigma^2} + \frac{(\mu_x - d_x)^2}{T_x^2} \right\}. \quad (4.46)$$

This formula is simply the multiplication of the function in (4.45) by the prior over the source locations.

Observe that the formula in (4.46) contains an exponential function which can

be simplified. In particular, this formula can be written in the form

$$f(\mu_x | \mathbf{X}_x, d_x, \sigma^2, T_x^2) = \frac{1}{\sqrt{(2\pi\epsilon_x^2)}} \exp \left\{ -\frac{(\mu_x - \theta_x)^2}{2\epsilon_x^2} \right\}. \quad (4.47)$$

To derive this formula observe that the exponent on the right hand side of (4.47) takes the form

$$\frac{\mu_x^2 - 2\mu_x\theta_x + \theta_x^2}{\epsilon_x^2} \quad (4.48)$$

whilst the exponent in (4.46) can be expressed as

$$\mu_x^2 \left(\frac{N}{\sigma^2} + \frac{1}{T_x^2} \right) - 2\mu_x \left(\frac{\sum x_i}{\sigma^2} + \frac{d_x}{T_x^2} \right) + \left(\frac{\sum x_i^2}{\sigma^2} + \frac{d_x^2}{T_x^2} \right). \quad (4.49)$$

Comparing the expansions in (4.48) and (4.49) yields the following identities

$$\epsilon_x^2 = \frac{1}{\frac{N}{\sigma^2} + \frac{1}{T_x^2}} \quad \theta_x = \frac{\frac{\sum x_i}{\sigma^2} + \frac{d_x}{T_x^2}}{\frac{N}{\sigma^2} + \frac{1}{T_x^2}}. \quad (4.50)$$

The significance of the identities derived in (4.50) is that the density from which μ_x is drawn can be obtained using the distribution

$$\mu_x \sim \mathcal{N} \left(\frac{\frac{\sum x_i}{\sigma^2} + \frac{d_x}{T_x^2}}{\frac{N}{\sigma^2} + \frac{1}{T_x^2}}, \frac{1}{\frac{N}{\sigma^2} + \frac{1}{T_x^2}} \right). \quad (4.51)$$

Since we made the simplifying assumption earlier that the dispersal along each spatial axis is independent the same formula in (4.51) holds for the y-dimension after a

change of notation

$$\mu_y \sim \mathcal{N} \left(\frac{\frac{\sum y_i}{\sigma^2} + \frac{d_y}{T_y^2}}{\frac{N}{\sigma^2} + \frac{1}{T_y^2}}, \frac{1}{\frac{N}{\sigma^2} + \frac{1}{T_y^2}} \right). \quad (4.52)$$

Recall that $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_N)$ denotes the vector of event data points, where, each $\mathbf{x}_i = (x_i, y_i)$ is a two dimensional point. Each run of the MCMC algorithm assigns to each data point \mathbf{x}_i a group via the index c_i as outlined in the previous Section 4.3.4. Moreover, each group defined by a set of the form $\{\mathbf{x}_i | c_i = j\}$ has a corresponding cardinality N_j . Let the pair (μ_x^j, μ_y^j) represent the x and y axes means for the event data assigned to the jth group. From (4.45) we have the following density formulas

$$f(x_i : c_i = j | \mu_x^j, \sigma^2) = \frac{1}{(2\pi\sigma^2)^{\frac{N_j}{2}}} \exp \left\{ - \sum_{i:c_i=j} \frac{(x_i - \mu_x^j)^2}{2\sigma^2} \right\}, \quad (4.53)$$

$$f(y_i : c_i = j | \mu_y^j, \sigma^2) = \frac{1}{(2\pi\sigma^2)^{\frac{N_j}{2}}} \exp \left\{ - \sum_{i:c_i=j} \frac{(y_i - \mu_y^j)^2}{2\sigma^2} \right\}. \quad (4.54)$$

Assuming that the number of groups found on the current step of the MCMC is k then by combining the density functions for each of the k groups via multiplication the formula for the complete density function describing the event data can be obtained via

$$f(\mathbf{X}_x | \mu_x^1, \dots, \mu_x^k, \sigma^2) = \frac{1}{(2\pi\sigma^2)^{\frac{N_1}{2}}} \cdots \frac{1}{(2\pi\sigma^2)^{\frac{N_k}{2}}} \exp \left\{ - \sum_{i:c_i=1} \frac{(x_i - \mu_x^1)^2}{2\sigma^2} \cdots - \sum_{i:c_i=k} \frac{(x_i - \mu_x^k)^2}{2\sigma^2} \right\} \quad (4.55)$$

and

$$f(\mathbf{X}_y | \mu_y^1, \dots, \mu_y^k, \sigma^2) = \frac{1}{(2\pi\sigma^2)^{\frac{N_1}{2}}} \cdots \frac{1}{(2\pi\sigma^2)^{\frac{N_k}{2}}} \exp \left\{ - \sum_{i:c_i=1} \frac{(y_i - \mu_y^1)^2}{2\sigma^2} \cdots - \sum_{i:c_i=k} \frac{(y_i - \mu_y^k)^2}{2\sigma^2} \right\}. \quad (4.56)$$

The final parameter that still remains undetermined in the equations derived thus far is the variance σ^2 . To draw this parameter an appeal is made to the theory of conjugate priors. A conjugate prior/posterior relationship simply means that the prior and posterior distribution are from the same family of distributions (Raiffa and Schlaifer, 2000). This is usually the case with densities which utilise exponential functions due to the additive property of the exponents.

It is a standard result that the conjugate prior distribution of a normal distribution with known mean and unknown variance is the inverse gamma distribution (Gelfand et al., 1990). In the context of the present study of the DPMM an inverse gamma (IG) distribution of the following form is fitted to the variance parameter σ^2

$$\sigma^2 \sim IG(\delta, \beta), \quad (4.57)$$

$$f(\sigma^2 | \delta, \beta) = \frac{\beta^\delta}{\Gamma(\delta)} \left(\frac{1}{\sigma^2} \right)^{\delta+1} \exp \left\{ - \frac{\beta}{\sigma^2} \right\},$$

where Γ is the gamma function as defined in Section 4.3.2 as $\Gamma(\alpha) = \int_0^\infty s^{\alpha-1} e^{-s} ds$.

Multiplying the densities along each axis given by (4.55) and (4.56) and the inverse gamma prior over the variance given by (4.57) yields the equation

$$f(\mathbf{X}, \sigma^2 | \mu_x^1, \dots, \mu_x^k, \mu_y^1, \dots, \mu_y^k) = \frac{\beta^\delta}{\Gamma(\delta)} \left(\frac{1}{\sigma^2}\right)^{\delta+1+N} \frac{1}{(2\pi)^N} \exp\left\{-\frac{1}{\sigma^2} \left[\beta + \sum_{j=1}^k \sum_{i:c_i=j} \frac{[(x_i - \mu_x^j)^2 + (y_i - \mu_y^j)^2]}{2}\right]\right\}. \quad (4.58)$$

Observe that the distribution in (4.58) is again an inverse gamma distribution. This was the motivating reason to select the IG conjugate prior to the normal distribution. As a result of this observation draws for the variance can now be constructed using the distribution defined by

$$\sigma^2 \sim IG\left(\delta + N + 1, \left[\beta + \sum_{j=1}^k \sum_{i:c_i=j} \frac{[(x_i - \mu_x^j)^2 + (y_i - \mu_y^j)^2]}{2}\right]\right). \quad (4.59)$$

An important remark concerning the result derived in (4.59) is that the inverse gamma distribution requires the specification of the parameters δ and β . A simple approach that can be taken to solve this problem is to set the parameters so that the inverse gamma distribution is uninformative as discussed in Stevenson (2013). However, in crime science much research has been conducted to numerically quantify distance-to-crime relationships so it would be preferable to set the inverse gamma parameters more precisely based on prior knowledge. In the context of the study of PIRA IED attacks research quantifying the distance travelled by PIRA members to conduct their attacks has been undertaken by Horgan et al. (2013). Specifically the

research of Horgan et al found that 63.4% of PIRA attackers travelled less than 4 miles to conduct attacks.

To adapt the model in (4.59) to use prior information about expected dispersal distance for the data studied, Faulkner et al. (2017) and Stevenson (2013) discuss a change of space from a prior over σ^2 to a one over σ . Firstly we use the formula discussed in the work of Glen (2011) which shows that the expected value for an inverse gamma distributed random variable X^r is given by the formula

$$\epsilon = \mathbb{E}(X^r) = \beta^r \frac{\Gamma(\delta - r)}{\Gamma(\delta)} \quad (4.60)$$

for $\delta > r$. Therefore, for a fixed expected value of X the following relationship exists

$$\beta = \left(\epsilon \frac{\Gamma(\delta)}{\Gamma(\delta - r)} \right)^{\frac{1}{r}}. \quad (4.61)$$

In addition to the expected value formula in (4.61) a transformation formula is also required. As found in Tanizaki (2004) if for random variables X and Y there is a one-to-one function ψ such that $Y = \psi^{-1}(X)$ then the following formula holds

$$f_Y(y) = |\psi'(y)| f_X(\psi(y)), \quad (4.62)$$

where, f_X and f_Y are the density functions for the random variables X and Y , respectively, and ψ' is the first derivative of ψ .

With the expectation formula from (4.61) and setting $\psi(y) = y^2$ in the change-of-variable formula from (4.62) it follows that the prior on the standard deviation is

given via the formula

$$f(\sigma|\delta, \epsilon) = \frac{2\beta^\delta}{\Gamma(\delta)} \left(\frac{1}{\sigma}\right)^{2\delta+1} \exp\left\{-\left(\frac{\epsilon\Gamma(\delta)}{\sigma\Gamma(\delta - \frac{1}{2})}\right)^2\right\}. \quad (4.63)$$

Here, the value of ϵ is interpretable as the expected dispersal distance for the data being studied. Observe that a user-specified value for δ is also required. As explained by Stevenson (2013) this parameter is less important than the value of ϵ and can be set to provide a diffuse fat-tailed distribution which carries little prior information. The value of δ needs to satisfy the condition that it is > 0.5 so that the model in (4.63) is well-defined. We follow the same convention as used in Stevenson (2013) and have the parameter $\delta = 1$.

On each iteration of the MCMC Gibbs sampler the standard deviation parameter σ can be now be updated so that the parameter's value is fitted automatically. Notice that in this section to obtain the prior on the standard deviation the group means for the multivariate normals used in the DPMM had to be realised. This is in comparison to the previous Section 4.3.4 where these means values were integrated out of the model. In a similar way the group allocations can be found on each iteration of the MCMC method by integrating out the means. However, prior to performing the update step of the standard deviation in (4.63) the mean values must be realised following the steps in this section.

In the research of Faulkner et al. (2017) and Stevenson (2013) the MCMC method for DPMM fitting was coded into a piece of R (Team, 2017) software. The R code is maintained on GitHub at <https://github.com/stevenlecomber/Rgeoprofile-1.1.0>. The analyses in the remainder of this chapter will be based on this R software which

will be adapted to include a study of temporal dynamics.

4.4 DPMM Analysis of PIRA IED Factories

In the previous sections of this chapter the mathematical foundation and implementation methods for using the DPMM have been introduced. Thus far the focus has been on utilising the DPMM to find groupings of datapoints based on their spatial configuration. Although this purely spatial model has proven useful in a number of research topics (Faulkner et al., 2015; Faulkner et al., 2017; Verity et al., 2014) there is scope for improving the model via the inclusion of temporal dynamics. In particular, temporal data can enrich insights from spatial patterns by providing a natural ordering for the data being studied thus allowing for a more detailed description of the underlying phenomena to be discovered.

As outlined already the outcome of applying the DPMM is a jeopardy surface illustrating the most likely source locations of a set of events based on grouping of their spatial coordinates. Here the objective is to look at this jeopardy surface for IED attacks committed by PIRA in Belfast during The Troubles in Northern Ireland. However, instead of studying IED attacks in isolation we will also consider how the spatial patterns evolved over time in relation to the closing of IED factories by British Security Forces.

Two datasets were required to perform the analyses in this chapter. One dataset covered the dates and locations of PIRA IED attacks whilst the second dataset provided the same information but for factory closures. The method devised to

incorporate events from both datasets was two-fold. Firstly, new datasets of IED attacks were constructed by extracting all events which occurred prior to the date of each recorded IED factory discovery. In this case the factory closure dates are employed as temporal boundaries for the entire time span of IED attacks. This generates one dataset corresponding to each factory. Secondly, the DPMM model is fitted to each of these datasets using the Gibbs method. This procedure yields a set of jeopardy surfaces delineating the timescale of PIRA attacks according to times when factory closures arose. By comparing and contrasting the final set of jeopardy surfaces the aim is to analyse the space-time impact of disrupting IED manufacturing on subsequent IED attacks.

4.5 Analysis

4.5.1 Model Calibration

Before implementing the spatio-temporal version of the DPMM one important model parameter requires calibration. In particular, in Section 4.3.5 a prior distribution was found to enable the MCMC procedure for the DPMM to automatically update the standard deviation used to model dispersion in the data. As part of the derivation of the formula describing this prior distribution (4.63) it was stated that the user must specify an initial expected value for the dispersal distance associated to the data being studied. In the context of PIRA it was found in the research of Horgan et al. (2013) that 63.4% of PIRA members travelled less than 4 miles to conduct attacks. Therefore, we shall take this value as an initialisation for the dispersal distance.

As well as defining an initial parameter value for the DPMM in the R code version of the model authored by Faulkner et al. (2017) and Stevenson (2013), and utilised in this chapter, the dispersal distance is required to be in units of latitude decimal degrees. To perform the conversion from miles a two step procedure was performed. Firstly from the R package “Measurements” by Birk (2016) the `conv_unit` function was employed to find the conversion of 1 mile to 1609.344 metres. Then from the R package “Geosphere” by Hijmans, Williams and Vennes (2016) the function `dism` was used to find the conversion between 1 latitude decimal degree and metres. To define 1 lat of distance in the input to `dism` the coordinates of the first factory location and another fictional location 1 latitude degree decimal away $\{(54.5961886, -5.9858976), (55.5961886, -5.9858976)\}$ was used. This resulted in a conversion of 1 latitude decimal degree to 110716.5 metres. It should also be noted that whilst utilising the `dism` function the method for computing distance was set to “`distVincentyEllipsoid`” which uses an ellipsoidal model for highly accurate and computationally efficient calculations (Vincenty, 1975). Moreover, the “`distVincentyEllipsoid`” method calculates “great-circle-distance” between inputs thus handling issues arising from the Earth’s curvature. Finally these two conversions were combined to find the distance conversion of 4 miles as $4 \left(\frac{1609.344}{110716.5} \right) = 0.05814285$ latitude decimal degrees, which, was the initialising value for equation (4.63).

4.5.2 Results

The results of applying the DPMM according to the steps outlined in the previous sections are now presented. In Table 4.3 information is presented which illustrates

the composition of the final datasets after using the IED factory closure dates as temporal boundaries. It should be observed from this table that there are a total of 16 datasets. This is in comparison to the value of 19 factory closures discussed in Section 4.1. Of the 19 factory closures 10, 12, 15 and 16 were found to contain no more additional IED events compared to the previous closure. Therefore, the influence of these factory closures are considered in datasets 11, 13 and 17. To avoid confusion the final datasets will be labelled 1-15. A final dataset denoted by 16 covers all IED attacks which appeared in the dataset before and after the final recorded factory closure. The final IED event in Belfast was recorded on 24/05/1998 whilst the final recorded factory find was 14/03/1994. A decision to have this additional dataset was taken to ensure that the impact of the last recorded factory closure could be analysed.

In Figures 4.1(a) - 4.1(o) and Figure 4.2 plots of the DPMM jeopardy surfaces are presented. The IED event locations are shown as red circles, the most recent factory closure is highlighted via a black triangle and the previous factory closures are displayed as blue squares. It should be noted that in these figures the map boundaries are determined from the IED locations. Therefore, when factory locations fall outside these boundaries they are not shown. The definition of the plot area is made in this way to ensure that the final model represents the best fit to the event data rather than imposing preconceived ideas about the source locations. Note that Figure 4.2 has no black triangle since this map is used to show the remaining IED events which occurred after the final factory closure in Figure 4.1(o). An overview of the datasets studied in each figure appears in Table 4.3.

In the plots of Figures 4.1(a) - 4.1(o) and Figure 4.2 likely source locations are determined via their hitscore. The hitscore is calculated by ranking the surface generated by the multivariate normals used in the DPMM and dividing by the size of the total search area. Each grouping of events found using the DPMM has an associated hitscore which can be interpreted as the percentage of space that must be searched to locate the source location. Therefore, the hitscore provides a measure of fit of the DPMM. A lower hitscore is associated with a better model fit and vice versa for a higher hitscore. As a benchmark value a 0.5 hitscore would be indicative of a random search over the given area. Graphically, a heatmap of the hitscores is presented in the plots of Figures 4.1(a) - 4.1(o) and Figure 4.2 with the lowest hitscores in white and the highest in dark red with shades of a white-yellow-dark red scale moving between the two extremes. Also in Table 4.3 the hitscore associated to the new factory closure being considered in each dataset is presented.

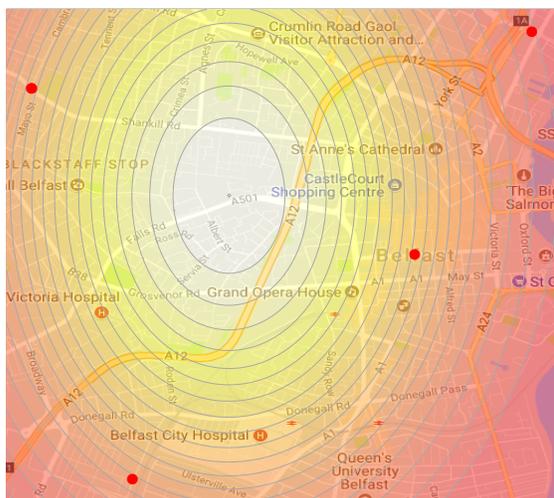
Table 4.3: Datasets Analysed using the DPMM

Factory Closure	Closure Date	Number of Data Points	Hitscore
1	1 / July / 1970	5	0.70
2	24 / November / 1972	17	0.90
3	11 / May / 1974	19	0.15
4	8 / March / 1979	61	0.28
5	11 / August / 1981	67	0.41
6	28 / August / 1982	72	0.95
7	21 / March / 1988	88	0.46
8	22 / August / 1988	89	0.17
9	7 / November / 1988	98	0.23
10	28 / April / 1990	135	0.40
11	18 / May / 1991	145	0.09
12	13 / January / 1992	158	0.85
13	15 / June / 1992	162	0.06
14	18 / August / 1993	172	0.49
15	14 / March / 1994	175	0.03
16	NA / NA / NA	192	NA

Jeopardy Surface Legend

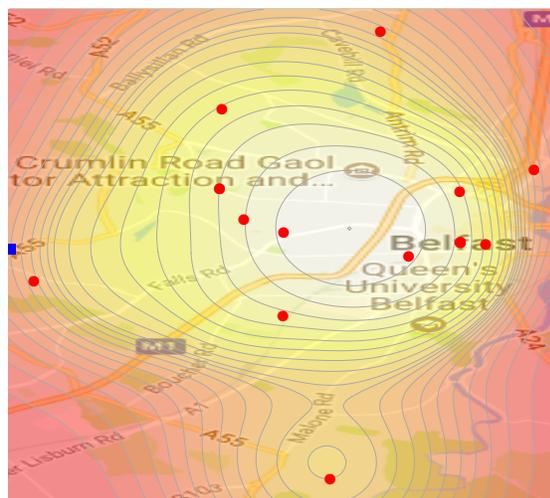
● IED Attack ■ Previous Factory ▲ New Factory ▲ Orientation

Phase 1, Factory 1



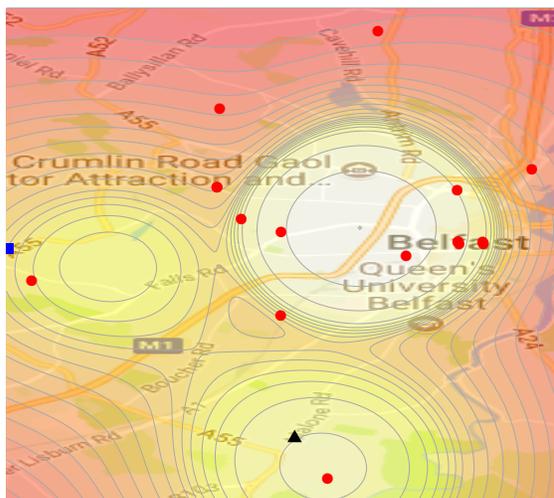
(a)

Phase 1, Factory 2



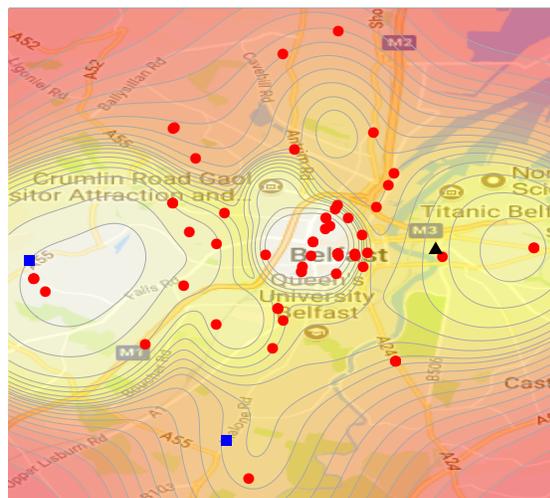
(b)

Phase 1, Factory 3



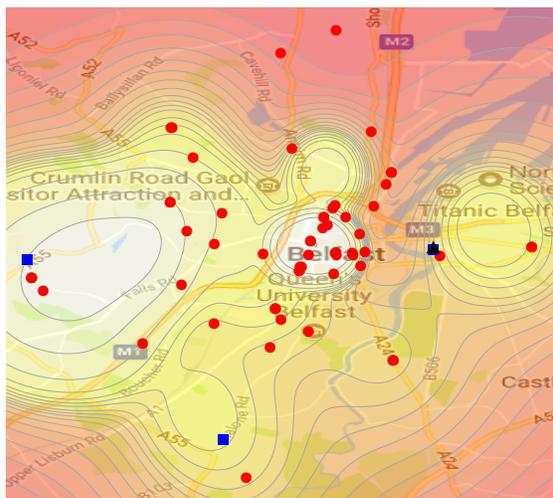
(c)

Phase 2, Factory 4



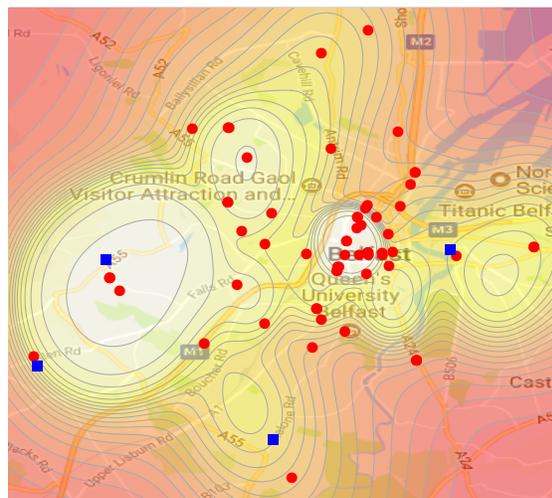
(d)

Phase 3, Factory 5



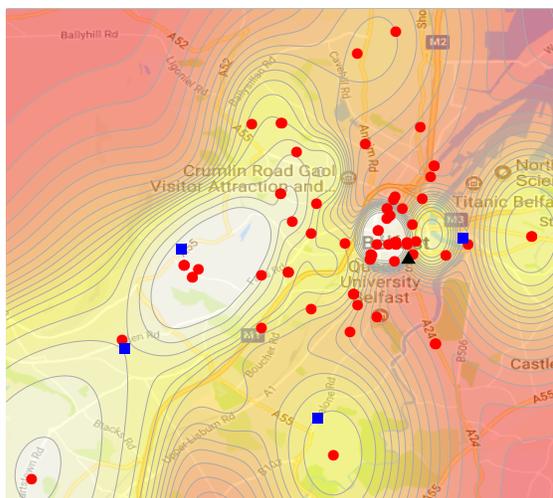
(e)

Phase 3, Factory 6



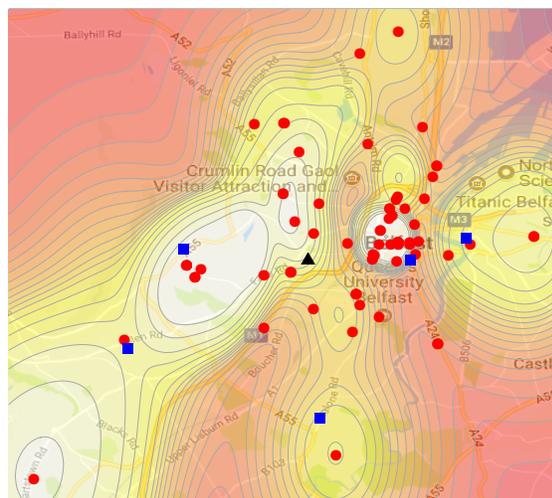
(f)

Phase 3, Factory 7



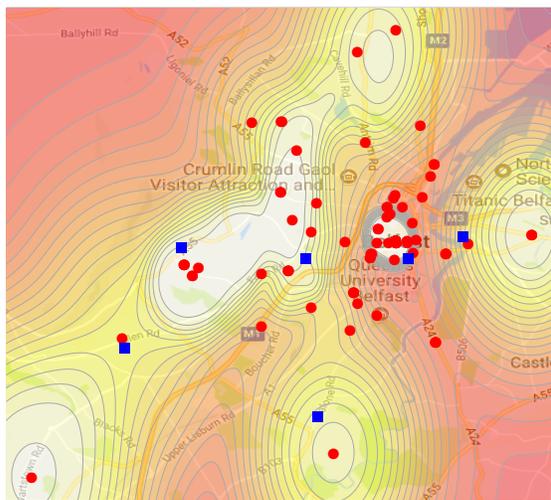
(g)

Phase 3, Factory 8



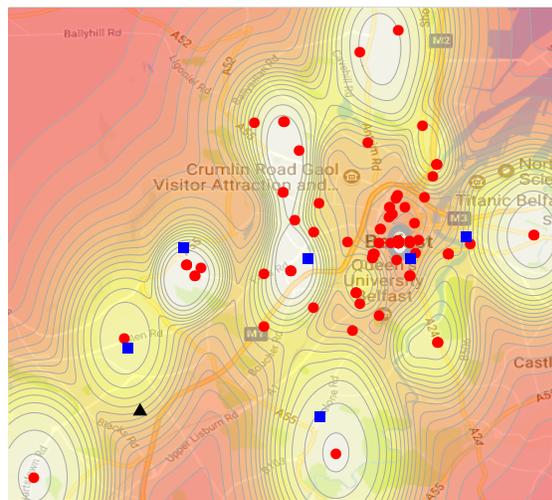
(h)

Phase 3, Factory 9



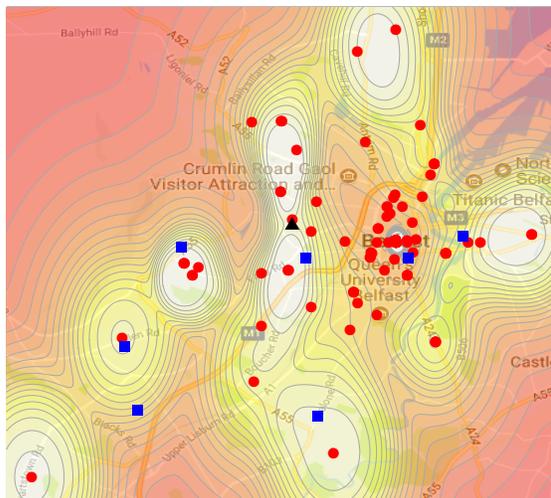
(i)

Phase 4, Factory 10



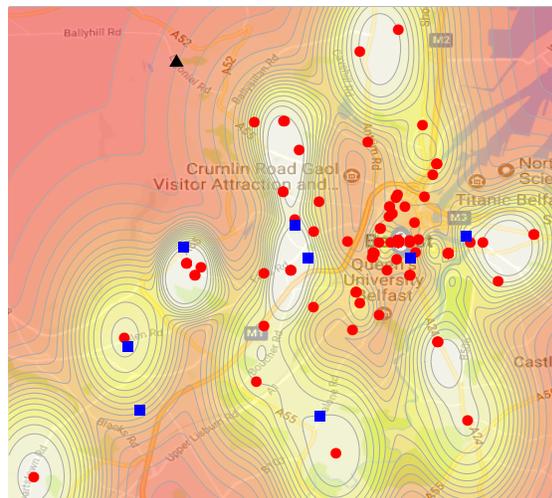
(j)

Phase 4, Factory 11



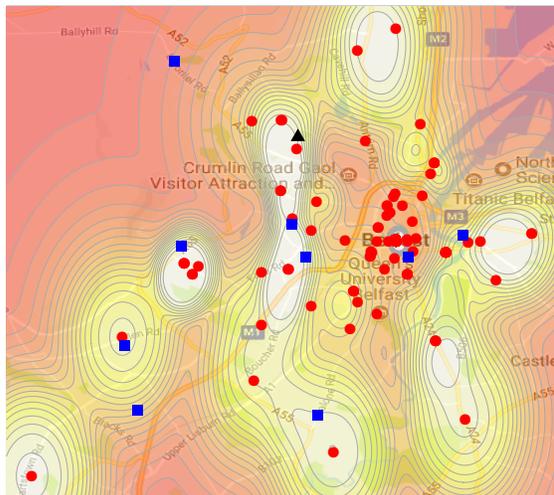
(k)

Phase 4, Factory 12



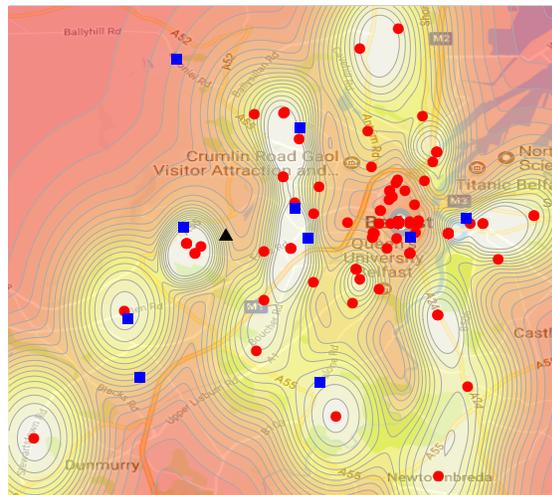
(l)

Phase 4, Factory 13



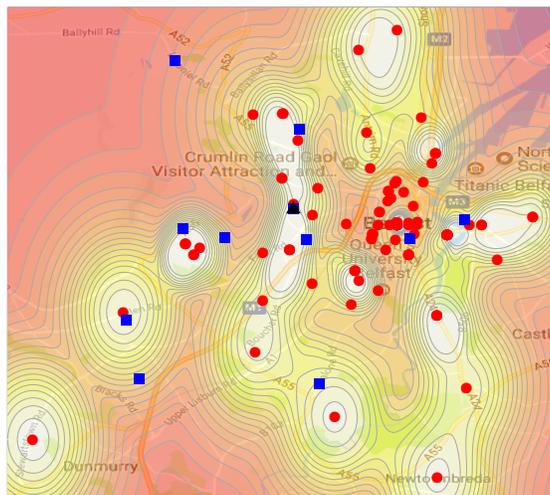
(m)

Phase 4, Factory 14



(n)

Phase 4, Factory 15



(o)

Figure 4.1: (a)-(o) Jeopardy surface distribution of PIRA IED attacks (red circles) before the closure of IED factories. Old factories closures are marked via blue squares whilst the new closure in each figure is marked with a black triangle. On the jeopardy surfaces white areas indicate likely source locations whilst red areas have a lower probability of containing a source.

Phase 5, All Factories

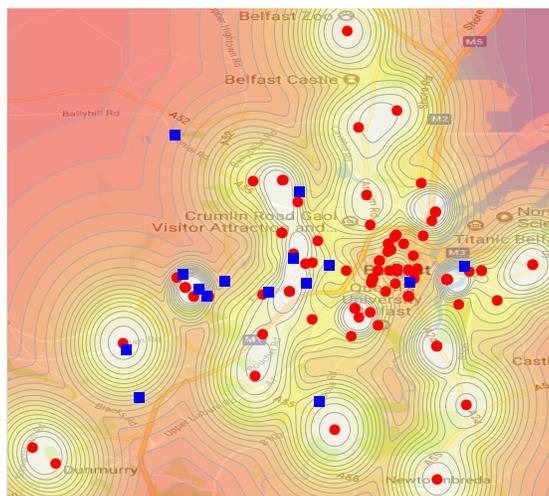


Figure 4.2: Jeopardy surface covering all PIRA IED attacks which occurred after the final recorded factory closure. Red circles highlight IED attacks and blue squares mark the locations of all previously identified factory closures. On the jeopardy surface white areas indicate likely source locations whilst red areas have a lower probability of containing a source.

4.6 Discussion

Before analysing the results in Section 4.5.2 it is necessary to recall the sociological phases of PIRA outlined in the literature review in Chapter 1. In the research of Asal et al. (2013) five distinct time periods of the conflict in Northern Ireland were identified which correspond to operational and tactical shifts of the PIRA. These five phases and summaries of the shifts observed in the organisation's structure are

- **1969-1976 - Phase 1:** During this phase the organisation was arranged in a military style consisting of brigades, battalions and companies. Within this army structure each unit of the organisation was given responsibility for a

specific geographical area of combat.

- **1977-1980 - Phase 2:** A cell-based structure was adopted. This approach was characterised by PIRA fracturing into small groups of members known as Active Service Units (ASUs) (Horgan and Taylor, 1997). The aim of this re-structuring was to improve the organisation's secrecy by making it harder to infiltrate. This change in structure was successful with 465 fewer charges for paramilitary activity within a year (Smith, 1997). During this phase new leaders were also appointed for the organisation including Gerry Adams and Martin McGuinness (Moloney, 2003).
- **1981-1989 - Phase 3:** This period began with the Hunger Strikes by Provisional IRA members protesting against the conditions of their incarceration. A catalyst moment during this period was the death of a PIRA member known as Bobby Sands who had been elected to Westminster whilst in prison and died on hunger strike (English, 2004). This incident resulted in a rise of sympathy for PIRA and its political wing Sinn Féin. As a result the Republican campaign moved into the political arena through the Sinn Féin party who now had similar levels of prestige as their militant wing, PIRA.
- **1990-1994 - Phase 4:** Secret meetings occurred involving top ranking PIRA leaders who were negotiating a ceasefire with the British Government.
- **1995-1998 - Phase 5:** Finally the peace talks were announced and a ceasefire ratified in the Good Friday Agreement. For many this signalled the end of "The Troubles".

Returning to the present chapter these qualitative descriptions of the phases of PIRA can be compared to the DPMM plots from Section 4.5.2. Specifically the timestamps of factory closures in Table 4.3 are related to the five phases in the following manner

- **1969-1976 - Phase 1:** Figures 4.1(a) - 4.1(c).
- **1977-1980 - Phase 2:** Figure 4.1(d).
- **1981-1989 - Phase 3:** Figures 4.1(e) - 4.1(i).
- **1990-1994 - Phase 4:** Figures 4.1(j) - 4.1(o).
- **1995-1998 - Phase 5:** Figure 4.2.

It can be seen that the jeopardy surfaces in Figures 4.1(a) - 4.1(c) are highly concentrated around a single area. This observation relates to the sociological description of PIRA in phase 1 as having a military style approach to the conflict. In particular, this military-type strategy could be leading to the source of attacks being centralised hence producing the concentration over one area found by the DPMM. In terms of the counter-terrorism responses in phase 1 it can be seen that the factory closures appeared away from the central mass of the jeopardy surfaces. One possible explanation for this distribution is that the centralised nature of PIRA in phase 1 acted as a form of protection for IED factories in more central parts of Belfast.

Moving to phase 2 of PIRA the organisation was seen in sociological research to have begun shifting to a cellular operation to thwart infiltration from British authorities. The DPMM model in Figure 4.1(d) demonstrates a similar pattern

where it can be observed that the concentration of mass on the jeopardy surface begins to disperse more widely over Belfast. An important shift in the discovery of bomb factories for phase 2 was that the recorded location is now closer to the central high probability source area found in the phase 1 plots. This indicates that the infiltration from the security forces was yielding better intelligence penetrating the protected area identified by the phase 1 DPMM results.

Phase 3 of PIRA corresponds to the longest period identified in the sociological research. During this phase there was an uptick in violence by PIRA with Moloney (2003) drawing an analogy between the tactics of PIRA in the 1980's and the waves of attacks conducted by the People's Army of Vietnam, known as the "Tet Offensive", in 1968. In Figures 4.1(e) - 4.1(i) the dispersal pattern observed in Figure 4.1(d) can be seen continuing in phase 3. Of particular interest in Figures 4.1(e) - 4.1(i) is the appearance of source areas in the North and North-West of Belfast. This highlights the expansionist nature of PIRA attacks in phase 3. The change in dispersal of possible source locations also resonates with the aim to wage a "war of attrition against enemy personnel which is aimed at causing as many casualties and deaths as possible so as to create a demand from their people at home for their withdrawal" stated as an objective in the PIRA "Green Book" (O'Brien, 1999). The increased diffusion of potential source locations as identified by the DPMM is also reflecting the relationship between the PIRA and security forces. In particular, as more factory closures occur the potential sources of new IED factories are shifted away from the previous factory closure indicating a certain level of displacement in factory locations.

In phases 4 and 5 PIRA began secret channels of communication with the British

Government to conduct peace negotiations ending in the announcement and ratification of the Good Friday Agreement. As noted by Asal et al. (2013) PIRA used IED attacks as a leverage in negotiations with the British authorities. Moreover, Coogan (2002) discusses the ability of PIRA to “turn...bombing[s] on and off like a tap”. These elements of the groups evolution are captured by the DPMM as further fracturing into numerous IED source locations. In this case the model could be seen to be providing evidence of PIRA using IEDs more strategically to strengthen the group’s hand during the peace deal discussions. Moreover, as PIRA was using IEDs more strategically as bargaining leverage the dispersal patterns observed might be the organisation utilising factories on a more temporary basis. This latter observation is particularly obvious in Figure 4.2 with the appearance of more disjoint source locations.

It is also interesting to consider the hitscore values presented in Table 4.3. These hitscores indicate how much of the search area identified by the DPMM must be searched to identify the closed factory locations considered in each dataset. Another interpretation of the hitscore is as a goodness of fit measure, whereby, a lower hitscore implies a better model fit. It can be seen from Table 4.3 that in general the hitscores for the closed factories are low. This observation suggests that the DPMM is good at locating the known factory locations. In addition, the low hitscores give evidence for Hypothesis 5 that the DPMM can find the locations of closed IED factories. When the hitscores are high this may be explained by data quality. In particular, it was described in Section 4.1 that the IED dataset had to be filtered to obtain data on a street level, which, reduced the number of data points available for use from

the original dataset. Since some of the IED events in the unfiltered dataset did not have street level data these entries would not have been analysed in this chapter. Therefore, some of the located IED factories may have been related to events not considered by the DPMM. Moreover, to improve computational efficiency, and due to limited data, the research in this chapter focused only on Belfast. However, it may be possible that some of the recorded IED factories were linked to IED attacks outside of the Belfast area. Without more data and expertise to link factories to IEDs this issue remains.

Combining these phase-by-phase observations it can be seen that this temporal version of the DPMM provides important lessons for the construction of effective counter-terrorism strategies. The jeopardy surfaces produced in this chapter demonstrate that over the course of “The Troubles” factory raids were occurring more frequently within source locations identified by the model. However, the DPMM was also revealing potential IED factory locations which do not appear to have been discovered by authorities working to counter PIRA attacks. In particular, observe that at the top of Figures 4.1(h) - 4.1(o) there is clearly a source being identified by the DPMM but with no follow-up factory raids. Therefore, had the framework as illustrated in this chapter been employed as part of the authorities tools to identify factories there is a chance that at least one further closure could have been achieved. Another factory identification may also have had further implications in terms of intelligence gathering.

Further to the contributions the DPMM can make to counter-terrorism practitioners the model can also be useful to academics. In particular, obtaining terrorism

data can be a difficult task due to security issues. This problem could become more significant when trying to conduct research on currently active terrorist groups. In this chapter the ability of the DPMM to realise IED sources without prior information concerning the number of sources can assist researchers in collating datasets. Specifically, the sources identified by the DPMM could act as a guide for researchers to understand geographical patterns of terrorist activity based on open source data of attacks.

Overall the DPMM has proven to be a very useful tool to further uncover shifting tactical and hierarchical structures used by PIRA during its IED campaign. This gives evidence to support Hypothesis 7 stating that the temporal framework for the DPMM could lead to new insights into spatio-temporal patterns of IED usage. The similarities between the findings of sociological research and this empirical space-time usage of the DPMM give credit to the model as a good method for counter-terrorism researchers and practitioners to understand terrorist groups. In addition to this support for the DPMM model it has been shown that the model can act as a testing tool to understand the impacts of counter-terrorism events aimed at closing IED factory locations. In particular, we have seen some level of displacement of areas most likely to contain the sources of IED attacks suggesting that following a factory closure there is a need to increase security operations to stifle the emergence of any new factories opening. Moreover, we saw that the negotiating period of PIRA was characterised by more temporary usage of disjoint IED factories. This suggests that if militant groups choose to begin moving away from conflict towards a ceasefire and political integration there is a need for security forces to increase intelligence

gathering to remain ahead of factory openings. These observations provide support for Hypothesis 6, which, considers the efficacy of the DPMM to provide a method of assessing count-terrorism policies aiming to close IED factories.

In future research the work in this chapter could be progressed by including temporal dynamics directly inside the framework of the DPMM model. In particular, having a specific time component as part of the model formulation could allow for greater usage of the model as a predictive tool by stepping forward the time component to uncover future potential dynamics. Another useful extension of the research presented in this chapter could be to find spatio-temporal boundaries for PIRA IED attacks. In particular, in Chapter 2 it was demonstrated how the Hawkes process can be used to determine temporal phase boundaries mathematically. Therefore, a coupling of the methods used in this chapter and those of Chapter 2 might present an opportunity to determine quantitatively both the spatial and temporal phase changes for terrorist group activity.

Finally, it should be noted that in Figures 4.1(a) - 4.1(o) and Figure 4.2 an implicit temporal memory is incorporated into the model since the effect of past factory closures is not considered. The impact of this treatment of time is that some of the IED events in later figures are possibly leading to false groupings of attacks. In a real-world situation this problem could be managed by utilising more detailed data about which IED events correspond to which factories. Having this additional detail would enable practitioners to remove IED events which are no longer relevant to the model. Hence, although this chapter provides a good prototype of how the DPMM can be used its full potential and accuracy is likely being underestimated.

CHAPTER 5

MANUAL VS AUTOMATIC DATA COLLECTION METHODS

5.1 Introduction

In a paper published by Moore (1965) an observation was made that the number of transistors in integrated circuitry was approximately modelled by the equation 2^y where y is the number of years since 1959. Qualitatively this observation, that the number of transistors doubles each year, has become known as Moore's Law. Due to a self-fulfilling nature, which observed the technology industry aiming to improve their products at this pace (Brock, 2006), Moore's Law is considered to have been a dominant theory underpinning much of the growth seen in technological advancements over the last five decades.

Alongside a rapid growth in technological capabilities has been the globally expanding reach and accessibility of the world wide web. First presented by Tim Berners Lee to CERN in 1989 (McPherson, 2010) the world wide web has expanded the ability of the human race to rapidly share information across the globe. This has led to advancements in scientific discovery and research as academics can seamlessly share and discuss their ideas. In the discussion of Coffman and Odlyzko (2002) the idea of a Moore's Law for the rate of growth of the internet is presented suggesting a link between the developments observed physically in the technology sector and wider usage of virtual data accessible via the web.

Combined new technologies and their use for data sharing are quickly beginning to evolve the landscape of 21st Century scientific research. The observance of this shift in research methods has been discussed by Hey, Tansley and Tolle (2009). In particular, the authors of the latter highlight four phases of progression in science

based research - empirical, theoretical, computational and data exploration. The empirical phase is described as the observation of natural phenomena and involved the manual collection of data approximately a thousand years ago concerning the changes observed in people's environments. As scientific reasoning and methods progressed the next shift was towards a theoretical approach to studying the world around a few hundred years ago. The theoretical phase saw the introduction of mathematical models and their use to abstract and generalise ideas and understandings obtained from empirical studies. More recently science has transitioned through the computational phase in the last few decades where the models which were introduced in the theoretical stage could now be simulated using newly available computational power afforded by the advancements in computer design. The current phase is that of data exploration. Now with new computational hardware and rapid sharing features provided by the web a combination of empirical, theoretical and computational research can be undertaken in large volumes on "big datasets".

An introduction to "Big Data" and its implications for modern society can be found in the book of Mayer-Schönberger and Cukier (2013). Big data is commonly described as exhibiting three properties known as volume, variety and velocity. Volume of data is simply a measure of the size of modern datasets which are now able to be stored in increasing magnitude due to the improved effectiveness and efficiency of modern hard drives. Within these large datasets it is possible to store a variety of information ranging from easily tabulated structured data, such as, numerical values, to unstructured data which requires more involved preprocessing, for example, collections of text documents. Finally, data is frequently moving either across the web on

external networks or through computers on internal networks leading to the a metric of data velocity. When referring to Big Data in scientific studies Demchenko et al. (2013) also include two more “v’s” of big data - veracity and value. Veracity relates to the need to have some method of data validation to ensure that the data being obtained provides a reliable basis for accurate analyses. Additionally, determining the added value provided by a dataset ensures that the scope of research remains focused to avoid scientific pitfalls, such as, linking correlation and causation.

With the beginning of a new phase in research centered on data exploration there is a requirement to develop new tools and techniques which can enable researchers to collate and analyse Big Data sources. There are three main branches of research which cover this area (Getoor and Machanavajjhala, 2012) - information retrieval, data mining and machine learning. Information retrieval (IR) is a term first coined by Mooers (1950) and aims to connect users with data. This can involve searching both structured datasets, such as, organised tabular databases and unstructured datasets, for example, searching text documents for key words or phrases (Manning, Raghavan and Schütze, 2008). A number of approaches to IR already exist including studies of Boolean, vector and probabilistic models (Baeza-Yates and Ribeiro-Neto, 2011; Manning, Raghavan and Schütze, 2008; Roelleke, 2013).

Having retrieved data via some combination of IR methods the next step is to obtain value from the data. A primary method to achieve this goal is to use the knowledge discovery in database (KDD) process. As described by Fayyad, Piatetsky-Shapiro and Smyth (1996) the KDD process is a method which segments the managing of large datasets into multiple steps moving from data retrieval to analysis

and then result validation. One of the main steps in the KDD process is data mining which is “a process of discovering useful patterns and trends in large datasets” (Larose and Larose, 2014). As with IR the field of data mining has begun rapidly growing as researchers develop tools to manage “Big” datasets. Some of the current methods of performing data mining are classification, regression, clustering, summarisation, dependency modelling and change and deviation detection (Fayyad, Piatetsky-Shapiro and Smyth, 1996).

Machine learning (ML) is an umbrella term which covers tools and techniques developed to enable computers to derive value from data via learned structures and patterns (Alpaydin, 2014). There are two main classes of machine learning - supervised and unsupervised (Alpaydin, 2014; Masashi, 2016; Müller and Guido, 2017). In the case of supervised learning some initial training input and observed output is used to train a ML algorithm. On the other hand, in unsupervised learning only inputs are provided to the ML algorithm. Methods of IR and data mining can be used as tools to develop machine learning approaches and vice versa (Aggarwal and Zhai, 2012; Craven et al., 2000; Liu and Motoda, 2001).

As the power of computing resources has grown, and the accessibility of data has increased, the applications of Big Data, and the tools required for its analysis, have become more widespread. Of particular interest in this chapter is the usage of these techniques as a basis for mathematical studies of terrorism and counter-terrorism. In particular, multiple books and research papers have been published examining a range of possibilities to use newly developed technological tools to monitor and analyse terrorist activity in real or near-real time (Baesens et al., 2009; Chen et al.,

2008; Leslie, 2016; Memon et al., 2009; Thuraisingham, 2003; Toure and Gangopadhyay, 2016). As in all areas of research these novel approaches to study terrorism must ensure that the data being utilised provides an accurate reflection of real-world events. If there is a failure to retrieve good data or a reliance on incomplete datasets then any resulting conclusions may be incorrect and lead to poor counter-terrorism strategy design.

In this chapter the aim is to address the issues of developing tools to automatically collate terrorism databases and methods to evaluate their usefulness. In particular, an important issue that the tools and techniques of information retrieval and machine learning can be used to tackle is information overload in manual data collection. For example, when collecting data concerning terrorist attacks there are a plethora of local and international news sources from which information can be gathered. Collating datasets from many sources manually is time intensive requiring a large outlay of human capital to be effective. However, with the speed and efficiency of modern computers and communication methods via the web there is potential to greatly simplify the task of data retrieval. Two main contributions will be made in this thesis. Firstly, an IR tool will be implemented in Python which can interact with online news articles to retrieve terrorism event data. Secondly, a unique evaluation dataframe will be designed combining methods of temporal and textual similarity measurements to compare the IR dataset to a manually collected database.

5.2 Datasets

5.2.1 Manually Collected Dataset

The manually collected dataset is derived from the Armed Conflict Location and Event Dataset (ACLED) project (Raleigh et al., 2010). As part of the ACLED project numerous research analysts are involved in collecting and storing data about conflicts in Africa as well as South and South East Asia. The database released from ACLED appears in two versions. The first dataset has complete year-on-year daily coverage of data since the project began in 1997. In the second dataset a real-time rolling coverage is provided which details recent weekly updates of events. The real-time data is annually reviewed and integrated into the complete dataset which is then released as a new full coverage archive.

In this chapter to complement the earlier research of Chapter 4 the focus will be the Al Shabaab militant group in Somalia. Moreover, to simulate a real-world setting for the tools developed the real-time ACLED data will be considered. Specifically, the 2017 ACLED rolling coverage dataset, which is manually collected, will be compared to an automatically retrieved dataset, which is detailed in Section 5.2.2. The final event date recorded in the ACLED database is 14/10/2017 . The data collected for the ACLED project covers the topics displayed in Table 5.1.

It should be noted that the source of the data for the ACLED database is split into two categories - non-local and local. Here the local sources are specific to ACLED which is connected to local organisations working regionally and reporting events to

Table 5.1: ACLED Data Headings

Heading	Data
Date	Date of the event (DD/MM/YYYY)
Location	Administrative areas and latitude/longitude coordinates
Event Type	{battle, civilian killing, riots, protests, recruitment activity}
Event Actors	{rebels, governments, militias, armed group, protesters, civilians}
Change in Territory	Details of any changes seen in area control
Source	Source of data {Non-Local, Local}
Notes	Notes covering details of the event
Fatalities	Estimated number of fatalities

the ACLED team. The non-local sources are more readily available articles which can be retrieved online. These two categories will be compared separately to the IR dataset as explained using the framework detailed in Section 5.4. The number of events across the two datasets are non-local sources - 268 and local sources - 881.

5.2.2 Automatically Collected Dataset

In the first unique contribution made in this chapter an automated information retrieval (IR) tool was written in a Python script. Specifically using the Python requests module (Chandra and Varanasi, 2015) an automated tool was constructed

Table 5.2: Automated Article Retrieval Data Headings

Heading	Data
Date	Date of the event (DD/MM/YYYY)
Description	Article title and first sentence from article

which searches Google news for user defined terms. Once the results have been retrieved for a search the Hypertext Markup Language (HTML) structure returned is saved to a csv file. To extract relevant data the saved HTML structures are parsed by searching for tags where important information is contained and retrieving the corresponding text. This textual data is then used to form an event database. In particular, for the database studied in this chapter the search term used was simply “al Shabaab” in an attempt to observe as many events as possible involving the group. In addition to a direct search for articles the code performs the search for each month of the year. The motivation for this temporally disaggregated searching method was that it resulted in a higher number of returned articles compared to longer time-frame searches. Data headings, and their descriptions, found in the final dataset are illustrated in Table 5.2.

In contrast to the ACLED dataset the IR data has only non-local sources having been collected entirely from online news articles. In total 960 articles were obtained using the Python IR tool.

5.3 Comparison Techniques

In this chapter the aim is to compare and contrast the data obtained from the manually collected ACLED data and an automatically retrieved dataset from online news sources. The two types of data appearing in Tables 5.1 - 5.2 are dates of events and textual data. Two tools from the current literature are employed in this chapter to analyse these types of data. The background of these pre-existing techniques are now introduced.

5.3.1 Dynamic Time Warping

To compare the event dates a technique known as Dynamic Time Warping (DTW) was employed to study the similarity between the datasets. In the research papers of Berndt and Clifford (1994), Ding et al. (2008) and Keogh and Ratanamahatana (2005) the fundamental motivations and assumptions for using the model are described alongside experimental applications demonstrating the approaches effectiveness. Principally, DTW is a search algorithm which provides a mapping between a template time series $T = \{t_1, t_2, \dots, t_m\} = \{t_j\}_{j=1}^m$ and a second comparison time series $S = \{s_1, s_2, \dots, s_n\} = \{s_i\}_{i=1}^n$. This mapping aims to find the minimum distance between the time series.

The initial step for constructing the DTW algorithm is to align the indices of the template time series against the second time series in a two dimensional grid. In this case the grid consists of points (i, j) which can be interpreted as a comparison between time points s_i and t_j . With this grid so constructed a vector of each of the

possible elements $w_k = (i, j)_k$ can be constructed. Selecting one element from each consecutive vector leads to a possible warping path

$$W = w_1, w_2, \dots, w_k, \dots, w_K. \quad (5.1)$$

For each of the points on this grid a comparison can be quantified using some distance function $d(s_i, t_j)$. A simple approach to measuring distance between time points is to use an absolute value measure as described by Berndt and Clifford (1994)

$$d(w_k) = d((i, j)_k) = |(s_i - t_j)_k|. \quad (5.2)$$

After defining an approach to measuring the distance between time points the main DTW problem can be introduced. In particular, the aim now is to find the optimal warping path between the series S and T by minimising the cumulative distance

$$DTW(S, T) = \min_W \sum_{k=1}^K (d(w_k)). \quad (5.3)$$

It should be observed that the cumulative summation in equation (5.3) is defined over the range of all possible warping paths. For large datasets this optimisation problem will be computationally expensive due to its combinatorial nature. Therefore, some assumptions are made about the DTW to improve efficiency (Berndt and Clifford, 1994)

1. Monotonicity - the time series used in the problem are assumed to be monotonically increasing in time, such that, the indices of S and T satisfy $i_{k-1} \leq i_k$

and $j_{k-1} \leq j_k$,

2. Continuity - each subsequent element in the warping path must only pass through neighbouring points so that $i_k - i_{k-1} \leq 1$ and $j_k - j_{k-1} \leq 1$,
3. Warping Window - when comparing points (i, j) from the grid there is a maximum distance of comparison $|i - j| \leq \omega$, where, ω is a positive integer value defining the window width,
4. Slope Constraint - to avoid large movements in any single direction it is possible to define a restriction on the steepness of possible warping paths,
5. Boundary Conditions - fixing the start of each warping path to $(1, 1)$ and the end to (n, m) ensures the path always moves between the lowest left-hand-side of the grid and finishes at the highest right-hand-side.

With the above assumptions the final DTW can be formulated as a dynamic programming problem using the function

$$\gamma(i, j) = d(i, j) + \min [\gamma(i - 1, j), \gamma(i - 1, j - 1), \gamma(i, j - 1)]. \quad (5.4)$$

From the equation in (5.4) the DTW is computed via a recurrence formula which computes the cumulative summation of the current grid element distance $d(i, j)$ and the minimum of the cumulative summations at neighbouring elements. The form of the DTW in (5.4) is known as the symmetric version of the problem due to its use of both previous points found diagonally on the grid. An asymmetric version of the equation in (5.4) also exists by choosing only one of the diagonals. However, the

research of Sakoe and Chiba (1978) found that the symmetric form of the problem to be the better choice.

The advantage of the approach to DTW using (5.4) is that the algorithm does not require re-computing previous path distances. In particular, at each step of the dynamic programming technique a table of values is computed which can be used to find the next entry of the cumulative sum in (5.4). To complete the DTW and find the minimised path matching series S and T this table can be traversed backwards finding the points corresponding to the lowest cumulative distance values.

In this chapter an implementation of the DTW in Python by Rouanet (2018) was used to perform temporal analyses.

5.3.2 Term Frequency - Inverse Document Frequency

The second area of overlap that is important to this thesis, and terrorism studies more broadly, is the spatial data of terrorist activity. Table 5.1 illustrates that the ACLED dataset contains location data whilst the IR dataset described in Table 5.2 has only textual information in the form of the article titles and the articles first sentence. A simple method that could be used to overcome this difference between the databases would be to simply scan each of the IR article titles and first sentences for the locations in the ACLED database. However, this approach presents multiple difficulties. In particular, difference in the spelling of place names, such as, (“Mogadisho”, “Mogadishu”) and (“Buulo Gaduud”, “Bulo-Gadud”) can prove difficult to cross-reference.

In addition to comparing spatial data between the datasets it is also of interest to

compare the text data contained in the “notes” section of ACLED to the “description” information in the IR events. This further analysis between the two datasets was used to provide insight into the usefulness of the IR data to offer more detail about the events being recorded.

To overcome the difficulties of comparing textual data and provide a comparison between location and notes data of the ACLED and IR datasets a method known as term frequency - inverse document frequency (tf-idf) was employed. As the name of the technique indicates tf-idf is a two step approach for comparing text documents.

Term frequency (tf) is the study of occurrence of a notion or combination of notions within a single document (Luhn, 1957). As noted by Luhn the more frequently a word or combination of words within a document the more important the word or combination is likely to be for the subject being studied. Hence a simple measure of the relative importance of different parts of text within a document can be obtained using a simple word frequency. However, the problem with such a simplistic result is that when comparing multiple documents according to simple word frequencies words, such as, “terrorism” in papers in the field of terrorism studies can lead to broad matching results making searches less useful.

To resolve the issue of matching documents with highly occurring keywords but low overall importance Spärck Jones (1972) suggested utilising a term weighting approach. In particular, when searching across a range of documents high frequency terms which likely have little meaning are given low weighting whilst terms of more rarity are given a greater weighting. This inverse relationship leads to the term inverse document frequency (idf). Under this scheme the issue of retrieving irrelevant

documents containing the word “terrorism” in a field studying terrorism is largely resolved. A commonly used formula to find the inverse document frequency for a term t_i occurring n_i times in a collection of N documents is (Robertson, 2004)

$$idf(t_i) = \log \left(\frac{N}{n_i} \right). \quad (5.5)$$

The term frequency (tf) multiplied by the inverse document frequency (idf) yields a measure of importance of a term to a document in a collection of documents.

As well as being useful as a matching tool for searching databases of documents for keywords the tf-idf can also be extended to a document comparison tool using a tf-idf cosine weighting (Peters, Braschler and Clough, 2012). The tf-idf can be considered a mapping from the space of the text documents to the d -dimensional real space \mathbb{R}^d . For example, one could form a vector of all words across all documents and for each vector element the tf-idf value could be computed. After the tf-idf calculation each document is now a vector of real numbers in space whose dimension is equal to the total number of words. Hence for two documents with tf-idf vectors denoted d_1 and d_2 it is possible to compute the cosine of their angle θ via

$$\cos(\theta) = \frac{d_1 \cdot d_2}{\|d_1\| \|d_2\|}. \quad (5.6)$$

For two similar documents the formula in (5.6) will have tf-idf vectors close in space and thus a small angle difference leading to a value of $\cos(\theta)$ near to 1. On the other hand, documents which are dissimilar will lead to values of $\cos(\theta)$ nearer 0. Note also that since tf-idf vectors contain only non-negative values we have that $\theta \in [0^\circ, 90^\circ]$.

In this chapter the tf-idf implementation from Python's scikit-learn module (Pedregosa et al., 2011) was used to perform the textual analyses.

5.4 Dataset Comparisons

To compare the ACLED and IR datasets a novel approach of combining dynamic time warping and the tf-idf methods was designed. In particular, a hybrid method was developed to evaluate the similarity between the two sets of data based on the event timings and textual descriptions.

5.4.1 Event Date Comparisons

For the temporal component of the data the standard DTW defined in Section 5.3.1 was adapted. Specifically for each of the dates observed in the IR database the set of closest dates in the ACLED database were identified. The cardinality of this set can range from a single element to multiple elements if there were multiple entries for a single date. However, at this top level of analysis the aim is to simply identify a closest date matching the IR event recording. In this case a perfectly matched date found in the IR and ACLED event recordings would produce a distance of 0 whilst any dissimilarity would be measured in the number of days difference.

5.4.2 Location Comparisons

As previously described in Section 5.2 the ACLED dataset has separate columns containing location data for the observed events. In the ACLED dataset the location

data is split into multiple columns covering different levels of spatial disaggregation. To compare this spatial information to that contained in the IR dataset the tf-idf method was employed. Specifically, each of entry of location data from the ACLED database was compared with the the text strings obtained from the IR dataset via tf-idf. To improve the efficiency of these textual based comparisons in the next level of analysis conducted the sets of closest event matches between the IR and ACLED databases from Section 5.4.1 were looped over and the IR elements in each set had their title and first sentence data compared to the corresponding ACLED event location data. To achieve this comparison the following procedure was performed for the elements of each closest events sets

1. Split text strings from IR data entry into a list of words in Python.
2. For each of these list elements extract all words beginning with an uppercase letter. This step follows from the observation that location names are written with capital letters.
3. For each upper case word lower all the characters, remove punctuation and perform a character comparison to the lower cased version of the location data from the ACLED event. The character comparison was performed using the tf-idf described in Section 5.3.2. In this context the tf-idf was set to perform a character comparison which considers similarity based on the number of letters matched between the inputs. This approach is used to overcome the issue of the different spelling of location names, such as, (“Mogadisho”, “Mogadishu”) and (“Buulo Gaduud”, “Bulo-Gadud”).

4. Return the highest matching value,
5. Repeat the above procedure for all location information of the event contained in the ACLED database.

The matching value from this procedure will be a number in the range $[0, 1]$, where, 0 represents no match and 1 represents a perfect match.

5.4.3 Event Note Comparisons

The final comparison technique performed between the databases was to compare the “notes” section of the ACLED events to the “description” data in the IR dataset. Similar to the location comparison the textual comparison was conducted using the tf-idf implementation of the scikit-learn module. However, instead of decomposing the text into characters the full sentences were passed as inputs to the function. The matching value, in the range $[0, 1]$, as before, was then returned. In addition, the note comparisons were made using the set of closest times found between the ACLED and IR datasets in Section 5.4.1 using the dtw method.

5.4.4 Comparisons Framework

These analyses yield three results, where, the IR data used is listed first and the data from the ACLED database is listed second

1. (Date) and (Date) time comparison
2. (Date, Description) and (Date, Locations) location comparison

Table 5.3: Example of Matching Events

Dataset	Event Date	Event Details
ACLED	15/08/2017	“Five police officers have been killed by Al-Shabaab militia while on patrol along Bothai-Ijara road, within Garissa County. One officer sustained injuries during the attack while another one managed to escape unhurt.”
IR	15/08/2017	“Nairobi - Five police officers have been killed by Al-Shabaab militia while on patrol along Bothai-Ijara road, within Garissa County. One officer sustained...”

3. (Date, Description) and (Date, Notes) event description comparison

As described in Section 5.2 the sources for the ACLED can be divided into local and non-local categories. To further disaggregate the comparisons made between the ACLED and IR databases the above analyses were repeated using the non-locally sourced ACLED data and only locally sourced ACLED data.

An example of the type of matches observed between the IR and ACLED databases is shown in Table 5.3. In particular, this table illustrates the best time-text match between the ACLED non-local dataset and IR dataset. It can be observed that the timestamps in Table 5.3 are identical whilst the textual descriptions of the events are almost identical. However, since the text in the IR data entry is only a snippet of the full article the final sentence is truncated reducing the overall matching according to the tf-idf to 68.2%.

5.5 Results

Sections 5.5.1 - 5.5.2 present the results of the analyses described in Sections 5.4.1 - 5.4.3. Two sets of results are presented according to comparisons between the IR dataset and the non-locally sourced ACLED data and the IR dataset and only locally sourced ACLED data.

Each set of results consists of six graphs. The first plot is a simple barcode of the event times occurring in each dataset. These event times are then compared using the DTW with plots provided showing the warping path between the ACLED and IR timeseries. After analysing temporal relationships the textual comparisons for each set of closest event times are made via the tf-idf method. The highest tf-idf values for these comparisons are plotted for both location and event text descriptions. Finally, cumulative percentages of the tf-idf values are shown illustrating the percentage of event comparisons with tf-idf exceeding threshold values in the range $\{0.1, 0.2, \dots, 1.0\}$. These cumulative plots indicate the distribution of tf-idf values.

5.5.1 ACLED Non-Local Source Comparison

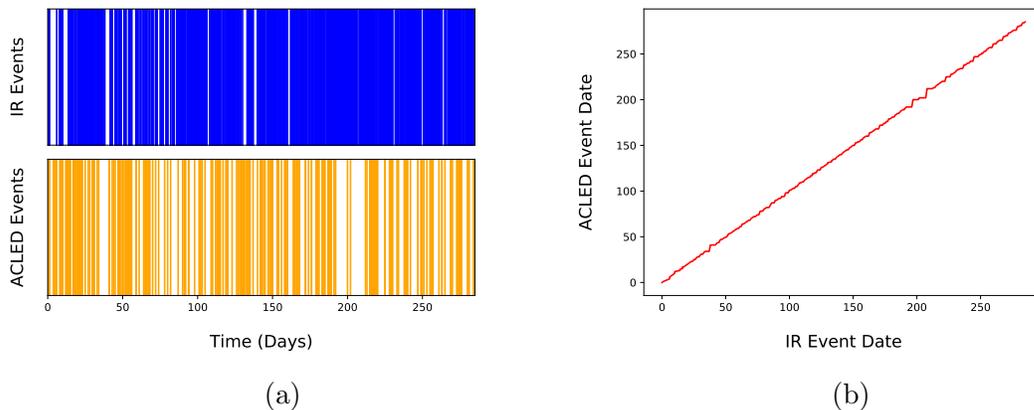


Figure 5.1: (a) Barcode plot comparing the timeseries of IR and ACLED events in days since each dataset's first event. (b) Time warping path between the IR and ACLED timeseries. This path is derived from the dynamic time warping (DTW) method.

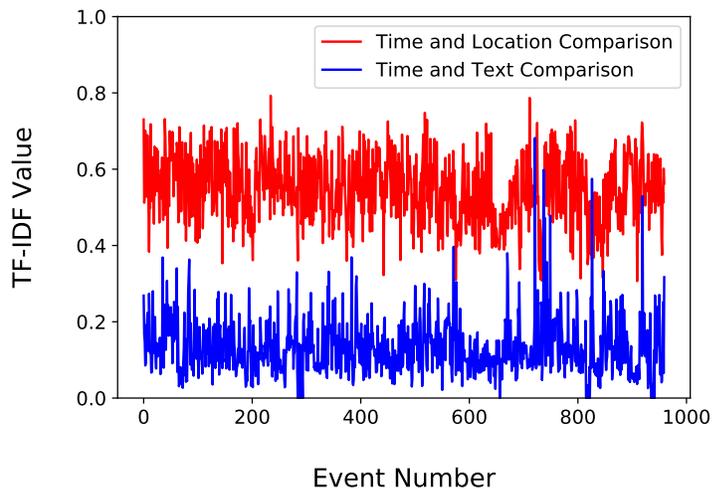


Figure 5.2: Location and event description comparison of the ACLED and IR datasets. For each event in the IR dataset the closest datapoint in the ACLED database in time and location is shown in red. Similar results are shown in blue for the closest event in time and event textual description.

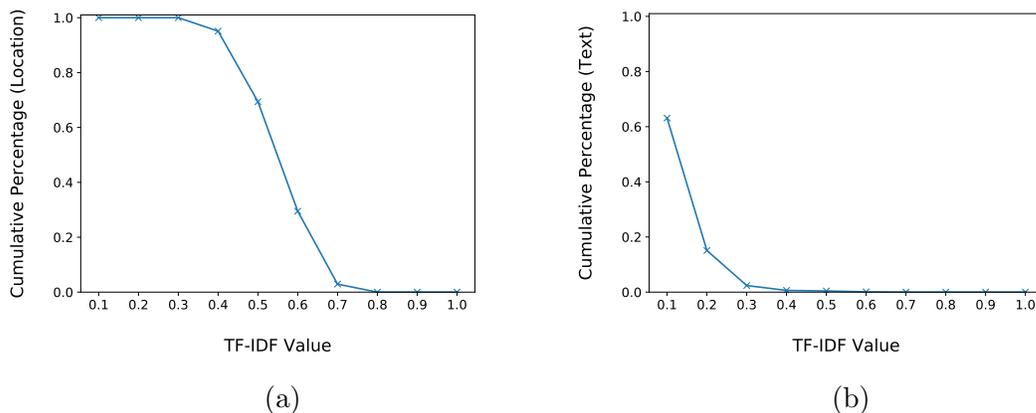


Figure 5.3: (a) Percentage of event locations matched with a tf-idf value greater than or equal to threshold values $\{0.1, 0.2, \dots, 1.0\}$. (b) Percentage of event descriptions matched with a tf-idf value greater than or equal to threshold values $\{0.1, 0.2, \dots, 1.0\}$.

5.5.2 ACLED Only Local Source Comparison

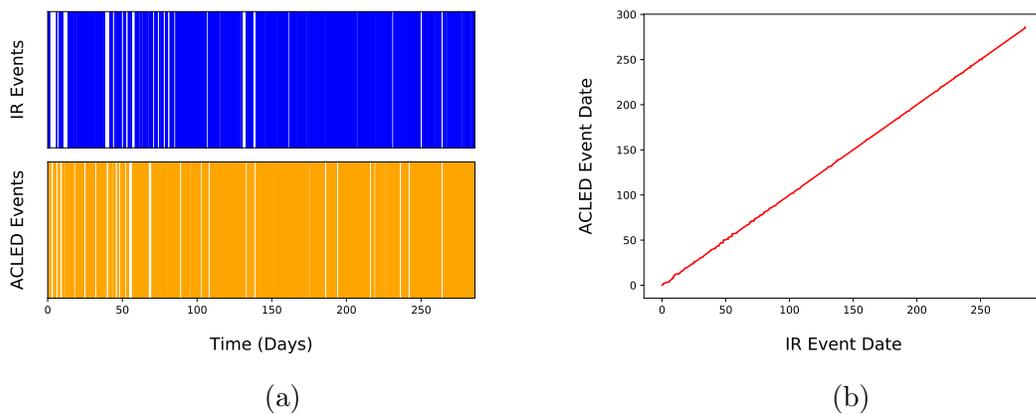


Figure 5.4: (a) Barcode plot comparing the timeseries of IR and ACLED events in days since each datasets first event. (b) Time warping path between the IR and ACLED timeseries. This path is derived from the dynamic time warping (DTW) method.

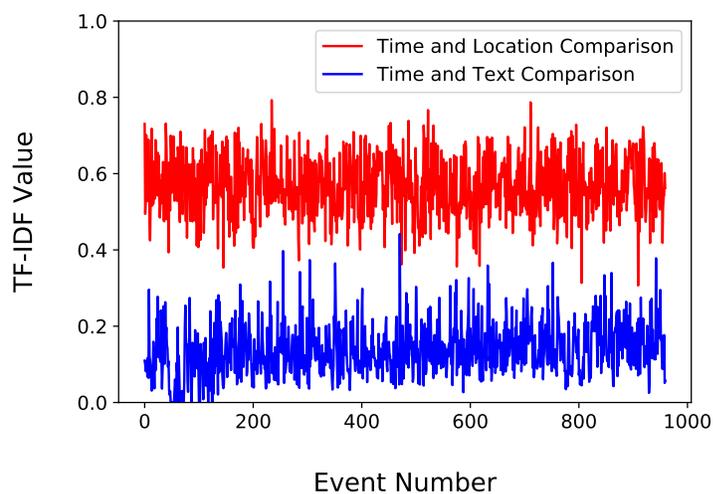


Figure 5.5: Location and event description comparison of the ACLED and IR datasets. For each event in the IR dataset the closest datapoint in the ACLED database in time and location is shown in red. Similar results are shown in blue for the closest event in time and event textual description.

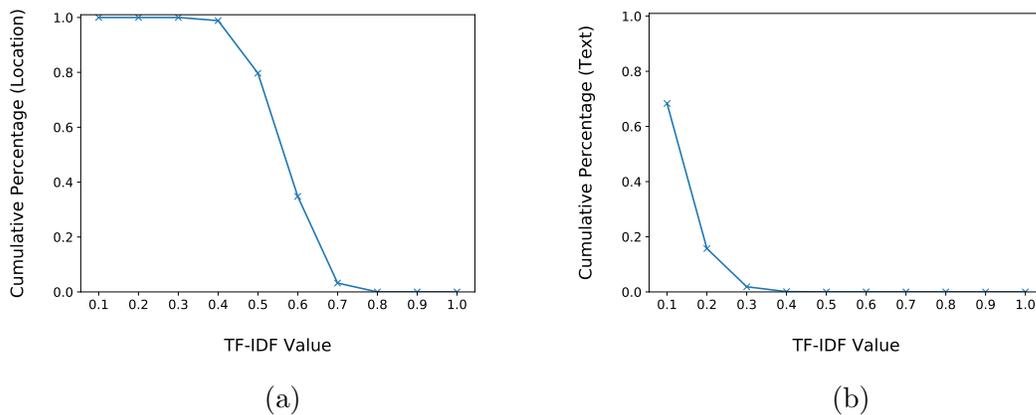


Figure 5.6: (a) Percentage of event locations matched with a tf-idf value greater than or equal to threshold values $\{0.1, 0.2, \dots, 1.0\}$. (b) Percentage of event descriptions matched with a tf-idf value greater than or equal to threshold values $\{0.1, 0.2, \dots, 1.0\}$.

5.6 Discussion

The graphs in Section 5.5 compare and contrast the ACLED dataset to the automatically mined IR dataset. There are two sets of results corresponding to a different disaggregation of the ACLED data according to the data's recorded source. In Figures 5.1 - 5.3 comparisons are made using only the data in the ACLED set derived from non-local sources, which, covers all articles and events manually read on the internet. The second set of graphs in Figures 5.4 - 5.6 are used to study the similarity between the ACLED data sourced from local projects, which, is reported to ACLED from local organisations, and the IR dataset.

To investigate the difference of temporal event records between the two datasets a Python implementation of the DTW method (Rouanet, 2018) of Section 5.3.1 was used. The results are presented in Figures 5.1(b) and 5.4(b). It can be observed from these plots that on a purely temporal basis the IR dataset events are more closely aligned to the locally source ACLED data. Moreover, it was found that 94% of event times had an exact match between the IR and local sourced ACLED data compared to only 56% for the IR and non-locally sourced ACLED data.

In addition to studying the temporal event records further depth of analysis can also be obtained by comparing locations data from the ACLED and IR datasets. In particular, for the set of events identified as having the closest distance in the previous temporal comparison a search for location data was performed. This comparison was undertaken using the tf-idf measure outlined in section 5.3.2. In this case the ACLED dataset was used as a benchmark with the source locations from the dataset

compared character-wise to the words in the event descriptions obtained using IR as discussed in 5.4.2. The higher the percentage from this comparison the more likely it is that the ACLED location recording was mentioned in the event description.

By comparing the plots in Figures 5.2 and 5.5 it can be seen that the location comparisons for both non-local and local sources produce similar results. In particular, the non-local analysis has a mean tf-idf of 0.55 and standard deviation of 0.09. On the other hand, the only local comparison has a mean tf-idf of 0.57 and standard deviation of 0.08.

The plots in Figures 5.3(a) and 5.6(a) show the cumulative percentage of tf-idf comparisons with value greater than values in the range $\{0.1, 0.2, \dots, 1.0\}$. Both non-local and local source data yielded similar patterns of high cumulative percentages which decrease rapidly around the 0.5 threshold.

The final comparison that was made between the ACLED data and the IR dataset was to analyse the difference between the textual content provided for the ACLED events and the event descriptions in the IR data. In a similar approach as the one taken to compare spatial details for each of the closest matches found in time between the events a match percentage between their event descriptions was calculated. As before the tf-idf metric was used to compare this textual data.

From the graphs in Figures 5.2 and 5.5 it can be seen that compared to locations comparison the textual descriptions of the ACLED and IR datasets have lower tf-idf values. This indicates a lower accuracy when considering the textual event data in each dataset. The mean tf-idf value for the non-local comparisons was 0.13 with a standard deviation of 0.07. On the other hand the local source data results had a

mean tf-idf of 0.14 and standard deviation 0.07.

In the graphs of Figures 5.3(b) and 5.6(b) it can be observed that the cumulative percentages of tf-idf values are similar for the non-local and only local results. Moreover, these plots show that the textual comparisons are worse than the results obtained for location comparisons.

It is interesting to note that the IR dataset has 960 data entries which is greater than the 268 non-local source entries in the ACLED database. In particular, this provides evidence that the IR technique presented in this chapter could be of benefit to supplement the ACLED data collection method. It can be observed from Figures 5.1(a) - 5.1(b) that there is divergence between the IR and non-local source ACLED time series. However, this may be explained by the use of less widely used news sources by the ACLED data gathering team, such as, online news organisations based in Somalia. As a consequence it may be that the ACLED non-local source data is obtaining some data that is unlikely to occur in a standard browser search. This issue could be resolved by redesigning the IR tool developed in this chapter to work with the specific websites used for the ACLED project.

Within the local source ACLED dataset there are 881 events. When compared to the local source ACLED data it can be seen from Figures 5.4(a) - 5.4(b) that there is a close relationship with the event times recorded in the IR dataset. Hence the IR tool is able to find news sources which provide information as effectively as local source organisations. This observation is promising since IR methods like those used in this chapter provide a cheaper and simpler option for obtaining data about terrorism related conflicts.

For both the non-local and local source ACLED data the textual comparisons with the IR dataset yield similar results. Location comparisons were matched with higher accuracy than event descriptions. This fall in accuracy may be the result of considering only first article sentences in the IR database which may not be providing a good summary of the main article text.

The results of the research conducted in this chapter indicate partial evidence for Hypothesis 8. In particular, this hypothesis described the possibility of producing an automatic data retrieval tool that could collect temporal, spatial and event description details comparable to a manually collected dataset. As has been discussed temporal data was matched with high accuracy between the IR and ACLED datasets. With lower matching some spatial and event description data was also found to be contained in the IR collected database.

As has been demonstrated in all the previous chapters of this thesis temporal dynamics can prove very useful for providing insights into terrorist groups and their activities. Therefore, having such quickly accessible data automatically collected yields great potential in the field of terrorism studies. A particularly useful application of this approach to data collection is the ability to monitor in near real time terrorist events and integrate the feed into mathematical models to generate a fully automated terrorism analysis tool. For example, the timestamps from the IR dataset could be used as input into the Hawkes point process model in Chapter 2 to generate a live alert system for monitoring shifts in a terrorist organisation's structure. The advantages of employing such a technique would likely be reaped by both academics and practitioners. In particular, in real world scenarios the ability

to constantly monitor terrorist activity is essential to provide quick and effective counter-measures. Moreover, since the research in this chapter is based on open source materials the IR tool developed provides a computationally cheap method without security constraints to analyse terrorist activity. The IR tool developed in this chapter also has a flexible search function meaning it can easily be used to gather data on emerging terrorist threats.

Although the location and textual data from the IR method were not as close to the ACLED data there is still some evidence that such an approach can be useful. An immediate extension of the research in this chapter would be to study natural language processing methodologies and to adapt them to generate a tool which can more accurately extract textual information from many different website layouts. This would likely provide a much richer automated dataset and would improve the depth of location and description data of events greatly. Again this additional automated data could be integrated directly into mathematical models yielding more in-depth real time event analyses.

CHAPTER 6

CONCLUSION AND FUTURE RESEARCH

The aim of this thesis was to demonstrate the effectiveness of mathematical modelling as a method of studying terrorism and providing new insights. In particular, via the the translation of the theory of crime science into mathematical language it was hoped that novel approaches to understanding terrorist activity with open sourced data could prove to be a fruitful area of research.

Initially in Chapter 1 the foundations of crime science and terrorism literature were provided to give contextual meaning to the thesis. Within this literature attention was drawn to two main groups - the Provisional Irish Republican Army (PIRA) and Al Shabaab (AS). These organisations provided the case studies which were used to test the models constructed throughout the later chapters.

Alongside a review of the literature in Chapter 1 was a brief introduction to the topics which would be covered in the thesis. The first of these areas was introduced as a study of the Hawkes process. Specifically, the Hawkes process was described as a self-exciting point process capable of capturing dynamics associated with an initial shock followed by a sequence of after-shocks. A foundation for applying this type of model had already been demonstrated in previous work by this thesis author, where, the Hawkes process had shown potential for capturing the temporal dynamics associated with the IED attacks conducted by the PIRA. One of the drawbacks of this previous research was its reliance on sociological theory as a basis for conducting analysis between the various phases changes observed with the PIRA organisation. To overcome this limitation in this thesis the research was extended so that the mathematical approach could be conducted independently from the predetermined boundaries. By considering the problem of change point detection in terms of an

optimisation problem it was shown in Chapter 2 that new insights could be obtained regarding the PIRA organisational structure and its attack patterns. To further this mathematical investigation a new sociological study could be conducted to try to reconcile the findings of the found change points with the broader theory surrounding the PIRA.

The Hawkes process studied in Chapter 2 had a simple constant background rate which it was found did not provide a good fitting model in the context of Al Shabaab attacks in Somalia. Therefore, to gain a further level of numerical detail in Chapter 3 a Hawkes process with time dependent background rate was analysed. Based on observed wet-dry weather seasonality in Somalia an attempt was made to model the influence of environmental factors on the attack patterns of the Al Shabaab group. The specific form used for the background rate was a Fourier expansion model for the underlying rate of attacks. Although some improvements were observed over the Hawkes process with a constant background rate the seasonal Hawkes process studied in Chapter 3 still failed to fully capture the underlying dynamics of the AS attacks. This result warrants further investigation of different forms of point processes in the modelling of the AS group. One possible extension would be to experiment with different kernels as opposed to the exponential decay utilised in the Hawkes processes of this thesis.

In the previous chapters the Hawkes processes studied were fitted only to temporal data with spatial information being encoded implicitly via the grouping of events according to geographical boundaries. To study the spatial nature of terrorism more exactly in Chapter 4 a Dirichlet Process Mixture Model (DPMM) was introduced.

The DPMM was used in an attempt to find a mathematical approach to determining the locations of PIRA bomb making factories based on the observed locations of IED events. Moreover, by using the timestamps of factory closures as temporal boundaries the DPMM was used to study the effects of these closures on terrorist activity in Northern Ireland. From the results of Chapter 4 it can be seen that the DPMM had some success in matching areas where bomb making plants were found. Moreover, in some cases the model discovered likely areas for bomb factories prior to their actual closure. These positive results provide evidence that the DPMM is effective at locating sources of IED manufacturing. In an applied setting such a model could be useful for highlighting areas most likely to be the source of terrorist IED factories thus reducing the required resources to trace and close down terrorist operations.

One of the main drawbacks of the research conducted using the DPMM was a lack of detail concerning the relationship between factories and IED attacks. In particular, since it was not possible to remove events according to each factory's closure the results obtained may have been altered by historical dependence in the data. Therefore, more precise data about the factories and IED events may greatly improve the accuracy of the results obtained with the DPMM. Another approach to improve the research of Chapter 4 would be to include a temporal component inside the model. Specifically, temporal dynamics have been treated implicitly with the use of factory closure timestamps. An explicit temporal treatment within the DPMM could allow more interesting findings to be obtained concerning the space-time evolution of the PIRA.

For the final research chapter of this thesis an important topic across all branches

of numerical studies has been explored. It was explained in Chapter 5 that data, and the methods used to collect and analyse its contents, have undergone a rapid evolution which is only set to increase in the future. In particular, with high volume and velocity “Big Data” provides many opportunities to expand human knowledge immensely in many different regards. However, pitfalls associated with analysing large datasets must be carefully guarded against. This is especially relevant in areas of critical importance, such as, terrorism studies whereby lives may be potentially put at risk by pursuing poor policy. Nonetheless, it was shown in Chapter 5 that tools can be constructed to provide information about terrorism automatically with the additional benefit of real-time analysis.

Comparing event timings found from this automated tool and those from a manually collected database, using the method of dynamic time warping, demonstrated a good match. Moreover, utilising the term frequency - inverse document frequency, there is evidence to suggest spatial information can also be obtained automatically with some degree of accuracy. Less accuracy was found when comparing event descriptions of the automatically and manually collected data. A possible explanation for the fall in accuracy compared to temporal data is that the automatic data collection method was only retrieving the first sentence of each online article. Due to the varying nature of website designs it can be difficult to find a single algorithmic approach to finding article text. Therefore, there is scope to extend the research in Chapter 5 by employing more sophisticated natural language processing techniques to enhance the detail of automated database recordings.

This thesis has highlighted several significant contributions mathematical mod-

elling can make to the field of terrorism research. However, it is also necessary to consider the limitations of the types of approaches which have been presented.

One of the major problems surrounding modelling techniques is accurately detecting the underlying distribution of the data being analysed. In particular, throughout each chapter in this thesis assumptions have been made concerning the precise analytical forms of the models studied. For example, within the Hawkes point process model an exponentially decaying kernel is used to model PIRA and AS attacks. Although the model choice is supported by current literature and statistical techniques other models may exist which have even greater potential to capture terrorism dynamics. Hence, the research in this thesis could likely be advanced by considering more general methods of model fitting. In particular, the field of machine learning is presently a very active area of research which can provide methods for finding generalised model frameworks which adapt to data inputs.

In addition to choosing the best models in each chapter the research outcomes also depend on the data utilised. When conducting terrorism research using only open source data there are limitations on attainable insights. This observation was made in relation to the DPMM in Chapter 4, where, allocating IED factories to sources was not possible. A similar situation may also be present in other chapters. In particular, the Hawkes process models the ability of past events to trigger future occurrences. Without further information attributing attacks to particular cells within PIRA or AS it was not possible to determine precisely which events were related. With this extra detail more refined research outcomes could have been achieved highlighting the size of the influence of different parts of terrorist groups on the overall attack

profile. Such findings have the potential to enable counter-terrorism practitioners to allocate resources most efficiently and effectively. These types of issues would likely be overcome in a real world scenario, where, practitioners have access to more detailed closed source datasets.

Alongside the data analysis and modelling results in each chapter is a discussion of the research outcomes. When presenting these discussions the arguments presented mainly focused on linking the quantitative findings to the present literature. However, the models and techniques developed in this thesis may have lead to new insights if considered from a social science perspective. An example of the unique insights offered by this thesis was observed in Chapter 2. Specifically, the change point detection methods developed using the Hawkes process model resulted in new dates for the phases of the PIRA. Therefore, engaging more widely the social science community with the research presented in this thesis would likely lead to much greater depth from the analysis outcomes.

In concluding this thesis it is hoped that the mathematical modelling utilised to study many different aspects of terrorism, terrorist groups and counter-terrorism strategies has demonstrated that mathematics has great potential in this field of study. Going forward the expansion of this type of approach in terrorism and crime science literature will enhance many future research endeavours and provide new and profound insights. In addition, extending these types of mathematical approaches beyond the confines of academia and into the tools employed by practitioners will likely prove to be highly beneficial and offer valuable depth of understanding.

APPENDIX A
SPATIO-TEMPORAL PATTERNS OF IED USAGE BY THE PROVISIONAL
IRISH REPUBLICAN ARMY (PUBLISHED PAPER)

Spatio-temporal patterns of IED usage by the Provisional Irish Republican Army

STEPHEN TENCH^{1,2}, HANNAH FRY² and PAUL GILL¹

¹ *UCL Jill Dando Institute of Security and Crime Science, 35 Tavistock Square, London, WC1H 9EZ*
email: stephen.tench.13@ucl.ac.uk

² *UCL Centre for Advanced Spatial Analysis, 90 Tottenham Court Road, London, W1T 4TJ*

In this paper a unique dataset of improvised explosive device (IED) attacks during “The Troubles” in Northern Ireland (NI) is analysed via a Hawkes process model. It is found that this past dependent model is a good fit to IED attacks yielding key insights about the nature of terrorism in NI. We also present a novel approach to quantitatively investigate some of the sociological theory surrounding the Provisional Irish Republican Army (PIRA) which challenges previously held assumptions concerning changes seen in the organisation. Finally we extend our use of the Hawkes process model by considering a multidimensional version which permits both self and mutual-excitations. This allows us to test how the PIRA responded to past IED attacks on different geographical scales from which we find evidence for the autonomy of the organisation over the six counties of NI and Belfast. By incorporating a second dataset concerning British Security Force (BSF) interventions, the multidimensional model allows us to test counter-terrorism (CT) operations in NI where we find subsequent increases in violence.

1 Introduction

Terrorism is a major international concern which shows little signs of abating. There is therefore great importance in developing scientific approaches to understand the behavioural underpinnings of terrorism in order to prevent and disrupt these activities. The ability to gain such insights through real world experimentation is questionable due to the risks associated with unsuccessful approaches. One of the cheapest and most adaptable methods of research in this area is mathematical modelling [33]. Such modelling provides not only a vast number of well-developed tools and techniques but also the opportunity to experiment freely without unnecessary safety risks or ethical concerns.

Several advances have been made in this field of late, leading to a variety of conclusions with policy making implications. Braithwaite and Johnson [7] studied the interactions of insurgent attacks and Coalition counter-insurgency operations in Iraq. The authors were able to conclude from space-time patterns that indiscriminate counter-insurgency operations resulted in a backlash effect by insurgents whilst discriminate operations had the opposite effect. Along similar lines Lewis et al. [29] apply self-exciting point process models to study violent civilian deaths in Iraq during the U.S.-led invasion. They found a two to six

lines Lewis et al. [29] apply self-exciting point process models to study violent civilian deaths in Iraq during the U.S.-led invasion. They found a two to six month timescale for violent deaths which correspond to a series of related attacks. Hence, quick interventions could help to drastically lower the problem of violent deaths in Iraq. In a similar spirit a study by Mohler [35] found evidence that terrorist attacks in Northern Ireland followed a pattern of self-excitation lasting 9.3 weeks. The line of investigation taken by these sorts of studies have important consequences for tackling the types of terrorism seen in NI where insurgency and civilian deaths were major issues [20]. These approaches marked a large departure from political science-inspired methodologies that generally linked the quantity of terrorist attacks to “root causes” like socio-economic indicators that are quasi-static and fail to provide insight into the triggers needed for a strategic intervention.

The methods employed in these terrorism studies [7, 20, 29, 35] along with several others [14, 22, 55] share similarities with the modelling of spatio-temporal phenomena in crime pattern theory [8]. In that context it is assumed that criminal activity forms a series of quantifiable patterns at the macro scale [16, 44]. Mathematical descriptions of these patterns, or more generally ‘crime hotspots’, can be explored and exploited in real-time. By targeting susceptible areas with preventative measures there is potential for great reductions in subsequent crime. This approach has since been shown to be effective in a number of real-world applications including policy [2, 17, 31], and predictive policing [6, 27].

It is from this point of view that we aim to approach the issue of terrorism during the conflict in NI. In this paper we seek to add to the literature concerning spatio-temporal patterns of terrorism by studying a unique dataset of IED attacks in NI between 1970-1998. Specifically for the case of NI this paper is focused on the group known as the Provisional Irish Republican Army (PIRA). Although there are extensive historical accounts and a growing body of social science research related to this group [3, 19, 25, 26, 51] there is a gap for a wider scope of mathematical investigations of their activities.

The PIRA was predominantly formed from members of the Catholic community in NI [56] and saw itself as “the legal representatives of the Irish people, [who] are morally justified in carrying out a campaign of resistance against foreign occupation forces and domestic collaborators” [39]. The active period of the PIRA between 1969-1998 can be traced out in five phases [3]. These phases and their historical context are described below.

- **1969-1976 - Phase 1:** During this phase the organisation was arranged in a military style having brigades, battalions and companies.
- **1977-1980 - Phase 2:** A cell-based structure was adopted. This approach was characterised by PIRA fracturing into small groups of members known as Active Service Units (ASUs) [25]. This approach aimed to improve the organisation’s secrecy by making it harder to infiltrate.
- **1981-1989 - Phase 3:** The Republican campaign moved into the political arena

through the Sinn Féin party who now had similar levels of prestige as their militant wing, PIRA.

- **1990-1994 - Phase 4:** Secret meetings involving top ranking PIRA leaders negotiating a ceasefire with the British Government.
- **1995-1998 - Phase 5:** Peace talks announced with a ceasefire ratified in the Good Friday Agreement signalling for many the end of “The Troubles”.

During its active phases the PIRA successfully developed a large arsenal of IEDs which it employed with devastating effects for both the security forces and civilians [53]. In particular, the impact of violence was felt heavily by the civilian population which constituted approximately 54% of all deaths [20]. Moreover, the deaths caused by the NI conflict were highly concentrated in Belfast, where approximately 47% of fatalities occurred [20].

In analogy to the studies of Braithwaite and Johnson [7], Lewis et al. [29] and Mohler [35] we make an attempt to understand the driving forces behind IED attacks in NI. Our extension to the present literature revolves around the access we have to a unique dataset of IED events which allows us to study a specific type of terrorist activity in great detail at a fine temporal scale. The dataset we use provides ample evidence for any past dependence on insurgent attacks and this study attempts to understand these dynamics in greater detail. The model chosen to explore this question is a Hawkes self-exciting point process. It makes use of a response function (or kernel) which holds information pertaining to the long-term influence of previous events and has been shown to well represent a number of past dependent processes including gang related violence [18, 24, 52], email exchanges to infer organisation leadership [21], burglary [34, 50] and violent deaths in conflicts as previously mentioned [29]. It also provides an opportunity to examine each of the five phases of PIRA activity separately, as is done in Model 1 and all subsequent models, offering direct quantitative insights into how the group behaved and reacted through the stages of the Republican campaign.

As well as studying temporal patterns of insurgency Lewis et al. [29] also compare Hawkes processes in different regions of Iraq to understand spatial influences seen during waves of violent attacks. Similarly in this paper we undertake spatial disaggregation of PIRA attacks according to the six counties of NI and we also separate Belfast due to its significance during “The Troubles” as discussed above [20]. This additional spatial information is to yield insights about the extent to which PIRA units in NI acted autonomously as suggested by Horgan and Taylor [25]. This latter study forms the theoretical basis for Models 2 and 3.

The PIRA did not act in isolation however, and in this contribution, we also aim to explore the interplay between PIRA and the British Security Forces (BSF) by employing a multidimensional Hawkes process. Previous studies into British counter-terrorism (CT) strategies in NI [12, 28, 40] found evidence that the actions of BSFs could undermine the effort to curb Republican terrorism and even result in a negative backlash increasing the number of attacks. For instance, in 1988 an operation by BSFs resulted in the deaths of three PIRA

members in Gibraltar. An analysis of subsequent PIRA attacks found positive increases related to the incident 36 months after it occurred [28]. The inclusion of additional mutually exciting terms in the Hawkes process has been seen to represent other interacting systems in the past including multiple gang networks [50], but, to the best of our knowledge, has yet to be applied to yield insights into CT strategies.

To test the CT strategies employed in NI two types of events will be assessed in this paper. Firstly, from the discussion above concerning the background of the PIRA it was pointed out the organisation drew the majority of its volunteers from the Catholic community [56]. Consequently deaths of Catholic civilians resulting from BSF operations may be expected to prompt a significant retaliation from the PIRA. We investigate this effect in Model 4. Moreover, following the case study findings of Lafree, Dugan and Korte [28], as discussed above, the backlash effect of BSF actions which killed PIRA members will be considered in Model 5. With such insights we aim to uncover a useful methodology by which CT practitioners and academics can judge the efficacy of past strategies to combat terrorism.

Alongside the contributions this paper aims to make concerning the use of the Hawkes process we also present a novel approach to deal with the issue we refer to as edge effects. This effect is the result of events outside the observation period influencing those inside [46]. Whereas in previous studies using the Hawkes process the data analysed is in a single time series [18, 29] with the segmentation of our data according to the five phases of PIRA it is possible that events in previous phases may have influenced those in future ones. As a consequence of this a moving time window approach was considered whereby the data points from adjacent phases were combined to find the quantitatively best fitting model.

The structure of presentation will be in 6 sections. In Section 2 a discussion of the datasets used for this paper will be provided. Then in Section 3 the mathematical models that have been studied will be introduced in more detail. This section will also contain further information about the method used to obtain model parameters. Next in Section 4 the numerical results of the paper will be provided. This will begin with a description of the novel approach we have taken to examining edge effects in the data analysed. Alongside the results a discussion will be provided about how the model parameters can be interpreted. Finally in Section 5 the findings of the paper will be discussed with their potential impacts for both the academic community and practitioners demonstrated as well as a brief overview of future research topics.

2 Data

This paper utilises a unique dataset of PIRA IED events from 1970 to 1998. This dataset was collected through an exhaustive coding of newspaper reports and other open source outlets. Please see Asal et al. [3] for a full outline of the data collection and verification process. In total, the dataset spans 5461 IED

Table 1. *PIRA IED Dataset Event Fields*

Field	Values
Date	Year (1970-1998)/Month (1-12)/Day (1-31)
Location	{Antrim, Armagh, Belfast, Derry, Down, Fermanagh, Tyrone}
Target	{Political, Military, Police, Paramilitary, Government, Transport, Civilian, Foreign}

Table 2. *BSF Dataset Event Fields*

Field	Values
Date	Year (1970-1976)/Month (1-12)/Day (1-31)
Religion	{Catholic, Protestant}
Status	{Civilian, British Security, Republican Paramilitary, Loyalist Paramilitary, Irish Security}
Organisation Responsible	{British Security, Republican Paramilitary, Loyalist Paramilitary, Irish Security, Unknown}
Geographical Location	{Belfast North/East/South/West, County Antrim/Armagh/Derry/Down/Fermanagh/Tyrone, Derry, Britain, Europe, Republic of Ireland}

events. For each event there are numerous details concerning the IED attack and groups involved. Details of the event fields relevant to our study are given in Table 1. It should be noted here that when discussing BSFs we refer to both military and police targets. For further details on all the information contained in the dataset the reader is referred to Asal et al. [3].

Alongside the study of PIRA related attacks in isolation, an additional investigation was made into how BSF attacks impacted upon further PIRA attacks. The dataset concerning BSF attacks was obtained from the Conflict Archive on the Internet (CAIN) [53]. In Table 2 the event fields and possible values for this dataset are presented. In this study only the events which occurred in NI were considered. There were a total of 131 Catholic civilian deaths recorded in the final dataset of BSF events. Of these entries 78 were found to correspond to Phase 1 of PIRA activity. This lead to a only a small number of data points being available in the other phases with 12, 34 and 7 points for Phases 2, 3 and 4 respectively and 0 for Phase 5. Due to this distribution of data, results were only obtainable for Phase 1. The number of PIRA IED events in Phase 1 targeting BSFs was 144. See Sections 4.8 and 4.9 for further details.

3 Methodology

3.1 Poisson Process

The first model studied serves as a baseline to compare with the Hawkes process [18]. The specific baseline model tested was a Poisson process model specified via a single parameter N/T , where N = number of events and T = time of the final event measured from time 0 [47]. The Poisson process assumes that each event is independent and thus that the system has no memory of attacks in the past.

3.2 Univariate Model

Moving to a more complex model this paper will consider the influence of self-excitations of PIRA IED attacks. The method we used is based on studying an intensity function which describes the rate of IED incidents as a function of time and conditional on the past history of events. The intensity function takes the following form [23] for a given set of event times $\{t_i\}_{i=1}^N$.

$$\lambda(t) = \mu + k_0 \sum_{t > t_i} g(t - t_i; \omega) \quad (3.2)$$

The response function g is taken to be of the form $g(t) = \omega e^{-\omega t}$. The exponential form for the response function is routinely used in studies of crime and insurgency data [18, 29, 52]. It makes good intuitive sense for events which are clustered in time and allows for a physical interpretation of each component of the intensity function. In the example of IED attacks the constant μ can be considered as a background rate at which IED events occur. After an initial IED attack there may be further attacks, for example, a PIRA unit may wish to follow up on the success of a previous attack, and the constant k_0 captures the jump in the IED event rate. However, an indefinitely higher rate is unrealistic and eventually the rate will return to the background rate. The rate of decay is controlled by the term ω . The additional ω preceding the exponential term acts as a normalisation constant so that the jump factor multiplied by the response function can be viewed as the number of offspring after an event and the density of the time interval for the increase in activity [46].

Parameter estimation for the intensity function can be undertaken via the method of maximum likelihood estimation (MLE) [42]. This process involves finding the parameters which maximise the following log-likelihood function derived by Rubin [48]:

$$\log L(\{t_i\}; \mu, k_0, \omega) = \sum_{i=1}^N \log(\lambda(t_i)) - \int_0^T \lambda(t) dt,$$

where $t_N = T$ will be taken as the final time of observation in a similar approach taken by Ozaki [42]. For the form of the intensity function given in

(3.2) the log-likelihood becomes

$$\log L = \sum_{i=1}^N \left[\log \left(\mu + k_0 \sum_{t_i > t_j} \omega e^{-\omega(t_i - t_j)} \right) + k_0 \left(e^{-\omega(T - t_i)} - 1 \right) \right] - \mu T.$$

A number of assumptions are built into this formulation. First, all parameters used in the intensity function should be positive [29] to ensure the model remains realistic. Second, the set of points $\{t_i\}_{i=1}^N$ should be measured from time zero. However, since the Hawkes process depends on the infinite past, this assumption is not achievable in a real world setting and it may be difficult to eradicate the influence of events outside the observation period on those inside [46]. More details on how this problem was handled for this paper will be discussed in Section 4.1.

A further assumption is that the set of times should be unique [30]. Our dataset contains several simultaneous events and thus to satisfy the requirement for uniqueness, events in the same county, or multiple events in Belfast, on the same day were regarded as a single event. However, to avoid losing too much detail, events in different counties, or events inside and outside of Belfast, on the same day were distinguished via the addition of a random timestamp (as in Bowsher [5]). This is justified by looking at the command and functional structure of the PIRA which reveals that at the county and Belfast levels IED attacks were fairly autonomous [25].

Finally, a restraint on the response function g ensures that the model is non-explosive (see Varadhan [54] for further details concerning explosive stochastic processes). This assumes that the integral of g over t should be strictly less than unity [42]. Our choice of g satisfies this condition.

3.3 Multidimensional Model

After examining the past dependent nature of IED attacks based on self-excitations the second type of model investigated will also include mutual-excitations. For example, such models will be used to consider the influence of PIRA attacks and BSF attacks on further PIRA attacks.

The multidimensional Hawkes process model can be defined in a similar way as was done for the one-dimensional case. Here, with two adversaries, we require a two-dimensional model. Now there are two sets of event times which will be labelled $\{t_a\}_{a=1}^N$ and $\{t_b\}_{b=1}^M$ and two counting processes, $N_r(t)$, $r \in \{1, 2\}$, which form a two-dimensional counting process $\mathbf{N}(t)$. Each individual process has intensity function defined by [23]

$$\lambda_r(t) = \mu_r + k_0 \sum_{t > t_a} g(t - t_a; \omega) + s_0 \sum_{t > t_b} h(t - t_b; \nu), \quad (3.3)$$

where the two response functions are defined by $g(t) = \omega e^{-\omega t}$ and $h(t) = \nu e^{-\nu t}$. The form chosen for the response functions is chosen by analogy to the research of Short et al. [50]. In particular, in extending the work of Egesdal et al. [18],

to study interactions between multiple gangs, results were obtained indicating that exponential response functions may prove useful to study mutual-excitations between gangs.

Similar to the univariate case described in Section 3.2 we can interpret the model in a real world setting. In the case of modelling influences on PIRA attacks one could, for example, take event times $\{t_a\}$ to represent times of PIRA IED attacks and times $\{t_b\}$ as BSF attacks. The background rate μ_r , jump rate k_0 and response function g have the same interpretation as that given in the one-dimensional case. Similarly the parameter s_0 represents the jump in IED attacks following a mutual-excitation, such as, retaliation against a BSF attack whilst ν controls the temporal scale over which this mutual-excitation persists. The parameter ν also acts as a normalisation constant for the response function h so that the product of the jump factor s_0 and the response function h can be interpreted as the number of offspring events and the density for the increase in activity following a mutually exciting incident [46]. In this two-dimensional system the second intensity function would model the influence of past BSF and PIRA IED attacks on BSF attacks.

To compute parameter estimates the MLE can again be employed. This MLE takes the following form [50]

$$\log L(\{t_a\}; \mu, k_0, \omega, s_0, \nu) = \sum_{a=1}^N \log(\lambda_1(t_a)) - \int_0^T \lambda_1(t) dt,$$

where, $T = \max\{\max_a\{t_a\}, \max_b\{t_b\}\}$. A similar formula holds for $\lambda_2(t)$ and $\{t_b\}$.

As in the one-dimensional case, all parameters must be positive [23] to make the model realistic and events in the infinite past should be considered [46]. Likewise, event times must be unique [30]. The condition necessary for the model to be non-explosive is reformulated in the higher dimensional case. In particular, consider the 2x2 matrix \mathbf{G} whose entries are formed of the integrals

$$\int_0^\infty g_{rs}(t) dt, \quad r, s \in \{1, 2\},$$

for each response function occurring in the definitions of λ_r . Then the condition for the model to be non-explosive is that the spectral radius defined as

$$\rho = \max_i \{|e_i|\} < 1,$$

where e_i represent the eigenvalues of \mathbf{G} [9]. Again this assumption has been checked and found to be satisfied for the models studied in this paper with the exception of Model 5 (see Table 3 below) which has a spectral radius of 1.0124. This case should therefore be treated with care and may be the result of a small dataset for this model (see the model analysis in Section 4.9 for further details).

Table 3. *Models*

Model Number	Event Times Interpreted	Dataset
0	Number of $\{t_i\}$ = IED events in NI	PIRA Events
1	$\{t_i\}$ = IED events in NI	PIRA Events
2	$\{t_i\}$ = IED events in Belfast	PIRA Events
3	$\{t_a\}$ = IED events in Belfast $\{t_b\}$ = IED events in the six counties of NI	PIRA Events PIRA Events
4	$\{t_a\}$ = IED events targeting BSFs in NI $\{t_b\}$ = BSF events which killed Catholic civilians	PIRA Events BSF Events
5	$\{t_a\}$ = IED events targeting BSFs in NI $\{t_b\}$ = BSF events which killed PIRA members	PIRA Events BSF Events

3.4 Table of Models

Having given the general form of the models in this paper the specific models that were studied are summarised in Table 3. In particular, this table presents an interpretation of the event times used in each model. For clarity the datasets being studied in each model are also listed corresponding to the information in Tables 1 and 2. Times t_i correspond to the univariate model presented in Section 3.2 whilst times t_a and t_b correspond to the multidimensional model from Section 3.3. In Section 4, where the numerical results of this paper are provided, the order of presentation will correspond to the ordering of models shown in Table 3.

3.5 Computational Methodology

Finding the parameters which maximise the log-likelihoods can be undertaken in numerous ways [42]. For this paper optimisation of the log-likelihood functions were undertaken in the Python programming language using the SciPy Optimize package Nelder-Mead [37, 49]. The Nelder-Mead algorithm was chosen based on previous observations of its effectiveness when applied to point process models [45] and also its performance during preliminary coding. It should be noted at this point that this optimisation procedure finds the minimum value, hence the equivalent problem of finding the minimising parameters of $-\log L$ was considered. To obtain further computational efficiency a recursive algorithm described by Liniger [30] was used to compute values of the intensity function.

Another important point, which is made by Egesdal et al. [18], is that due to the nonlinear nature of the minimisation it is not guaranteed that a global minimum will be found. Therefore, there is a need to begin the optimisation

procedure at multiple points and take the parameters yielding the lowest value of $-\log L$ and subject to the conditions given earlier.

Having derived each model's parameters we then go on to assess its goodness of fit. To determine the overall model fit residual analysis was employed. The basic ideas of this approach can be found in Brown et al. [10] and are also summarised below.

Consider a point process formed of the set of event times $\{t_i\}$ with intensity function λ . Perform the following integrals which transform the set $\{t_i\}$ to the set $\{\tau_i\}$

$$\tau_i = \int_0^{t_i} \lambda(t) dt.$$

If the model is a good fit then the residuals $\{\tau_i\}$ are independent and distributed according to a stationary Poisson process with unit rate [43]. Therefore, the inter-arrival times given by

$$Y_k = \tau_k - \tau_{k-1}$$

are exponentially distributed. Setting $\tau_0 = 0$ [10] and applying these procedures for the intensity function given in (3.2) the following formula is obtained

$$Y_1 = \mu t_1,$$

$$Y_k = \mu(t_k - t_{k-1}) - k_0 \sum_{i=1}^{k-1} e^{-\omega(t_k - t_i)} - e^{-\omega(t_{k-1} - t_i)}, \quad 1 < k \leq N.$$

If the inter-arrival times are exponentially distributed then

$$U_k = 1 - e^{-Y_k} \tag{3.5}$$

form a set of independent uniform random variables over $[0, 1)$. Therefore, to test the goodness of fit of the Hawkes process it remains to determine if the corresponding U_k do indeed come from a uniform distribution.

A quantitative test that can be used to check this assumption on the distribution of the U_k values is the Kolmogorov-Smirnov (KS) test [32]. The KS test in this case works by comparing the value of the test statistic $D_n = \max_k (|U_k - \frac{k-1}{N}|, |\frac{k}{N} - U_k|)$ [57] to a critical value D_α (see O'Connor and Kleyner [41] for a table of critical values). Statistical significance is obtained if the condition $D_n < D_\alpha$ is found to hold in which case there is evidence to suggest goodness of fit of the model.

Another method we used, which compares the fit of different models, is the Akaike Information Criterion (AIC) [1]:

$$AIC = 2k - 2 \log L,$$

where, k is the number of parameters being fitted in the model and $\log L$ is the maximum of the log-likelihood function. The model yielding the lowest value for AIC is deemed the better fit: more parameters are penalised whilst a greater value for the log-likelihood is rewarded. Burnham and Anderson [11] point out the AIC difference is not a significance test in the sense of critical values and

requires some judgement. For general guidance Burnham and Anderson [11] suggest that $0 - 2$ shows little difference between models, $4 - 7$ considerably more evidence for a difference and > 10 is classified as a significant difference. It is important that the AIC should only be used to compare models which are fitted with the same dataset [11], as was done for the comparisons in this paper.

4 Numerical Results

4.1 Determining Phase Boundaries

Initially we will be focused on PIRA IED attacks across all of the counties of NI and Belfast. Each phase of the events as outlined in Asal et al. [3] will be explored separately to offer a quantitative description of the changes in the Republican campaign.

The phases may not be treated in isolation, however. As discussed in Section 3.2, the historical dependence of the Hawkes process means that events outside of the observation period may influence those inside leading to spurious parameter values [46]. Thus, Phase 1 may influence events in Phase 2, Phase 2 may influence Phase 3 and so on. To avoid this edge effects issue, a systematic approach is required to deal with the phase boundaries.

We propose a novel approach to examine the impact of edge effects (see also Nichols and Schoenberg [38] for another similar approach in the field of seismology). Specifically, a moving time window was used to include the influence of events from Phase i in Phase $i + 1$ for $i = 1, 2, 3, 4$. This method is best illustrated via example. Consider Phases 1 and 2. First, we calculated the MLE parameters that result from the dataset consisting of Phase 2 only. Then adding one point from Phase 1 the MLE parameters were recalculated. This was then repeated with two points from Phase 1 and so on until the MLE parameters for all the data in Phases 1 and 2 combined had been calculated. Finally we determined which of these models provided the best fit compared to the others using the KS test. It was decided that a more positive difference $\Delta = D_\alpha - D_n$ gives more certainty that statistical significance has been reached. Thus the choice of the best fitting model, and hence the most sensible phase boundaries to use, was decided based on maximising Δ .

Since Phase 1 was effectively the start of the conflict, and the beginning of the PIRA as an organisation, we do not assume there will be any substantial effects from previous events outside the dataset. Therefore, the boundaries for Phase 1 were unchanged from their original definition as in Asal et al. [3]. In Table 4 the number of points that were required to maximise the difference between the critical value and KS test statistic are given for the remaining four phases. It can be seen that very few points were required to fix the correct mathematical boundaries for Phases 2 and 3 and it was found that with so few points there was little change in the parameter values compared to the unadjusted phases. However, in Phase 4 it was necessary to include 65% of data points from Phase 3 which significantly changed the parameter values from

Table 4. *Edge Effects Results*

Phase	Number of Data Points Added	New Boundary
1	-	27/01/1970
2	11	24/10/1976
3	11	22/08/1980
4	494	11/4/1984
5	41	22/5/1994

those found for the original Phase 4 boundary. This is illustrated clearly by the plots in Figure 1 which show the variation in the values of Δ and the three model parameters as each data point is added from Phase 3 to Phase 4. A similar observation was also made for the case of edge effects in Phase 5.

The boundaries found from this edge effects analyses will be used throughout the remainder of this paper. This ensures consistency and enables comparisons between models. From here on in reference to a model's MLE parameters means the parameters found using the boundaries stated in Table 4.

4.2 Comparing Models

4.3 Table of Results

For brevity, all model results are listed concurrently for each phase in Tables 5 to 7, but care should be taken when making a comparison of the different models. A comparison may be made within each of the three investigations: PIRA events across NI, PIRA events in Belfast and outside, PIRA and BSF events in NI, but not directly between them, except for model fit. See Section 3.5 for more detail.

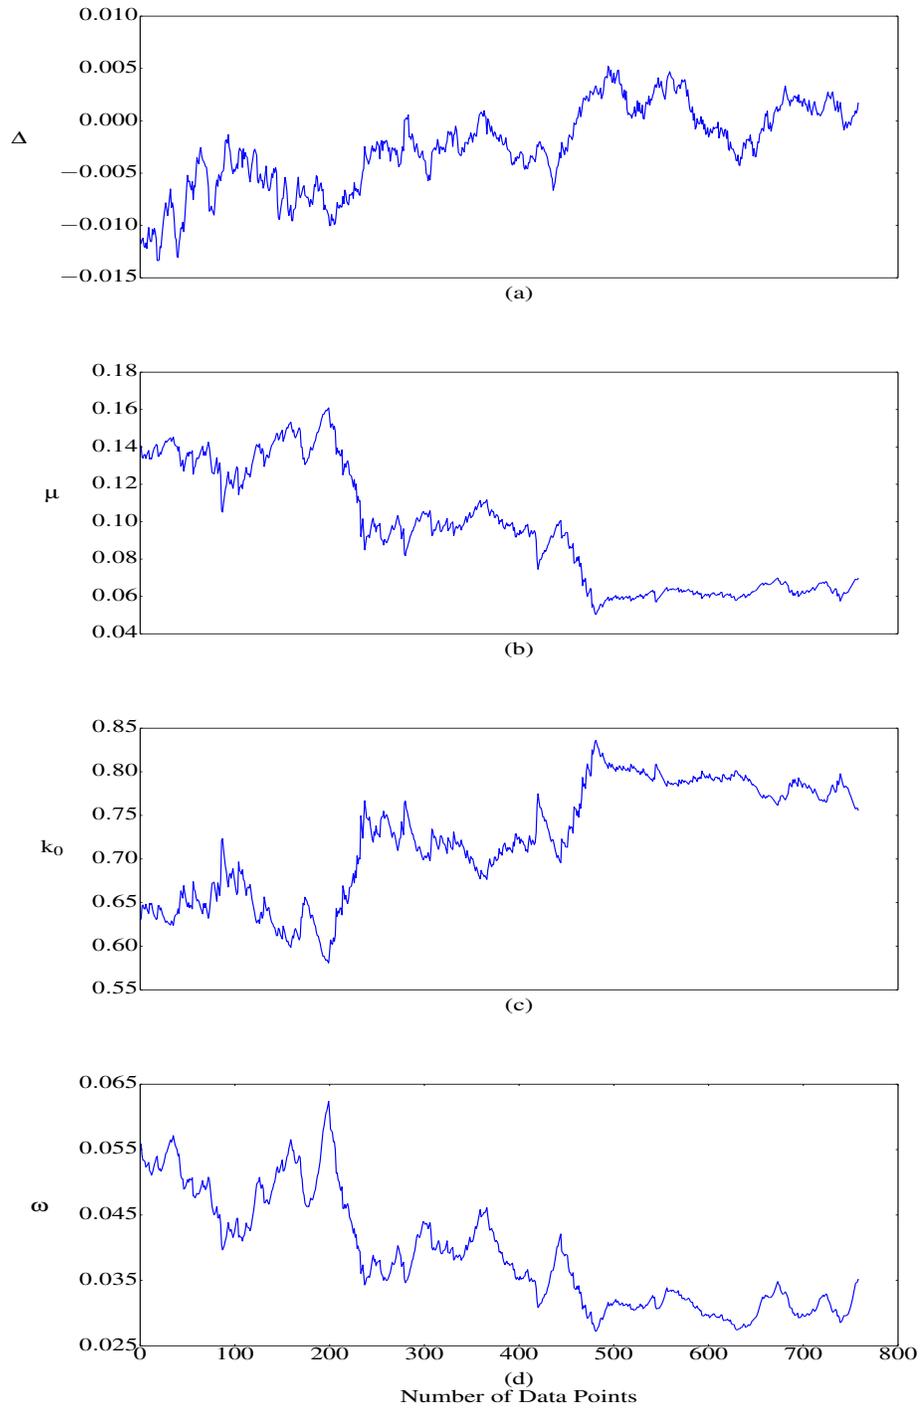


FIGURE 1. Figure (a) shows the values of $\Delta = D_\alpha - D_n$ for each point added from Phase 3 to Phase 4. Figures (b)-(d) show the corresponding changes in the MLE parameter values.

Table 5. *PIRA events across NI.*

		Phase 1	Phase 2	Phase 3	Phase 4	Phase 5
Model 0	μ	0.3020	0.2541	0.2250	0.3028	0.0957
	AIC	3359.6	1784.4	3834.3	5190.5	1079.7
Model 1	μ	0.0543	0.1721	0.0800	0.0597	0.0270
	k_0	0.8241	0.3233	0.6529	0.8040	0.7231
	ω	0.0542	0.7685	0.0426	0.0316	0.0901
	KS Test	0.0686	0.0528*	0.0465*	0.0343*	0.0455*
	KS Critical 95%	0.0492	0.0701	0.0490	0.0396	0.1072
	KS Critical 99%	0.0590	-	-	-	-
	AIC	3083.7	1717.6	3750.5	5004	987

* Significant at 95% level.

Table 6. *PIRA events in Belfast and outside.*

		Phase 1	Phase 2	Phase 3	Phase 4	Phase 5
Model 2	μ	0.0375	0.0397	0.0222	0.0175	0.0178
	k_0	0.7364	0.6456	0.7123	0.8766	0.5925
	ω	0.0246	0.0298	0.0135	0.0103	0.0874
	KS Test	0.0715*	0.0996*	0.0477*	0.0679**	0.0693*
	KS Critical 95%	0.0726	0.1096	0.0878	0.0640	0.1626
	KS Critical 99%	-	-	-	0.0767	-
	AIC	2015.1	967	1724.8	2707.8	556
	Model 3	μ	0.0396	0.0299	0.0189	0.0099
k_0		0.6441	0.5783	0.7126	0.7449	0.3842
ω		0.0272	0.0339	0.0135	0.0138	0.1001
s_0		0.0647	0.0949	0.0208	0.1085	0.2336
ν		0.7840	0.3976	1.1336	0.1099	0.1681
KS Test		0.0561*	0.0961*	0.0667*	0.0562*	0.1745**
KS Critical 95%		0.0726	0.1096	0.0878	0.0640	0.1626
KS Critical 99%		-	-	-	-	0.1948
AIC		2013.6	973	1726.1	2716.4	552.3

* Significant at 95% level. ** Significant at 99% level.

4.4 Model 0

For this model we look at just a Poisson process applied to IED attacks across NI. Even this simple framework manages to capture a difference in rate across the five phases, showing Phase 4 as having the highest probability of a random event in a given time window with $\mu = 0.3028$. This is in contrast to the Hawkes process in Model 1 where Phase 2 is found to have the highest background rate whilst μ is much lower in Phase 4, suggesting that events in Phase 4 were heavily dependent on the past. Historical dependence is also seen for the other phases since the background parameter values found for Model 1 are all significantly lower than those of the Poisson process in Model 0. According to the AIC comparisons Model 1 is also shown to provide a better model fit in each phase.

Table 7. *PIRA and BSF events in NI.*

		Phase 1
Model 4	μ	0.0189
	k_0	0.1138
	ω	0.9716
	s_0	1.0694
	ν	0.0137
	KS Test	0.1122*
	KS Critical 95%	0.1133
	KS Critical 99%	-
	AIC	1057.8
Model 5	μ	0.0212
	k_0	0.2411
	ω	0.2757
	s_0	0.9774
	ν	0.0087
	KS Test	0.1205**
	KS Critical 95%	0.1133
	KS Critical 99%	0.1358
	AIC	1070.4

* Significant at 95% level. ** Significant at 99% level.

4.5 Model 1

The next model studied was a Hawkes process with a single self-exciting term applied to IED events across NI.

4.5.1 Goodness of Fit

Qualitative evidence for the model's ability to represent the data is gained from visualisation of the intensity function. In Figure 2 a plot of the intensity function over time for the Hawkes process in Phase 1 is presented. Here the peaks and troughs of the model are seen to follow closely the patterns of event times observed in the actual data.

The overall goodness of fit of Model 1 in each phase is determined via the KS test results. The critical values for the 95% confidence level have $D_\alpha = \frac{1.36}{\sqrt{N}}$ [41]. For Phase 1 it was found that the KS test statistic exceeded the critical value even when considering the 99% confidence level with $D_\alpha = \frac{1.63}{\sqrt{N}}$ [41]. Hence there is insufficient evidence in this case to conclude that the model is accurately capturing the dynamics of the data.

Although the KS test results for Phase 1 gave a negative result for the model fit there is some evidence for goodness of fit in a KS plot. Following the method outlined in Brown et al. [10] U_k , as defined in (3.5), is plotted against the hypothesized cumulative distribution, evaluated at $\frac{k-0.5}{N}$. If U_k is indeed uniformly distributed, the resulting graph should be a 45° line. The results of performing this procedure are shown in Figure 3 for Model 1 in each phase of the PIRA. Upper and lower bounds are also shown, obtained by plotting the lines $y = x \pm D_\alpha$ [13]. In cases where points do deviate from the best fit

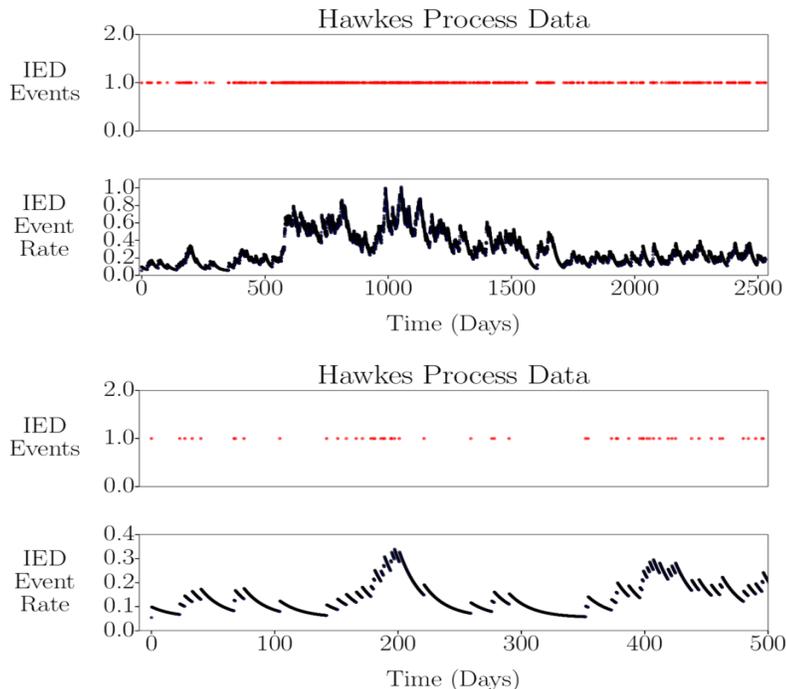


FIGURE 2. Top graphs show event times of IED attacks and the intensity function for Model 1 in Phase 1. The bottom graphs are the same but only for the first 500 days.

line they still remain within or close to the error bounds, suggesting that the Hawkes process generally represents the data well.

A final test is to inspect a plot of U_{k+1} against U_k to ensure that the U_k are independent [4]. If serial correlation occurs it is most likely to be between adjacent time intervals, hence if these plots reveal some patterning it suggests that the transformed times are not independent. The graphs in Figure 4 present this analysis for each phase of the IED data. Although there is some patterning occurring in the plots for Phases 1 and 4 overall these plots appear to show little correlation between the neighbouring points of the sequence $\{U_k\}$. This serves as reassurance that independence exists and the Hawkes process is a good representation of IED attacks.

An AIC comparison to the simple Poisson process in Model 0 shows that Model 1 provides a better fit in each phase.

4.5.2 Interpreting Parameters

The parameter ω^{-1} gives information concerning the average length of time a series of attacks persists. In Phase 1 the average attack window is $\frac{1}{0.0542} = 18.5$ days. The rate of decay then has a large increase in Phase 2 yielding an av-

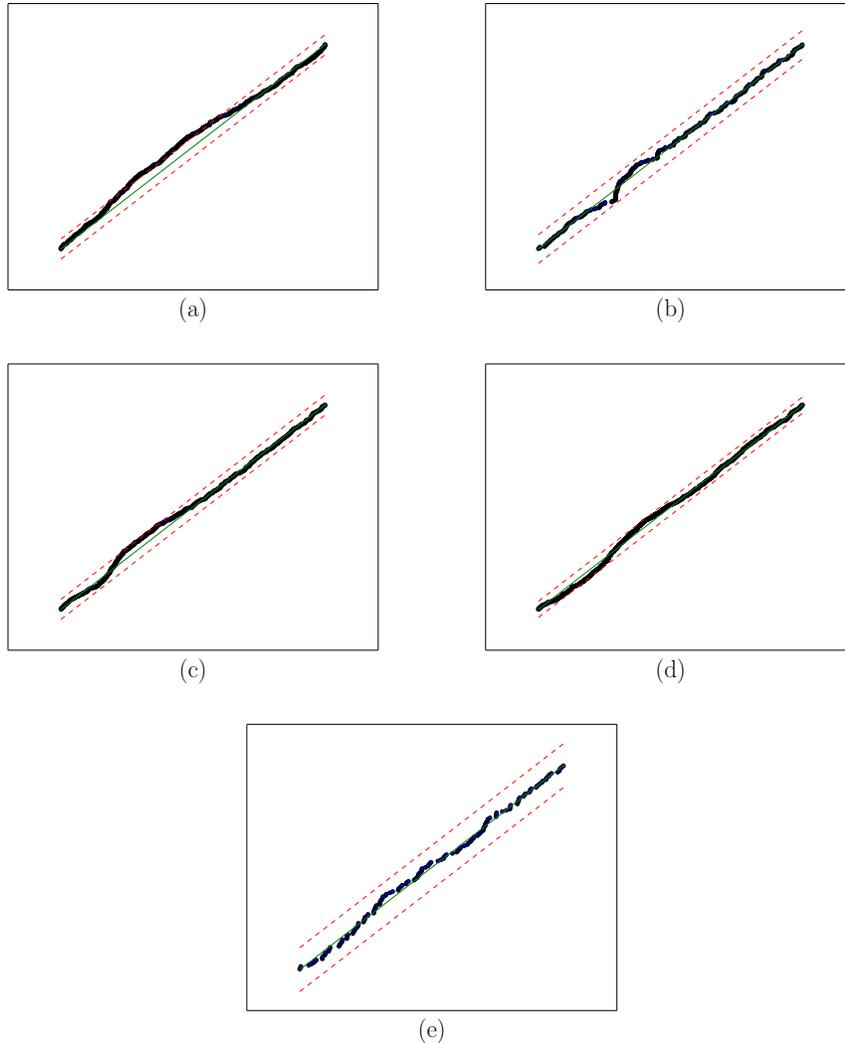


FIGURE 3. Figures (a) - (e) show KS plots for Hawkes process Model 1 in Phases 1-5 respectively. Data points falling on the solid goodness of fit line imply a perfect model fit with the dashed lines representing 95% error bounds.

erage time window of 1.3 days. The work of Asal et al. [3] points out that in Phase 1 there was a more militaristic style of operation within the PIRA which could suggest better attack coordination allowing for longer periods of related waves of attacks. However, this organisational structure made the PIRA susceptible to infiltration by Security Forces thus prompting a shift to a cellular based approach in Phase 2 [3]. One of the consequences of infiltration could be reflected in the shorter attack window, which might be representing the fact

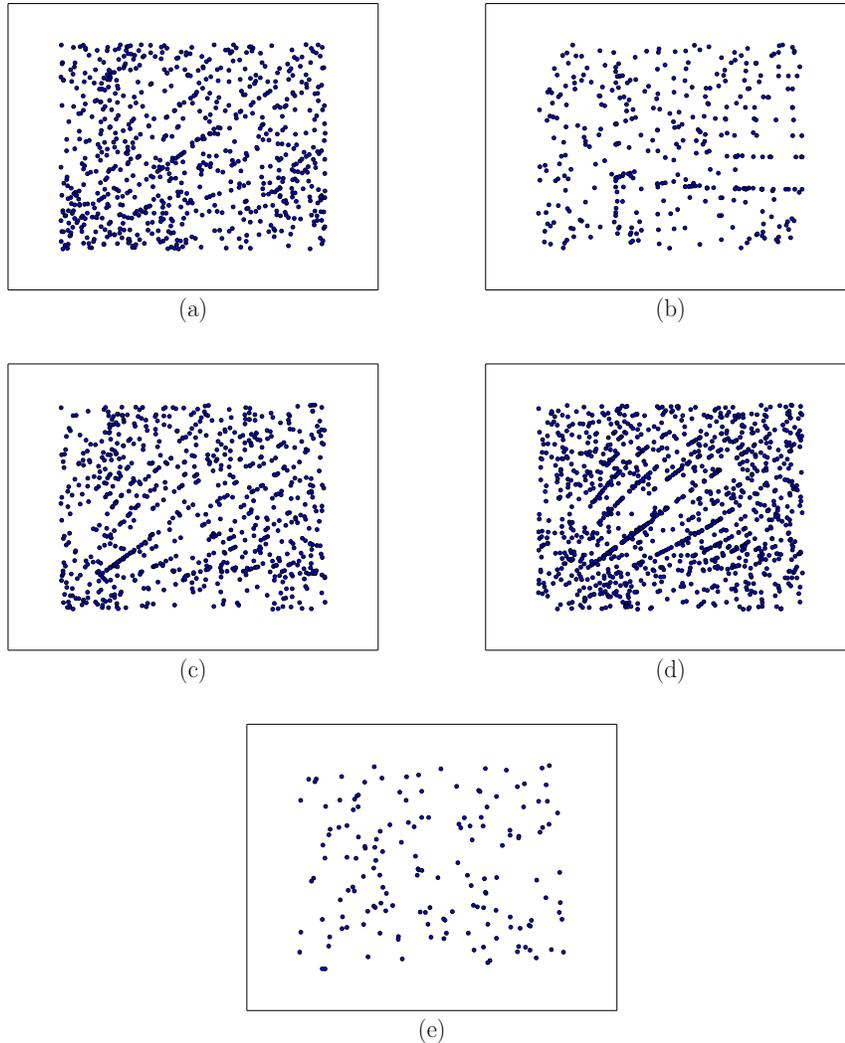


FIGURE 4. Figures (a) - (e) show serial correlation plots for Hawkes process Model 1 in Phases 1-5 respectively. A random dispersion of data points indicates goodness of fit of the model.

that many PIRA members were imprisoned [19], thus, there would have been fewer members to carry out attacks. In Phase 3 the attack window becomes 23.5 days. This phase saw a resurgence of violence by the PIRA with Moloney [36] describing the 1980's as a period of escalating violence similar to the "Tet Offensive" launched by the People's Army of Vietnam in 1968. In particular, the decay rate found could represent the PIRA using sustained attacks to weaken British resolve to remain in NI. This also links to the PIRA objectives described in the organisation's Green Book [39]. One of these objectives was to

use a “bombing campaign aimed at making the enemy’s financial interests in our country unprofitable”. Another objective was to wage a “war of attrition against enemy personnel which is aimed at causing as many casualties and deaths as possible so as to create a demand from their people at home for their withdrawal”.

Phases 4 and 5 then show an increasing trend for the decay rate. In Phase 4 the average time window for attacks is 31.6 days and in Phase 5 it is 11.1 days. These periods were characterised by secret meetings and negotiations that eventually led to the Good Friday Agreement [3]. The PIRA used IED attacks as a bargaining tool with the British Government [3] and as noted by Coogan [15] PIRA had the ability to “turn ... bombing[s] on and off like a tap”. So these shortening periods of IED usage may have been the PIRA using its capabilities as a way to achieve leverage with the British Government during peace negotiations rather than for a war of attrition.

The parameter μ can be interpreted as the background rate at which new events randomly occurred. The trend appears to be for the parameter to increase from Phase 1 to 2 and then fall in the remaining three phases. The rise in Phase 2 may be related to the fact that the organisation of PIRA was shifting and hence with less control members were conducting attacks more randomly. However, as the “Tet Offensive” campaign began in Phase 3 attacks became more systematic. Finally de-escalation of violence over Phases 4 and 5 explain the decreases in the value of μ .

The final parameter k_0 can be interpreted as the jump in the rate of events following an initial event. Phase 1 has the highest value for this parameter. This could again be a result of the military structure of the PIRA leading to more flexibility to escalate events. As before the drop in Phase 2 may be related to imprisonment of PIRA members hindering the extent to which attacks could occur. The rise in Phase 3 may also be interpreted as PIRA adopting a “Tet Offensive” approach and the Green Book objectives both explained above. The final phases have relatively high values of k_0 . This can be linked to the PIRA using IED attacks to demonstrate its capabilities during peace negotiations also as described above.

4.6 Model 2

The work of Fay, Morrissey and Smyth [20] demonstrates that violence during “The Troubles” was highly concentrated in Belfast. Thus in Model 2 we refine the geographical scale of investigation to this region considering a single term self-exciting Hawkes process applied to PIRA IED events.

4.6.1 Goodness of Fit

The results in Table 6 show that there is strong quantitative evidence for Model 2 providing a good fit to the data. The KS plot for the model in Phase 2 is shown in Figure 5(a). It was not felt that the data stayed close enough to the

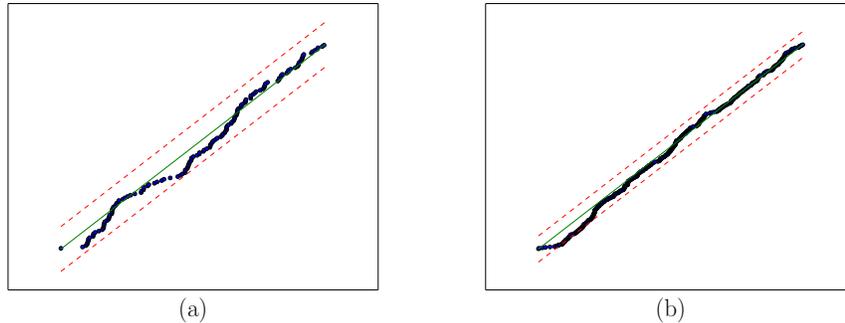


FIGURE 5. Figures (a)-(b) show KS plots for Hawkes process Model 2 in Phases 2 and 4 respectively. Data points falling on the solid goodness of fit line imply a perfect model fit with the dashed lines representing 95% error bounds.

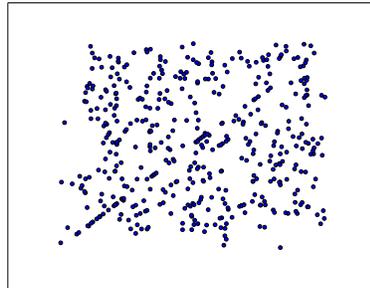


FIGURE 6. Serial correlation plot for Hawkes process Model 2 in Phase 4. A random dispersion of data points indicates goodness of fit of the model.

line of best fit to conclude a good fit of the Hawkes process. For the model in Phase 4 the KS test had to be conducted at the 99% level of confidence to obtain a significant outcome. Despite this the KS plot for Phase 4, as shown in Figure 5(b), gives some qualitative evidence for goodness of fit of the model even at the 95% level. There also appeared to be some patterning in the serial correlation plot for Model 2 in Phase 4 as shown in Figure 6.

4.6.2 Interpreting Parameters

Since Belfast was such a central stage in the NI conflict it is not too surprising that the trends for the parameter values in Model 2 are similar to those seen for Model 1. However, Model 2 in Phase 2 does not share the same significant changes in parameter values as observed for Model 1. This implies the PIRA attacks in Belfast were less susceptible to internal and external changes. One explanation for this observation is the existence of a Northern Command Unit being based in Belfast since 1969 [25]. As such it may have been easier for

PIRA to quickly adapt to internal and external events, such as, the mass imprisonment of PIRA members which led to the organisation becoming cell-based in Phase 2 [19].

4.7 Model 3

The next model considered was a multivariate Hawkes process. With this model we aimed to capture the influence on PIRA IED attacks in Belfast based on self-excitations of past PIRA attacks in Belfast and mutual-excitations of past PIRA attacks in the six counties of NI.

4.7.1 Goodness of Fit

It can be seen from the results table presented earlier that the increased complexity from using a multivariate Hawkes process in Model 3 does not yield a better fitting model compared to Model 2. This may be due to the autonomy of ASUs in the counties of NI and Belfast [25].

Nonetheless, in all but one case there is quantitative evidence for goodness of fit of the models. For the model in Phase 5 goodness of fit was found only after the KS test was conducted at the 99% level. The models in Phases 2, 3 and 5 did not have strong qualitative evidence for the goodness of fit of the model as shown by KS plots in Figure 7. Nonetheless, the transformed time data for each model appeared to be independent as measured by a serial correlation plot.

4.7.2 Interpreting Parameters

From the parameter values presented in Table 6 more evidence is gained for the relative autonomy of the Belfast Brigade from the other PIRA units. In particular, it can be seen that the value of the jump from self-excitations, k_0 , is much higher than that for mutual-excitations, s_0 . Whilst the opposite is true for the decay rate of self-excitations, ω , and those of mutual-excitations, ν . This suggests that the events in the six counties of NI had little impact on IED attacks in Belfast and the impact they did have was short lived. Also it can be seen that the self-excitation part of the model is very similar to that of the Belfast only case in Model 2 suggesting that Model 2 is sufficient for studying the internal dynamics of PIRA attacks in Belfast.

4.8 Model 4

In the final two models examined in this paper the focus will be on the influence of actions by BSFs in NI. These models are only examined in Phase 1 and on spatially aggregated data across NI due to a lack of data concerning BSF attacks. The first of these models is a multivariate Hawkes process considering the influence on PIRA attacks against BSFs based on self-excitations

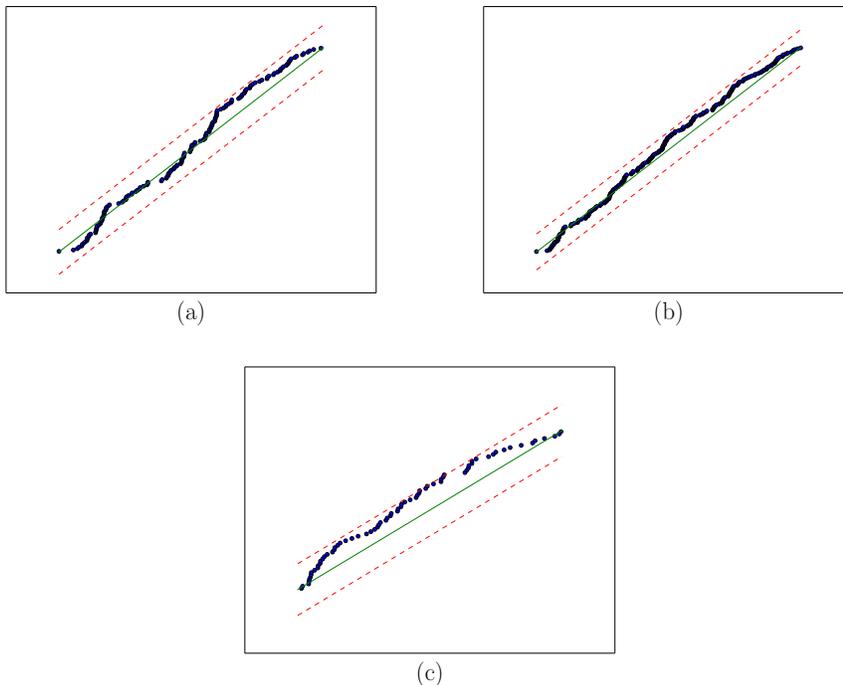


FIGURE 7. Figures (a) - (c) show KS plots for Hawkes process Model 3 in Phases 2, 3 and 5 respectively. Data points falling on the solid goodness of fit line imply a perfect model fit with the dashed lines representing 95% error bounds.

of past IED attacks against BSFs and mutual-excitations of past BSF attacks which killed Catholic civilians. The decision to study the impact of the deaths of Catholic civilians resulted from the fact that the PIRA were heavily rooted in the Catholic community [56]. In our datasets the number of IED attacks targeting BSFs is 144 and the number of Catholic civilian deaths resulting from BSF events is 78.

4.8.1 Goodness of Fit

Results of applying goodness of fit tests to Model 4 are shown in Table 7. Quantitatively Model 4 appears to provide a good fit to the IED data. Qualitatively, however, it should be noted that the data points in the KS plot do not lie on the line of best fit, as can be seen in Figure 8. Nonetheless, the serial correlation plot, although not included here, does suggest independence of the data points.

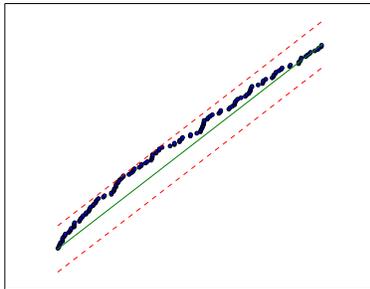


FIGURE 8. KS plot for multidimensional Hawkes process Model 4. Data points falling on the solid goodness of fit line imply a perfect model fit with the dashed lines representing 95% error bounds.

4.8.2 Interpreting Parameters

Interpreting the MLE parameter values for Model 4 it can be seen that BSF attacks, which lead to the death of Catholic civilians, actually caused a backlash in terms of leading to an increase in IED attacks. In particular, the jump parameter, s_0 , is high compared to the other values seen in this paper implying that following an incident involving the death of a Catholic civilian the PIRA were likely to respond with a large increase in IED attacks targeting BSFs. Also the decay rate of this increase in attacks, ν , is small suggesting a lengthy period of increased violence equivalent to an average of 73 days. This prolonged retaliation by PIRA may be seen as an attempt by the organisation to obtain public legitimacy by acting as defenders of the Catholic population. Such results also give support for the sort of findings made by Braithwaite and Johnson [7] where less discriminatory counter-insurgency operations were found to result in an increase in violence.

4.9 Model 5

The final model studied in this paper is similar to Model 4 except now the influence of BSF attacks which killed PIRA members is considered. In our dataset there are 58 recorded incidents of BSF events resulting in PIRA member deaths.

4.9.1 Goodness of Fit

Goodness of fit test results are shown in Table 7. These results show that, based on the AIC, Model 5 performs worse than Model 4 at modelling PIRA attacks targeting BSFs. This result indicates that Catholic civilian deaths were better predictors of a backlash by the PIRA. Moreover, the KS test is only significant for Model 5 when considered at the 99% level. The KS plot for Model 5, shown in Figure 9, suggests that the data points are not falling on the line of best fit so there is not enough qualitative evidence to declare a

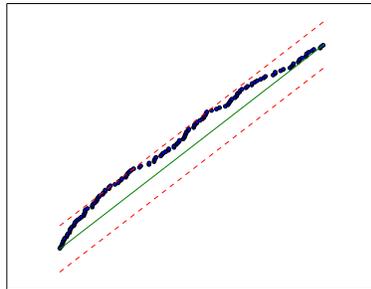


FIGURE 9. KS plot for multidimensional Hawkes process Model 5. Data points falling on the solid goodness of fit line imply a perfect model fit with the dashed lines representing 95% error bounds.

significantly good fit. However, the serial correlation plot did suggest that the data points were independent.

4.9.2 Interpreting Parameters

The parameter values for Model 5 also suggest an increase in violence against BSFs following a BSF event leading to the death of a PIRA member. The increase in the rate of IED attacks, given by s_0 , following such an operation is high compared to other values seen in this paper and the length of time this increase is sustained, given by ν , is quite long at an average of 115 days. These observations may be the result of PIRA trying to revenge the death of its members as well as demonstrating defiance. This is similar to the findings of Lafree, Dugan and Korte [28] which suggested an increase in PIRA attacks following BSF operations which killed PIRA members. When planning CT strategies these observations, and similar ones for Model 4, imply that after a civilian or terrorist death quick interventions are necessary to avoid long time periods of increased violence.

5 Discussion

We began this paper by outlining the importance of mathematical modelling in terrorism and stressing in particular the contributions such modelling could make to improving CT responses. From a firm foundation of criminological theory concerning spatio-temporal patterns of crime and a mathematical model known as a Hawkes process we then proceeded to model IED attacks during the NI conflict.

The first step taken in this paper was to divide the IED dataset we used into five phases corresponding to organisational changes within the PIRA as described by Asal et al. [3]. Although this allowed for a greater depth of analysis it did introduce the issue of edge effects where events in one phase influence those in the next phase. To account for this a novel approach was found using a

moving time frame to incorporate events from the previous phase where an improvement was then judged using the KS test. It was found that in Phases 2 and 3 very few data points were needed to achieve the best fitting model. However in Phases 4 and 5 more significant edge effects were found. For each phase new mathematical boundaries were fixed for the analyses in this paper. Such findings raise interesting questions relating to the timing of tactical and organisational shifts within the PIRA that may have previously been missed.

Having defined the phase boundaries of PIRA we then moved to analyse six models aimed at capturing different temporal patterns of IED usage by the organisation at different geographical scales.

Initially in Model 1 we examined the self-exciting nature of IED attacks across the whole of NI. For this case it was found that the model outperforms a simple Poisson process, defined by Model 0, as measured by the AIC. Moreover, quantitative and qualitative evidence suggested that this model was capturing the temporal patterns of IED attacks. With the model specified we were then able to compare its predictions to historical accounts of the PIRA illustrating how such models could be used in practise to determine how terrorist groups respond to past events.

Being the center of much of the violence seen during “The Troubles” [20] we then decided to refine the geographical scale to focus on Belfast. Here it was found that there was quantitative and qualitative evidence to suggest the model is good at capturing past influences on further IED attacks. In addition, a study of the model parameters revealed a similar pattern to Model 1 although with some difference in Phase 2. But again using historical accounts of the PIRA we were able to account for this difference. This shows that by adjusting the spatial scale it is possible to gain more refined information about a terrorist organisation demonstrating the depth of insights that can be gained from Hawkes process models.

Having studied univariate models the next model considered was a multidimensional model which aimed to examine the interplay between IED attacks in Belfast and those in the six counties of NI. It was found that the additional complexity did not yield significant improvements over the self-exciting model in Belfast only. However, it was found that this could be accounted for by examining the autonomous nature of Active Service Units of the PIRA in the counties of NI and Belfast. Hence this gives an example of how Hawkes processes can also be used to uncover simultaneous influences on different components of terrorist organisations over a range of spatial scales. Quantitative and some qualitative evidence also exists confirming these models are capturing some of the IED dynamics studied.

For the final two models our focus was on determining the effectiveness of multidimensional Hawkes processes for testing CT strategies. The first of these models was used to investigate how BSF attacks which lead to the death of Catholic civilians influenced PIRA attacks. On the other hand the second model examined the influence of BSF attacks leading to PIRA member deaths. The former model proved to be a better fit for the IED data both in terms of a

direct comparison using the AIC and using quantitative and qualitative goodness of fit tests. Although with small datasets for this investigation some care should be taken when drawing conclusions from these models. Despite this the parameter values for both models were consistent with previous research showing the retaliatory nature of terrorist groups. This indicates that important lessons can be learned from Hawkes processes concerning how terrorist groups will respond to different events.

In conclusion, it is hoped that this paper has shown the adaptability of Hawkes process models to study a range of areas within CT. For future research it is planned to extend these models further using an explicit spatial component as opposed to the implicit approach taken here. This should allow further details concerning the patterning of PIRAs attacks to be uncovered and also enable more depth to be gained concerning hotspots of terrorism during the NI conflict. It would also be interesting to take a sociological point of view on the results found here and in particular examine again the changes that occurred in the PIRA to align theory with the mathematical phase boundaries that we have found.

Acknowledgement

This project has been funded by the EPSRC through the UCL Security Science Doctoral Training Centre, reference EP/G037264/1.

References

- [1] Akaike, H. (1974) A new look at the statistical model identification. *IEEE T. Automat. Contr.*, **19**(6), 716–723.
- [2] Andresen, M. A. and Malleson, N. (2011) Testing the Stability of Crime Patterns: Implications for Theory and Policy. *J. Res. Crime Delinq.*, **48**(1), 58–82.
- [3] Asal, V., Gill, P., Rethemeyer, R. K. and Horgan, J. (2013) Killing Range: Explaining Lethality Variance within a Terrorist Organization. *J. Conflict Resolut.* **59**(3), 401–427.
- [4] Berman, M. (1983) Comment on “Likelihood Analysis of Point Processes and Its Applications to Seismological Data”, by Y. Ogata. In: *Bulletin of the International Statistical Institute*, **50** (3), 412–418.
- [5] Bowsher, C. G. (2007) Modelling security market events in continuous time: Intensity based, multivariate point process models. *J. Econometrics*, **141**(2), 876–912.
- [6] Braga, A. A. (2001) The Effects of Hot Spots Policing on Crime. *Ann. Am. Acad. Polit. Ss.*, **578**(1), 104–125.
- [7] Braithwaite, A. and Johnson, S. D. (2012) Space–Time Modeling of Insurgency and Counterinsurgency in Iraq. *J. Quant. Criminol.*, **28**(1), 31–48.
- [8] Brantingham, P. J. and Brantingham, P. L. (1984) *Patterns in Crime*. New York: Macmillan.

- [9] Brémaud, P. and Massoulié, L. (1996) Stability of Nonlinear Hawkes Processes. *Ann. Probab.* **24**(3), 1563–1588.
- [10] Brown, E. N., Barbieri, R., Ventura, V., Kass, R. E. and Frank, L. M. (2002) The Time-Rescaling Theorem and Its Application to Neural Spike Train Data Analysis. *Neural Comput.*, **14**(2), 325–346.
- [11] Burnham, K. P. and Anderson, D. R. (2002) *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach*. 2nd ed. New York: Springer Science & Business Media.
- [12] Campbell, C. and Connolly, I. (2003) A Model for the “War Against Terrorism”? Military Intervention in Northern Ireland and the 1970 Falls Curfew. *J. Law Soc.*, **30**(3), 341–375.
- [13] Chen, Z., Purdon, P., Brown, E. N. and Barbieri, R. (2012) A unified point process probabilistic framework to assess heartbeat dynamics and autonomic cardiovascular control. *Front. Physiol.*, **3**.
- [14] Clauset, A. and Woodard, R. (2013) Estimating the historical and future probabilities of large terrorist events. *J. Appl. Stat.*, **7**(4), 1838–1865.
- [15] Coogan, T. P. (2002) *The IRA*. New York: Palgrave.
- [16] Eck, J. E., Chainey, S., Cameron, J. G., Leitner, M. and Wilson, R. E. (2005) *Mapping Crime: Understanding Hot Spots*. Washington: U.S. Department of Justice.
- [17] Eck, J. E. and Weisburd, D. (1995) *Crime Prevention Series, Volume 4: Crime and Place*. New York: Criminal Justice Press.
- [18] Egesdal, M., Fathauer, C., Louie, K. and Neuman, J. (2010) Statistical and Stochastic Modeling of Gang Rivalries in Los Angeles. *SIURO*, **3**, 72–94.
- [19] English, R. (2004) *Armed Struggle: The History of the IRA*. London: Pan Macmillan.
- [20] Fay, M.-T., Morrissey, M. and Smyth, M. (1999) *Northern Ireland’s Troubles: The Human Costs*. London: Pluto Press.
- [21] Fox, E. W., Short, M. B., Schoenberg, F. P., Coronges, K. D. and Bertozzi, A. L. (2014) Modeling E-mail Networks and Inferring Leadership Using Self-Exciting Point Processes. [online] Available at: <<http://people.math.gatech.edu/~mshort9/papers/IkeNet.pdf>> [Accessed 22 May 2015].
- [22] Gao, P., Guo, D., Liao, K., Webb, J. J. and Cutter, S. L. (2013) Early Detection of Terrorism Outbreaks Using Prospective Space–Time Scan Statistics. *Prof. Geogr.*, **65**(4), 676–691.
- [23] Hawkes, A. G. (1971) Spectra of Some Self-Exciting and Mutually Exciting Point Processes. *Biometrika*, **58**(1), 83–90.
- [24] Hegemann, R. A., Lewis, E. A. and Bertozzi, A. L. (2013) An “Estimate & Score Algorithm” for simultaneous parameter estimation and reconstruction of incomplete data on social networks. *Security Informatics*, **2**(1).
- [25] Horgan, J. and Taylor, M. (1997) The Provisional Irish Republican Army: Command and Functional Structure. *Terror Polit. Violenc.*, **9**(3), 1–32.
- [26] Johnson, N. F., Medina, P., Zhao, G., Messinger, D. S., Horgan, J., Gill, P., Bohorquez, J. C., Mattson, W., Gangi, D., Qi, H., Manrique, P., Velasquez, N., Morgenstern, A., Restrepo, E., Johnson, N., Spagat, M. and Zarama, R.

- (2013) Simple mathematical law benchmarks human confrontations. *Sci. Rep.*, **3**.
- [27] Johnson, S. D. and Bowers, K. J. (2004) The Burglary as Clue to the Future The Beginnings of Prospective Hot-Spotting. *Eur. J. Criminol.*, **1**(2), 237–255.
- [28] Lafree, G., Dugan, L. and Korte, R. (2009) The Impact of British Counterterrorist Strategies on Political Violence in Northern Ireland: Comparing Deterrence and Backlash Models. *Criminology*, **47**(1), 17–45.
- [29] Lewis, E., Mohler, G., Brantingham, P. J. and Bertozzi, A. L. (2012) Self-exciting point process models of civilian deaths in Iraq. *Security Journal*, **25**(3), 244–264.
- [30] Liniger, T. (2009) “Multivariate Hawkes Processes”. PhD thesis. Zurich: Swiss Federal Institute of Technology.
- [31] Lum, C. and Kennedy, L. W. (2012) *Evidence-Based Counterterrorism Policy*. New York: Springer Science + Business Media.
- [32] Massey Jr., F. J. (1951) The Kolmogorov-Smirnov Test for Goodness of Fit. *J. Am. Stat. Assoc.*, **46**(253), 68–78.
- [33] Memon, N., Farley, J. D., Hicks, D. L. and Rosenorn, T. (2009) *Mathematical Methods in Counterterrorism*. New York: Springer Science + Business Media.
- [34] Mohler, G. O., Short, M. B., Brantingham, P. J., Schoenberg, F. P. and Tita, G. E. (2011) Self-Exciting Point Process Modeling of Crime. *J. Am. Stat. Assoc.*, **106**(493), 100–108.
- [35] Mohler, G. (2013) Modeling and estimation of multi-source clustering in crime and security data. *Ann. Appl. Stat.* **7**(3), 1525–1539.
- [36] Moloney, E. (2003) *A Secret History of the IRA*. New York: W. W. Norton.
- [37] Nelder, J. A. and Mead, R. (1965) A Simplex Method for Function Minimization. *Comput. J.*, **7**(4), 308–313.
- [38] Nichols, K. and Schoenberg, F. P. (2014) Assessing the dependency between the magnitudes of earthquakes and the magnitudes of their aftershocks. *Environmetrics*, **25**(3), 143–151.
- [39] O’Brien, B. (1999) *The Long War: The IRA and Sinn Féin*. 2nd ed. New York: Syracuse University Press.
- [40] O’Connor, M. P. and Rumann, C. (2003) Into the Fire: How to Avoid Getting Burned by the Same Mistakes Made Fighting Terrorism in Northern Ireland. *Cardozo Law Review*, **24**(4).
- [41] O’Connor, P. D. T. and Kleyner, A. (2012) *Practical Reliability Engineering*. 5th ed. West Sussex: Wiley.
- [42] Ozaki, T. (1979) Maximum likelihood estimation of Hawkes’ self-exciting point processes. *Ann. I. Stat. Math.*, **31**(1), 145–155.
- [43] Papangelou, F. (1972) Integrability of Expected Increments of Point Processes and a Related Random Change of Scale. *T. Am. Math. Soc.*, **165**, 483–506.
- [44] Pease, K. (1998) *Repeat Victimisation: Taking Stock*. London: Home Office Police Research Group.
- [45] Peng, R. (2003) Multi-dimensional Point Process Models in R. *J. Stat. Softw.*, **8**(16).

- [46] Rasmussen, J. G. (2013) Bayesian Inference for Hawkes Processes. *Methodol. Comput. Appl.*, **15**(3), 623–642.
- [47] Ross, S. (2010) *A First Course in Probability*. 8th ed. London: Pearson.
- [48] Rubin, I. (1972) Regular Point Processes and Their Detection. *IEEE T. Inform. Theory*, **IT-18**(5), 547–557.
- [49] Scipy Optimize Minimize (2014) [online] Available at: <<http://docs.scipy.org/doc/scipy-0.14.0/reference/generated/scipy.optimize.minimize.html#scipy.optimize.minimize>> [Accessed 17 Dec. 2014].
- [50] Short, M. B., Mohler, G. O., Brantingham, P. J. and Tita, G. E. (2014) Gang rivalry dynamics via coupled point process networks. *Discrete Cont. Dyn.-B*, **19**(5), 1459–1477.
- [51] Stevenson, R. and Crossley, N. (2014) Change in Covert Social Movement Networks: The “Inner Circle” of the Provisional Irish Republican Army. *Social Movement Studies*, **13**(1), 70–91.
- [52] Stomakhin, A., Short, M. B. and Bertozzi, A. L. (2011) Reconstruction of missing data in social networks based on temporal patterns of interactions. *Inverse Probl.* **27**(11).
- [53] Sutton, M. (1994) *An Index of Deaths from the Conflict in Ireland 1969-1993*. Belfast: Beyond the Pale.
- [54] Varadhan, S. R. S. (2007) *Stochastic Processes (Courant Lecture Notes)*. Rhode Island: American Mathematical Soc.
- [55] White, G., Porter, M. D. and Mazerolle, L. (2012) Terrorism Risk, Resilience and Volatility: A Comparison of Terrorism Patterns in Three South-east Asian Countries. *J. Quant. Criminol.*, **29**(2), 295–320.
- [56] White, R. W. (1993) *Provisional Irish Republicans: An Oral and Interpretive History*. Westport, CT: Greenwood Press.
- [57] Zar, J. H. (2014) *Biostatistical Analysis*. 5th ed. Essex: Pearson Education Limited.

BIBLIOGRAPHY

- Abell, P. (1971). *Model building in sociology*. London: Weidenfeld and Nicolson.
- Abramowitz, M. and Stegun, I. A. (1964). *Handbook of Mathematical Functions*. New York: Dover.
- ACLEDD (2016). URL: <http://www.acleddata.com/> (visited on 25/10/2016).
- Adamopoulos, L. (1975). “Some Counting and Interval Properties of the Mutually-Exciting Processes”. *Journal of Applied Probability*, **12**(1), pp. 78–86.
- Agbiboa, D. (2014). “Terrorism without Borders: Somalia’s Al-Shabaab and the global jihad network”. *Journal of Terrorism Research*, **5**(1).
- Aggarwal, C. C. and Zhai, C. (2012). *Mining text data*. New York: Springer Science & Business Media.
- Akaike, H. (Dec. 1974). “A new look at the statistical model identification”. *IEEE Transactions on Automatic Control*, **19**(6), pp. 716–723.
- Alpaydin, E. (2014). *Introduction to Machine Learning*. 3rd ed. Massachusetts: The MIT Press.
- Anderson, D. and McKnight, J. (2015a). “Understanding al-Shabaab: clan, Islam and insurgency in Kenya”. *Journal of Eastern African Studies*, **9**(3), pp. 536–557.

-
- Anderson, D. M. and McKnight, J. (1st Jan. 2015b). “Kenya at war: Al-Shabaab and its enemies in Eastern Africa”. *African Affairs*, **114**(454), pp. 1–27.
- Antoniak, C. (1974). “Mixtures of Dirichlet Processes with Applications to Bayesian Nonparametric Problems”. *The Annals of Statistics*, **2**(6), pp. 1152–1174.
- Arney, D. and Arney, K. (2012). “Modeling insurgency, counter-insurgency, and coalition strategies and operations”. *Defense Modeling and Simulation: Applications, Methodology, Technology*, **10**(1), pp. 57–73.
- Artin, E. and Butler, M. (2015). *The Gamma Function*. New York: Dover Publications Inc.
- Asal, V., Gill, P., Rethemeyer, R. and Horgan, J. (2013). “Killing Range Explaining Lethality Variance within a Terrorist Organization”. *Journal of Conflict Resolution*, **59**(3), pp. 401–427.
- Atlas, R. I. (2008). *21st century security and CPTED : designing for critical infrastructure protection and crime prevention*. Boca Raton: CRC Press.
- Atlas, R. (1st Jan. 2003). “Designing Against Terror, Violence and Crime”. ASSE Professional Development Conference and Exposition. American Society of Safety Engineers.
- Baesens, B., Mues, C., Martens, D. and Vanthienen, J. (2009). “50 years of data mining and OR: upcoming trends and challenges”. *Journal of the Operational Research Society*, **60**(1), pp. 16–23.
- Baeza-Yates, R. and Ribeiro-Neto, B. (2011). *Modern Information Retrieval: The Concepts and Technology Behind Search*. 2nd ed. Essex: Pearson.
- Ball, P. (2012). *Why Society is a Complex Matter*. Berlin: Springer-Verlag.

-
- Barabási, A.-L. (12th May 2005). “The origin of bursts and heavy tails in human dynamics”. *Nature*, **435**(7039), pp. 207–211.
- Beavon, D., Brantingham, P. L. and Brantingham, P. J. (1994). “The influence of street networks on the patterning of property offenses.” *Crime Prevention Studies*, **2**, pp. 115–148.
- Bell, E. T. (1934). “Exponential Polynomials”. *Annals of Mathematics*, **35**(2), pp. 258–277.
- Bell, E. T. (1938). “The Iterated Exponential Integers”. *Annals of Mathematics*, **39**(3), pp. 539–557.
- Bell, J. B. (1st Sept. 2000). *The IRA, 1968-2000: An Analysis of a Secret Army*. 1st ed. Oxon: Frank Cass.
- Berg, B. (2004). *Markov chain Monte Carlo simulations and their statistical analysis: with web-based fortran code*. Singapore: World Scientific Publishing.
- Bernasco, W. (2008). “Them again? Same-offender involvement in repeat and near repeat burglaries.” *European Journal of Criminology*, **5**(4), pp. 411–431.
- Bernasco, W., Johnson, S. and Ruiter, S. (2015). “Learning where to offend: Effects of past on future burglary locations”. *Applied Geography*, **60**, pp. 120–129.
- Berndt, D. and Clifford, J. (1994). *Using Dynamic Time Warping to Find Patterns in Time Series*. AAAI-94 Workshop on Knowledge Discovery in Databases. California, pp. 359–370.
- Berrebi, C. and Lakdawalla, D. (1st Apr. 2007). “How Does Terrorism Risk Vary Across Space and Time? An Analysis Based on the Israeli Experience”. *Defence and Peace Economics*, **18**(2), pp. 113–131.

-
- Birk, M. (2016). *Tools for Units of Measurement*. URL: <https://cran.r-project.org/web/packages/measurements/measurements.pdf>.
- Bohorquez, J. C., Gourley, S., Dixon, A. R., Spagat, M. and Johnson, N. F. (17th Dec. 2009). “Common ecology quantifies human insurgency”. *Nature*, **462**(7275), pp. 911–914.
- Bowers, K. (2014). “Risky facilities: Crime radiators or crime absorbers? A comparison of internal and external levels of theft.” *Journal of Quantitative Criminology*, **30**(3), pp. 389–414.
- Bowers, K. and Guerette, R. (2014). “Effectiveness of Situational Crime Prevention”. *Encyclopedia of Criminology and Criminal Justice*. New York: Springer Science+Business Media, pp. 1318–1329.
- Bowsher, C. G. (Dec. 2007). “Modelling security market events in continuous time: Intensity based, multivariate point process models”. *Journal of Econometrics*, **141**(2), pp. 876–912.
- Braithwaite, A. and Johnson, S. D. (1st Mar. 2012). “Space–Time Modeling of Insurgency and Counterinsurgency in Iraq”. *Journal of Quantitative Criminology*, **28**(1), pp. 31–48.
- Brantingham, P. J. and Brantingham, P. L. (1981). *Environmental criminology*. California: Sage.
- Brantingham, P. L. and Brantingham, P. J. (1993a). “Environment, Routine, and Situation: Toward a Pattern Theory of Crime”. *Routine Activity and Rational Choice*. Ed. by Clarke, R. V. and Felson, M. Advances in Criminological Theory Volume 5. New Brunswick: Transaction, pp. 259–294.

-
- Brantingham, P. L. and Brantingham, P. J. (1993b). “Nodes, Paths and Edges: Considerations on the Complexity of Crime and the Physical Environment”. *Journal of Environmental Psychology*, **13**(1), pp. 3–28.
- Brock, D. (2006). *Understanding Moore’s Law: Four Decades of Innovation*. Philadelphia: Chemical Heritage Press.
- Brown, E. N., Barbieri, R., Ventura, V., Kass, R. E. and Frank, L. M. (Feb. 2002). “The Time-Rescaling Theorem and Its Application to Neural Spike Train Data Analysis”. *Neural Computation*, **14**(2), pp. 325–346.
- Burnham, K. P. and Anderson, D. R. (2002). *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach*. 2nd ed. New York: Springer Science & Business Media.
- Bush, R. (1956). *Mathematics for psychologists : examples and problems*. New York: Social Science Research Council.
- Canter, D., Coffey, T., Huntley, M. and Missen, C. (Dec. 2000). “Predicting Serial Killers’ Home Base Using a Decision Support System”. *Journal of Quantitative Criminology*, **16**(4), pp. 457–478.
- Caplan, B. (2006). “Terrorism: The relevance of the rational choice model”. *Public Choice*, **128**(1), pp. 91–107.
- Chandra, R. V. and Varanasi, B. S. (2015). *Python Requests Essentials*. Birmingham: Packt Publishing.
- Chen, H., Reid, E., Sinai, J., Silke, A. and Ganor, B. (2008). *Terrorism Informatics: Knowledge Management and Data Mining for Homeland Security*. New York: Springer Science+Business Media.

-
- Clarke, R. V. (1997). *Situational crime prevention: successful case studies*. 2nd ed. Monsey, NY: Criminal Justice Press.
- Clarke, R. and Eck, J. (2003). *Become a Problem Solving Crime Analyst*. London: Jill Dando Institute of Crime Science.
- Clarke, R. and Newman, G. R. (2006). *Outsmarting the Terrorists*. Connecticut: Praeger Security International.
- Clarke, R. V. and Cornish, D. B. (1st Jan. 1985). “Modeling Offenders’ Decisions: A Framework for Research and Policy”. *Crime and Justice*, **6**, pp. 147–185.
- Clauset, A. and Gleditsch, K. (2012). “The Developmental Dynamics of Terrorist Organizations”. *PLoS One*, **7**(11).
- Clauset, A. and Woodard, R. (2013). “Estimating the historical and future probabilities of large terrorist events”. *Annals of Applied Statistics*, **7**(4), pp. 1838–1865.
- Clauset, A., Young, M. and Gleditsch, K. (2007). “On the Frequency of Severe Terrorist Events”. *Journal of Conflict Resolution*, **51**(1), pp. 58–87.
- Coenen, F. (2011). “Data Mining: Past, Present and Future”. *The Knowledge Engineering Review*, **26**(1), pp. 25–29.
- Coffman, K. and Odlyzko, A. (2002). “Internet Growth: Is There a ”Moore’s Law” for Data Traffic?” *Handbook of Massive Data Sets*. Ed. by Abello, J., Pardalos, P. M. and Resende, M. G. C. Dordrecht: Kluwer Academic, pp. 47–93.
- Cohen, L. and Felson, M. (1979). “Social Change and Crime Rate Trends: A Routine Activity Approach”. *American Sociological Review*, **44**(4), pp. 588–608.

-
- Cohn, E. G. (21st Dec. 1990). "Weather and Crime". *British Journal of Criminology*, **30**(1), pp. 51–64.
- Coogan, T. (2002). *The IRA*. New York: Palgrave.
- Cooley, J. W. and Tukey, J. W. (1965). "An Algorithm for the Machine Calculation of Complex Fourier Series". *Mathematics of Computation*, **19**(90), pp. 297–301.
- Cornish, D. B. and Clarke, R. (1986). *The Reasoning Criminal*. New York: Springer-Verlag.
- Cothren, J., Smith, B., Roberts, P. and Damphousse, K. (2008). "Geospatial and Temporal Patterns of Preparatory Conduct among American Terrorists". *International Journal of Comparative and Applied Criminal Justice*, **32**(1), pp. 23–41.
- Craven, M., DiPasquo, D., Freitag, D., McCallum, A., Mitchell, T., Nigam, K. and Slattery, S. (2000). "Learning to construct knowledge bases from the World Wide Web". *Artificial intelligence*, **118**(1-2), pp. 69–113.
- Crowe, T. D. (2000). *Crime prevention through environmental design : applications of architectural design and space management concepts*. In collab. with National Crime Prevention Institute. 2nd ed. Boston: Butterworth-Heinemann. xii+333.
- Daley, D. and Vere-Jones, D. (2003). *An Introduction to the Theory of Point Processes: Volume I: Elementary Theory and Methods*. 2nd ed. New York: Springer.
- Davies, T. and Johnson, S. (2015). "Examining the relationship between road structure and burglary risk via quantitative network analysis." *Journal of Quantitative Criminology*, **31**(3), pp. 481–507.

-
- Davies, T. and Bishop, S. (2013). “Modelling patterns of burglary on street networks”. *Crime Science*, **2**(10).
- Davis, H. F. (1989). *Fourier Series and Orthogonal Functions*. New York: Dover Books.
- Davis, P. J. and Rabinowitz, P. (30th Nov. 2007). *Methods of Numerical Integration*. 2nd edition. Mineola, N.Y: Dover Publications Inc.
- Demchenko, Y., Grosso, P., Laat, C. de and Membrey, P. (2013). “Addressing big data issues in Scientific Data Infrastructure”. *2013 International Conference on Collaboration Technologies and Systems (CTS)*, pp. 48–55.
- Ding, H., Trajcevski, G., Scheuermann, P., Wang, X. and Keogh, E. (2008). “Querying and Mining of Time Series Data: Experimental Comparison of Representations and Distance Measures”. *Proceedings of the VLDB Endowment*, **1**(2), pp. 1542–1552.
- Draper, R. and Cadzow, E. (1st Oct. 2004). *Crime Prevention Through Environmental Design*. PEB Exchange, Programme on Educational Building. Paris: Organisation for Economic Co-operation and Development.
- Dyke, P. (2014). *An Introduction to Laplace Transforms and Fourier Series*. 2nd ed. London: Springer-Verlag.
- Eck, J. (1994). “Drug markets and drug places: A case-control study of the spatial structure of illicit drug dealing”. PhD thesis. Maryland: University of Maryland.
- Eck, J. (2010). “Places and the Crime Triangle”. *Encyclopedia of Criminological Theory*. Thousand Oaks: SAGE Publications, Inc.

-
- Egesdal, M., Fathauer, C., Louie, K. and Neuman, J. (2010). *Statistical and Stochastic Modeling of Gang Rivalries in Los Angeles*. Vol. 3.
- Embrechts, P., Liniger, T. and Lin, L. (2011). “Multivariate Hawkes Processes: An Application to Financial Data”. *Journal of Applied Probability*, **48A**, pp. 367–378.
- Enders, W., Parise, G. F. and Sandler, T. (1st Nov. 1992). “A time-series analysis of transnational terrorism: Trends and cycles”. *Defence Economics*, **3**(4), pp. 305–320.
- English, R. (2004). *Armed Struggle: The History of the IRA*. London: Pan Macmillan.
- Epstein, J. (2008). “Why model?” *Journal of Artificial Societies and Social Simulation*, **11**(4), p. 12.
- Escobar, M. and West, M. (1995). “Bayesian Density Estimation and Inference Using Mixtures”. *Journal of the American Statistical Association*, **90**(430), pp. 577–588.
- FAO SWALIM (2016). URL: <http://www.faoswalim.org/> (visited on 02/11/2016).
- Farley, J. (2007). “Evolutionary Dynamics of the Insurgency in Iraq: A Mathematical Model of the Battle for Hearts and Minds”. *Studies in Conflict & Terrorism*, **30**(11), pp. 947–962.
- Faulkner, S. C., Stevenson, M. D., Verity, R., Mustari, A. H., Semple, S., Tosh, D. G. and Le Comber, S. C. (1st Apr. 2015). “Using geographic profiling to locate elusive nocturnal animals: a case study with spectral tarsiers”. *Journal of Zoology*, **295**(4), pp. 261–268.
- Faulkner, S. C., Verity, R., Roberts, D., Roy, S. S., Robertson, P. A., Stevenson, M. D. and Le Comber, S. C. (1st Jan. 2017). “Using geographic profiling to

-
- compare the value of sightings vs trap data in a biological invasion”. *Diversity and Distributions*, **23**(1), pp. 104–112.
- Fay, M.-T., Morrissey, M. and Smyth, M. (1999). *Northern Ireland’s Troubles: The Human Costs*. London: Pluto Press.
- Fayyad, U., Piatetsky-Shapiro, P. and Smyth, P. (1996). “From Data Mining to Knowledge Discovery in Databases”. *American Association for Artificial Intelligence*, **17**(3), pp. 37–54.
- Ferguson, T. (1983). “Bayesian density estimation by mixtures of normal distributions”. *Recent Advances in Statistics*, pp. 287–303.
- Frigo, M. and Johnson, S. (2005). “The Design and Implementation of FFTW3”. *Proceedings of the IEEE*, **93**(2), pp. 216–231.
- Gamerman, D. (1997). *Markov Chain Monte Carlo: Stochastic Simulation for Bayesian Inference*. Boca Raton: Chapman & Hall.
- Gelfand, A., Hills, S., Racine-Poon, A. and Smith, A. (1990). “Illustration of Bayesian Inference in Normal Data Models Using Gibbs Sampling”. *Journal of the American Statistical Association*, **85**(412), pp. 972–985.
- Gelman, A., Carlin, J., Stern, H., Dunson, D., Vehtari, A. and Rubin, D. (2014). *Bayesian Data Analysis*. 3rd ed. Florida: CRC Press.
- Getoor, L. and Machanavajjhala, A. (2012). “Entity Resolution: Theory, Practice & Open Challenges”. *Proceedings of the VLDB Endowment*, **5**(12), pp. 2018–2019.
- Gilks, W., Richardson, S. and Spiegelhalter, D. (1996). *Markov Chain Monte Carlo In Practice*. Boca Raton: Chapman & Hall.

-
- Glen, A. G. (2011). “On the Inverse Gamma as a Survival Distribution”. *Journal of Quality Technology*, **43**(2), pp. 158–166.
- Goldstein, H. (1990). *Problem-oriented policing*. New York: McGraw Hill.
- Green, P. and Richardson, S. (2001). “Modelling Heterogeneity with and without the Dirichlet Process”. *Scandinavian Journal of Statistics*, **28**(2), pp. 355–375.
- Grenander, U. (1959). *Probability and statistics; the Harald Cramér volume*. New York: Wiley.
- Groff, E. and Lockwood, B. (2014). “Criminogenic facilities and crime across street segments in Philadelphia uncovering evidence about the spatial extent of facility influence.” *Journal of Research in Crime and Delinquency*, **51**(3), pp. 277–314.
- Han, J., Kamber, M. and Pei, J. (2012). *Data Mining: Concepts and Techniques*. 3rd ed. Massachusetts: Morgan Kaufmann.
- Hanhimäki, J. and Blumenau, B. (2013). *An International History of Terrorism: Western and Non-Western Experience*. Oxon: Routledge.
- Hansen, P. C., Pereyra, V. and Scherer, G. (5th Dec. 2012). *Least Squares Data Fitting with Applications*. Baltimore, Md: Johns Hopkins University Press.
- Hansen, S. (2013). *Al-Shabaab in Somalia: The History and Ideology of a Militant Islamist Group, 2005-2012*. New York: Oxford University Press.
- Hardiman, S. J., Bercot, N. and Bouchaud, J.-P. (28th Oct. 2013). “Critical reflexivity in financial markets: a Hawkes process analysis”. *The European Physical Journal B*, **86**(10), pp. 1–9.
- Harris, F. (1978). “On the use of windows for harmonic analysis with the discrete Fourier transform”. *Proceedings of the IEEE*, **66**(1), pp. 51–83.

-
- Hawkes, A. G. (1st Apr. 1971). "Spectra of Some Self-Exciting and Mutually Exciting Point Processes". *Biometrika*, **58**(1), pp. 83–90.
- Hawkes, A. G. and Oakes, D. (1st Sept. 1974). "A Cluster Process Representation of a Self-Exciting Process". *Journal of Applied Probability*, **11**(3), pp. 493–503.
- Heal, K., Laycock, G. K. and Great Britain. Home Office. Research and Planning Unit (1986). *Situational crime prevention : from theory into practice*. London: HMSO.
- Hegemann, R., Lewis, E. and Bertozzi, A. (2013). "An "Estimate & Score Algorithm" for simultaneous parameter estimation and reconstruction of incomplete data on social networks". *Security Informatics*, **2**(1).
- Hey, T., Tansley, S. and Tolle, K. (2009). *The Fourth Paradigm*. Washington: Microsoft Research.
- Hijmans, R., Williams, E. and Vennes, C. (2016). *Spherical Trigonometry*. URL: <https://cran.r-project.org/web/packages/geosphere/geosphere.pdf>.
- Horgan, J. and Braddock, K. (2012). *Terrorism Studies: A Reader*. Oxon: Routledge.
- Horgan, J., Gill, P., Hunter, S., Piazza, J., Cushenbery, L., Lovelace, J., Foo, S. C., Asal, V., Rethemeyer, K., Cousins, K., Anderson, I., Young, J. and Johnson, N. (2013). *From Bomb to Bomb Maker: A Social Network Analysis Model of the Socio-Psychological and Cultural Dynamics of the IED Process*. Virginia: International Center for the Study of Terrorism.
- Horgan, J. and Taylor, M. (1997). "The provisional Irish republican army: Command and functional structure". *Terrorism and Political Violence*, **9**(3), pp. 1–32.

-
- IBM (2015). *The Four V's of Big Data*, IBM. URL: http://www.ibmbigdatahub.com/sites/default/files/infographic_file/4-Vs-of-big-data.jpg.
- Johnson, N., Medina, P., Zhao, G., Messinger, D., Horgan, J., Gill, P., Bohorquez, J., Mattson, W., Gangi, D., Qi, H., Manrique, P., Velasquez, N., Morgenstern, A., Restrepo, E., Johnson, N., Spagat, M. and Zarama, R. (2013). "Simple mathematical law benchmarks human confrontations". *Scientific Reports*, **3**.
- Johnson, N. F., Restrepo, E. M. and Johnson, D. E. (2015). "Modeling Human Conflict and Terrorism Across Geographic Scales". *Social Phenomena*. Ed. by Gonçalves, B. and Perra, N. Computational Social Sciences. Springer International Publishing, pp. 209–233.
- Johnson, N., Carran, S., Botner, J., Fontaine, K., Laxague, N., Nuetzel, P., Turnley, J. and Tivnan, B. (1st July 2011). "Pattern in Escalations in Insurgent and Terrorist Activity". *Science*, **333**(6038), pp. 81–84.
- Johnson, S. (2008). "Repeat burglary victimisation: a tale of two theories". *Experimental Criminology*, **4**(3), pp. 215–240.
- Johnson, S. and Bowers, K. (2010). "Permeability and burglary risk: Are cul-de-sacs safer?" *Journal of Quantitative Criminology*, **26**(1), pp. 89–111.
- Johnson, S. and Braithwaite, A. (2009). "Spatio-temporal modeling of insurgency in Iraq." *Crime Prevention Studies*, **25**, pp. 9–32.
- Johnson, S., Summers, L. and Pease, K. (2009). "Offender as forager? A direct test of the boost account of victimization." *Journal of Quantitative Criminology*, **25**(2), pp. 181–200.

-
- Johnson, S. (2017). *The NLOpt nonlinear-optimization package*. URL: <http://ab-initio.mit.edu/nlopt> (visited on 14/02/2017).
- Johnson, S. D. and Bowers, K. J. (1st Apr. 2004). “The Burglary as Clue to the Future: The Beginnings of Prospective Hot-Spotting”. *European Journal of Criminology*, **1**(2), pp. 237–255.
- Kempf, W. and Repp, B. (1977). *Mathematical Models for Social Psychology*. London: Wiley.
- Kendall, M. (1968). *Mathematical model building in economics and industry*. London: Griffin.
- Keogh, E. and Ratanamahatana, C. A. (2005). “Exact indexing of dynamic time warping”. *Knowledge and Information Systems*, **7**(3), pp. 358–386.
- Kliot, N. and Charney, I. (9th Nov. 2006). “The geography of suicide terrorism in Israel”. *GeoJournal*, **66**(4), pp. 353–373.
- Lafree, G. and Dugan, L. (2007). “Introducing the Global Terrorism Database”. *Terrorism and Political Violence*, **19**(2), pp. 181–204.
- Lafree, G., Dugan, L., Xie, M. and Singh, P. (2012). “Spatial and Temporal Patterns of Terrorist Attacks by ETA 1970 to 2007”. *Journal of Quantitative Criminology*, **28**(1), pp. 7–29.
- Lafree, G., Morris, N. A. and Dugan, L. (1st July 2010). “Cross-National Patterns of Terrorism Comparing Trajectories for Total, Attributed and Fatal Attacks, 1970–2006”. *British Journal of Criminology*, **50**(4), pp. 622–649.
- Larose, D. and Larose, C. (2014). *Discovering Knowledge in Data: An Introduction to Data Mining*. 2nd ed. New Jersey: John Wiley & Sons, Inc.

-
- Lenhard, J., Küppers, G. and Shinn, T. (2006). *Simulation : pragmatic construction of reality*. Dordrecht: Springer.
- Leslie, B. (2016). “Automating social network analysis: A power tool for counter-terrorism”. *Security Journal*, **29**(2), pp. 147–168.
- Levine, N. (2006). “Crime Mapping and the Crimestat Program”. *Geographical Analysis*, **38**(1), pp. 41–56.
- Lewis, E., Mohler, G., Brantingham, P. J. and Bertozzi, A. L. (July 2012). “Self-exciting point process models of civilian deaths in Iraq”. *Security Journal*, **25**(3), pp. 244–264.
- Lewis, P. and Shedler, G. (1979). “Simulation of nonhomogeneous Poisson processes by thinning”. *Naval Research Logistics Quarterly*, **26**(3), pp. 403–413.
- Liniger, T. (2009). “Multivariate Hawkes Processes”. PhD thesis. Zurich: Swiss Federal Institute of Technology.
- Liu, H. and Motoda, H. (2001). *Feature extraction, construction and selection: A data mining perspective*. Massachusetts: Kluwer Academic Publishers.
- Liu, L. and Eck, J. (2008). *Artificial Crime Analysis Systems: Using Computer Simulations and Geographic Information Systems*. New York: Information Science Reference.
- López, M. (2005). “The Spatial Behavior of Residential Burglars”. *Proceedings of the 5th International Space Syntax Symposium TU Delft*. Delft: Techne Press.
- Luhn, H. (1957). “A Statistical Approach to Mechanized Encoding and Searching of Literary Information”. *IBM Journal of Research and Development*, **1**(4), pp. 309–317.

-
- Lum, C., Kennedy, L. W. and Sherley, A. (1st Nov. 2006). “Are counter-terrorism strategies effective? The results of the Campbell systematic review on counter-terrorism evaluation research”. *Journal of Experimental Criminology*, **2**(4), pp. 489–516.
- Lyons, T., Marlowe, J., Mickler, D. and Thornton, A. (June 2015). “Denizens, ‘artists’, and terrorists - disarmament, development and diagnosis: Understanding Africa in Australasia and the Pacific”. *Australasian Review of African Studies*, **36**(1), p. 3.
- Manning, C., Raghavan, P. and Schütze, H. (2008). *Introduction to Information Retrieval*. New York: Cambridge University Press.
- Marchal, R. (1st Nov. 2009). “A tentative assessment of the Somali Harakat Al-Shabaab”. *Journal of Eastern African Studies*, **3**(3), pp. 381–404.
- Masashi, S. (2016). *Introduction to Statistical Machine Learning*. Massachusetts: Morgan Kaufmann.
- Massey Jr., F. J. (1st Mar. 1951). “The Kolmogorov-Smirnov Test for Goodness of Fit”. *Journal of the American Statistical Association*, **46**(253), pp. 68–78.
- Maszka, J. (2017). “A Strategic Analysis of Al Shabaab”. PhD thesis. Bournemouth University.
- Matusitz, J. (2013). *Terrorism and Communication: A Critical Introduction*. California: SAGE Publications, Inc.
- Mayer-Schönberger, V. and Cukier, K. (2013). *Big Data: A Revolution that Will Transform how We Live, Work, and Think*. New York: Houghton Mifflin Harcourt.

-
- McGarrell, E. F., Freilich, J. D. and Chermak, S. (1st May 2007). “Intelligence-Led Policing as a Framework for Responding to Terrorism”. *Journal of Contemporary Criminal Justice*, **23**(2), pp. 142–158.
- McKittrick, D. (2007). *Lost Lives: The stories of the men, women and children who died as a result of the Northern Ireland troubles*. Edinburgh: Mainstream Publishing.
- McPherson, S. (2010). *Tim Berners-Lee: Inventor of the World Wide Web*. Minneapolis: Twenty-First Century Books.
- Memon, N., Farley, J. D., Hicks, D. L. and Rosenorn, T. (Aug. 2009). *Mathematical Methods in Counterterrorism*. New York: Springer Science + Business Media.
- Miller, C. (2003). “Geographic Profiling Serial Offenses with ECRI’s Rigel”. *Law Enforcement Technology*, **30**(7), pp. 130–135.
- Mitchell, L. and Cates, M. E. (23rd Dec. 2009). “Hawkes process as a model of social interactions: a view on video dynamics”. *Journal of Physics A: Mathematical and Theoretical*, **43**(4).
- Mohler, G. (2013). “Modeling and Estimation of Multi-Source Clustering in Crime and Security Data”. *The Annals of Applied Statistics*, **7**(3), pp. 1525–1539.
- Moloney, E. (2003). *A Secret History of the IRA*. New York: W. W. Norton.
- Mooers, C. (1950). “The theory of digital handling of non-numerical information and its implications to machine economics”. *Proceedings of the meeting of the Association for Computing Machinery at Rutgers University*.
- Moore, G. (1965). “Cramming more components onto integrated circuits”. *Electronics*, **38**(8), pp. 114–117.

-
- Muchiri, P. (2007). *Climate of Somalia*. W-01. Nairobi: FAO.
- Müller, A. and Guido, S. (2017). *Introduction to Machine Learning with Python: A Guide for Data Scientists*. California: O'Reilly Media.
- Neal, R. (2000). "Markov Chain Sampling Methods for Dirichlet Process Mixture Models". *Journal of Computational and Graphical Statistics*, **9**(2), pp. 249–265.
- Nelder, J. A. and Mead, R. (1st Jan. 1965). "A Simplex Method for Function Minimization". *The Computer Journal*, **7**(4), pp. 308–313.
- O'Brien, B. (1999). *The Long War: The IRA and Sinn Féin*. 2nd ed. New York: Syracuse University Press.
- O'Connor, P. and Kleyner, A. (2012). *Practical Reliability Engineering*. 5th ed. West Sussex: Wiley.
- O'Leary, M. (2009). "The mathematics of geographic profiling". *Journal of Investigative Psychology and Offender Profiling*, **6**(3), pp. 253–265.
- O'Leary, M. (2010). "Implementing a Bayesian approach to criminal geographic profiling". First International Conference on Computing for Geospatial Research and Application, June 21-23. Washington.
- O'Leary, M. (2012). *New Mathematical Approach to Geographic Profiling*. Washington: National Institute of Justice.
- O'Loughlin, J., Witmer, F. and Linke, A. (2010). "The Afghanistan-Pakistan Wars, 2008-2009: Micro-geographies, Conflict Diffusion, and Clusters of Violence". *Eurasian Geography and Economics*, **51**(4), pp. 437–471.
- O'Loughlin, J., Holland, E. C. and Witmer, F. D. W. (1st Sept. 2011). "The Changing Geography of Violence in Russia's North Caucasus, 1999-2011: Regional Trends

-
- and Local Dynamics in Dagestan, Ingushetia, and Kabardino-Balkaria”. *Eurasian Geography and Economics*, **52**(5), pp. 596–630.
- O’Loughlin, J. and Witmer, F. D. W. (1st Oct. 2012). “The Diffusion of Violence in the North Caucasus of Russia, 1999–2010”. *Environment and Planning A*, **44**(10), pp. 2379–2396.
- Ogata, Y. (1981). “On Lewis’ simulation method for point processes”. *IEEE Transactions on Information Theory*, **IT-27**(1), pp. 23–31.
- Ogata, Y. (1983). “Likelihood analysis of point processes and its applications to seismological data”. *Bulletin of the International Statistical Institute*, **50**(2), pp. 943–961.
- Ogata, Y. (1999). “Seismicity Analysis through Point-process Modeling: A Review”. *Pure and Applied Geophysics*, **155**(2), pp. 471–507.
- Ogata, Y., Akaike, H. and Katsura, K. (1982). “The application of linear intensity models to the investigation of causal relations between a point process and another stochastic process”. *Annals of the Institute of Statistical Mathematics*, **24**(1), pp. 373–387.
- Ogata, Y. (1st Mar. 1988). “Statistical Models for Earthquake Occurrences and Residual Analysis for Point Processes”. *Journal of the American Statistical Association*, **83**(401), pp. 9–27.
- Ogata, Y. and Akaike, H. (1982). “On Linear Intensity Models for Mixed Doubly Stochastic Poisson and Self- Exciting Point Processes”. *Journal of the Royal Statistical Society. Series B (Methodological)*, **44**(1), pp. 102–107.

-
- Oppenheimer, A. R. (2009). *IRA: The Bombs and the Bullets: A History of Deadly Ingenuity*. Newbridge: Irish Academic Press Ltd.
- Ostaszewski, A. (1993). *Mathematics in Economics: Models and Methods*. Oxford: Blackwell.
- Ozaki, T. (1st Dec. 1979). “Maximum likelihood estimation of Hawkes’ self-exciting point processes”. *Annals of the Institute of Statistical Mathematics*, **31**(1), pp. 145–155.
- Pakiam, J. E. and Lim, J. (1st Jan. 1984). “Crime and Weather in Singapore”. *International Journal of Comparative and Applied Criminal Justice*, **8**(1), pp. 209–220.
- Papangelou, F. (1st Mar. 1972). “Integrability of Expected Increments of Point Processes and a Related Random Change of Scale”. *Transactions of the American Mathematical Society*, **165**, pp. 483–506.
- Park, A., Tsang, H., Sun, M. and Glässer, U. (2012). “An agent-based model and computational framework for counter-terrorism and public safety based on swarm intelligence”. *Security Informatics*, **1**(23), p. 9.
- Patterson, H. (1997). *The Politics of Illusion: Political History of the I.R.A.* 2nd revised ed. London: Serif.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M. and Duchesnay, É. (2011). “Scikit-learn: Machine Learning in Python”. *Journal of Machine Learning Research*, **12**, pp. 2825–2830.

-
- Peters, C., Braschler, M. and Clough, P. (2012). *Multilingual Information Retrieval: From Research To Practice*. Berlin: Springer-Verlag.
- Picoli, S., Castillo-Mussot, M. del, Ribeiro, H. V., Lenzi, E. K. and Mendes, R. S. (2014). “Universal bursty behaviour in human violent conflicts”. *Scientific Reports*, **4**(4773), pp. 1–3.
- Pinkus, A. and Zafrany, S. (1997). *Fourier Series and Integral Transforms*. Cambridge: Cambridge University Press.
- Pitcher, A. and Johnson, S. (2011). “Exploring theories of victimization using a mathematical model of burglary.” *Journal of Research in Crime and Delinquency*, **48**(1), pp. 83–109.
- Porter, M. D. and White, G. (Mar. 2012). “Self-exciting hurdle models for terrorist activity”. *The Annals of Applied Statistics*, **6**(1), pp. 106–124.
- Problem-oriented policing* (2003). *Problem-oriented policing : from innovation to mainstream*. In collab. with Knutsson, J. Crime prevention studies ; v. 15. Monsey, NY: Criminal Justice Press.
- Raiffa, H. and Schlaifer, R. (2000). *Applied Statistical Decision Theory*. Canda: Wiley Classics Library.
- Raleigh, C., Linke, A., Hegre, H. and Karlsen, J. (2010). “Introducing ACLED - Armed Conflict Location and Event Data”. *Journal of Peace Research*, **47**(5), pp. 651–660.
- Rambaldi, M., Pennesi, P. and Lillo, F. (2015). “Modeling foreign exchange market activity around macroeconomic news: Hawkes-process approach”. *Physical Review E*, **91**(1).

-
- Ranson, M. (May 2014). “Crime, weather, and climate change”. *Journal of Environmental Economics and Management*, **67**(3), pp. 274–302.
- Rasmussen, J. G. (1st Sept. 2013). “Bayesian Inference for Hawkes Processes”. *Methodology and Computing in Applied Probability*, **15**(3), pp. 623–642.
- Restle, F. (1971). *Mathematical models in psychology : an introduction*. Harmondsworth: Penguin.
- Reynaud-Bouret, P. and Schbath, S. (2010). “Adaptive Estimation for Hawkes Processes: Application to Genome Analysis”. *The Annals of Statistics*, **38**(5), pp. 2781–2822.
- Rhodes, W. and Conly, C. (1981). “Crime and mobility: An empirical study.” *Environmental Criminology*. Ed. by Brantingham, P. J. and Brantingham, P. L. Illinois: Waveland Press, pp. 167–188.
- Robertson, S. (2004). “Understanding inverse document frequency: on theoretical arguments for IDF”. *Journal of Documentation*, **60**(5), pp. 503–520.
- Roelleke, T. (2013). *Information Retrieval Models: Foundations and Relationships*. San Rafael, Calif.: Morgan & Claypool Publishers.
- Ross, G. (2015). “Parametric and Non-parametric Sequential Change Detection in R: The cpm Package”. *Journal of Statistical Software*, **66**(3), pp. 1–20.
- Ross, S. (2010). *A First Course in Probability*. 8th ed. London: Pearson.
- Rossmo, D. (2000). *Geographic Profiling*. Florida: CRC Press.
- Rouanet, P. (2018). *DTW (Dynamic Time Warping) python module*. URL: <https://github.com/pierre-rouanet/dtw>.

-
- Rubin, I. (1972). “Regular Point Processes and Their Detection”. *IEEE Transactions on Information Theory*, **IT-18**(5), pp. 547–557.
- Ryder, C. (2005). *A Special Kind of Courage: Bomb Disposal and the Inside Story of 321 EOD Squadron*. London: Methuen Publishing Ltd.
- Sagovsky, A. and Johnson, S. D. (1st Apr. 2007). “When Does Repeat Burglary Victimization Occur?” *Australian & New Zealand Journal of Criminology*, **40**(1), pp. 1–26.
- Sakoe, H. and Chiba, S. (1978). “Dynamic programming algorithm optimization for spoken word recognition”. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, **26**(1), pp. 43–49.
- SciPy Optimize* (2015). URL: <http://docs.scipy.org/doc/scipy-0.15.1/reference/optimize.html> (visited on 16/11/2015).
- Scipy.Stats* (2016). URL: http://docs.scipy.org/doc/scipy/reference/generated/scipy.stats.gaussian_kde.html (visited on 24/02/2016).
- Scott, M. S. (2000). *Problem-oriented policing : reflections on the first 20 years*. In collab. with United States. Dept. of Justice. Office of Community Oriented Policing Services. COPS publication. Washington, DC: US Department of Justice, Office of Community Oriented Policing Services.
- Shinn, D. (2011). “Al Shabaab’s Foreign Threat to Somalia”. *Orbis*, **55**(2), pp. 203–215.
- Short, M. B., Mohler, G. O., Brantingham, P. J. and Tita, G. E. (Apr. 2014). “Gang rivalry dynamics via coupled point process networks”. *Discrete and Continuous Dynamical Systems - Series B*, **19**(5), pp. 1459–1477.

-
- Short, M., D'Orsogna, M., Brantingham, P. and Tita, G. (2009). "Measuring and Modeling Repeat and Near-Repeat Burglary Effects". *Quantitative Criminology*, **25**(3), pp. 325–339.
- Shubik, M. (1967). *Essays in mathematical economics in honor of Oskar Morgenstern*. New Jersey: Princeton University Press.
- Silke, A. (2001). "The Devil You Know: Continuing Problems with Research on Terrorism". *Terrorism and Political Violence*, **13**(4), pp. 1–14.
- Simon, H. (1990). "Invariants of Human Behavior". *Annual Review of Psychology*, **41**, pp. 1–19.
- Smith, B. (2008). "A look at terrorist behavior: how they prepare, where they strike". *National Institute of Justice Journal*.
- Smith, M. J. and Cornish, D. B. (2003). *Theory for practice in situational crime prevention*. Crime prevention studies v. 16. Monsey, NY: Criminal Justice Press.
- Smith, M. L. R. (1997). *Fighting for Ireland?: Military Strategy of the Irish Republican Movement*. London: Routledge.
- Smith, M. J. and Tilley, N., eds. (2005). *Crime Science: New Approaches to Preventing and Detecting Crime*. Cullompton: Willan.
- Sorensen, A. (1978). "Mathematical Models in Sociology". *Annual Review of Sociology*, **41**(1), pp. 345–371.
- Spärck Jones, K. (1972). "A statistical interpretation of term specificity and its application in retrieval". *Journal of Documentation*, **28**(1), pp. 11–21.
- Stade, E. (2005). *Fourier Analysis*. New Jersey.: John Wiley & Sons, Inc.

-
- Staudt, V. (1998). “Effects of window functions explained by signals typical to power electronics”. *8th International Conference on Harmonics and Quality of Power. Proceedings*, pp. 952–957.
- Stein, E. M. and Shakarchi, R. (2003). *Fourier Analysis: An Introduction*. New Jersey: Princeton University Press.
- Stevenson, M. (2013). “Geographic profiling in biology.” PhD thesis. Queen Mary University of London.
- Stomakhin, A., Short, M. and Bertozzi, A. (2011). “Reconstruction of missing data in social networks based on temporal patterns of interactions”. *Inverse Problems*, **27**(11).
- Storn, R. and Price, K. (1997). “Differential Evolution – A Simple and Efficient Heuristic for Global Optimization over Continuous Spaces”. *Journal of Global Optimization*, **11**(4), pp. 341–359.
- Strang, G. (1st Jan. 1986). *Introduction to Applied Mathematics*. Wellesley, Mass: Wellesley-Cambridge Press.
- Sutton, M. (1994). *An Index of Deaths from the Conflict in Ireland 1969-1993*. Belfast: Beyond the Pale.
- Tanizaki, H. (2004). *Computational Methods in Statistics and Econometrics*. New York: Marcel Dekker.
- Taylor, P. (24th Sept. 1998). *The Provos: The IRA and Sinn Fein*. New edition edition. London: Bloomsbury Publishing PLC. 416 pp.
- Team, R. C. (2017). *R: A Language and Environment for Statistical Computing*. URL: <https://www.R-project.org>.

-
- Tench, S. (2014). “Temporal Patterns of IED usage by the Provisional Irish Republican Army”. Master’s Thesis. University College London (UCL).
- Tench, S., Fry, H. and Gill, P. (Jan. 2016). “Spatio-temporal patterns of IED usage by the Provisional Irish Republican Army”. *European Journal of Applied Mathematics*, **27**(3), pp. 377–402.
- Thuraisingham, B. (2003). *Web Data Mining and Applications in Business Intelligence and Counter-Terrorism*. Florida: CRC Press.
- Tilley, N. (2012). *Crime Prevention*. Oxon: Routledge.
- Tolstov, G. P. (1977). *Fourier Series*. Trans. by Silverman, R. A. Mineola: Dover Publications Inc.
- Toure, I. and Gangopadhyay, A. (2016). “Real time big data analytics for predicting terrorist incidents”. 2016 IEEE Symposium on Technologies for Homeland Security (HST). Waltham, MA, pp. 1–6.
- Townsley, M., Homel, R. and Chaseling, J. (1st Apr. 2000). “Repeat Burglary Victimization: Spatial and Temporal Patterns”. *Australian & New Zealand Journal of Criminology*, **33**(1), pp. 37–63.
- Townsley, M., Johnson, S. D. and Ratcliffe, J. H. (2008). “Space Time Dynamics of Insurgent Activity in Iraq”. *Security Journal*, **21**(3), pp. 139–146.
- Verity, R., Stevenson, M., Rossmo, D., Nichols, R. and Le Comber, S. (2014). “Spatial targeting of infectious disease control: Identifying multiple, unknown sources”. *Methods in Ecology and Evolution*, **5**(7), pp. 647–655.

-
- Vincenty, T. (1st Apr. 1975). “Direct and Inverse Solutions of Geodesics on the Ellipsoid with Application of Nested Equations”. *Survey Review*, **23**(176), pp. 88–93.
- Wang, T., Bebbington, M. and Harte, D. (30th Dec. 2010). “Markov-modulated Hawkes process with stepwise decay”. *Annals of the Institute of Statistical Mathematics*, **64**(3), pp. 521–544.
- Ward, J. S. and Barker, A. (2013). “Undefined By Data: A Survey of Big Data Definitions”. *CoRR*, **abs/1309.5821**. arXiv: 1309.5821.
- Weimann, G. and Brosius, H.-B. (1st Jan. 1988). “The predictability of international terrorism: A time-series analysis”. *Terrorism*, **11**(6), pp. 491–502.
- White, G., Porter, M. and Mazerolle, L. (2012). “Terrorism Risk, Resilience and Volatility: A Comparison of Terrorism Patterns in Three Southeast Asian Countries”. *Journal of Quantitative Criminology*, **29**(2), pp. 295–320.
- Williams, P. D. (1st July 2014). “After Westgate: opportunities and challenges in the war against Al-Shabaab”. *International Affairs*, **90**(4), pp. 907–923.
- Wise, R. (2011). *Al Shabaab*. Washington: Center for Strategic and International Studies.
- Wu, X., Zhu, X., Wu, G. and Ding, W. (2014). “Data mining with big data”. *IEEE Transactions on Knowledge and Data Engineering*, **26**(1), pp. 97–107.
- Zammit-Mangion, A., Dewar, M., Kadiramanathan, V. and Sanguinetti, G. (31st July 2012). “Point process modelling of the Afghan War Diary”. *Proceedings of the National Academy of Sciences*, **109**(31), pp. 12414–12419.
- Zar, J. H. (2014). *Biostatistical Analysis*. 5th ed. Essex: Pearson Education Limited.