Examining the extent to which hotspot analysis can support spatial predictions of crime

Spencer Paul Chainey

Thesis submitted in accordance with the requirements of the Degree of Doctor of Philosophy

University College London
Department of Security and Crime Science
London. UK

I, Spencer Paul Chainey, 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.

Acknowledgements

Over the past seven years I have received a huge amount of support from a number of individuals who have encouraged me to complete this thesis. Firstly, I am extremely grateful to my two supervisors, Professor Richard Wortley and Professor Paul Longley, for their support, guidance and feedback on the research as it has progressed and for reviewing drafts of this thesis. In addition, I would like to extend further thanks to Richard Wortley for relieving me of many departmental duties in order for me to have the time to complete this research. I would also like to thank Professor Gloria Laycock for her continual encouragement to complete this research, and for her support as a supervisor prior to her retiring as director of the Department of Security and Crime Science.

I also wish to thank the Metropolitan Police and Northumbria Police, and particularly the Safer Newcastle Partnership for supplying me with the data I have used in this research.

Finally, I wish to thank my wife, Victoria Chainey for her support, patience and encouragement over the seven years it has taken to complete this work. She more than anyone has seen what this research has meant to me. I dedicate this thesis to Victoria, and our two sons Oscar and Isaac.

Abstract

The premise that where crime has occurred previously, informs where crime is likely to occur in the future has long been used for geographically targeting police and public safety services. Hotspot analysis is the most applied technique that is based on this premise – using crime data to identify areas of crime concentration, and in turn predict where crime is likely to occur. However, the extent to which hotspot analysis can accurately predict spatial patterns of crime has not been comprehensively examined. The current research involves an examination of hotspot analysis techniques, measuring the extent to which these techniques accurately predict spatial patterns of crime. The research includes comparing the prediction performance of hotspot analysis techniques that are commonly used in policing and public safety, such as kernel density estimation, to spatial significance mapping techniques such as the Gi* statistic. The research also considers how different retrospective periods of crime data influence the accuracy of the predictions made by spatial analysis techniques, for different periods of the future. In addition to considering the sole use of recorded crime data for informing spatial predictions of crime, the research examines the use of geographically weighted regression for determining variables that statistically correlate with crime, and how these variables can be used to inform spatial crime prediction. The findings from the research result in introducing the crime prediction framework for aiding spatial crime prediction. The crime prediction framework illustrates the importance of aligning predictions for different periods of the future to different police and prevention response activities, with each future time period informed by different spatial analysis techniques and different retrospective crime data, underpinned with different theoretical explanations for predicting where crime is likely to occur.

Table of contents

List of figures

List of tables

1. Introduction

Hotspot mapping is a popular analytical technique that is used for identifying concentrations of crime. In practice, hotspot mapping is used as a basic form of crime prediction – it uses data on past incidents to determine where crime may occur in the future to inform where to target police and public safety resources. Yet, to date, very little research has been conducted that statistically examines the ability to predict spatial patterns of crime using hotspot mapping. Additionally, in recent years several new predictive crime mapping techniques have been developed, with the creators of these claiming they offer better predictions than hotspot mapping. However, as no accurate statistical benchmark exists that allows any new technique to be compared to hotspot mapping it is difficult to validate whether these new techniques are actually any better than hotspot mapping. The aim of this PhD research is to establish this benchmark by identifying how well hotspot mapping can produce accurate spatial predictions of crime. That is, the primary question the research aims to answer is to what extent can hotspot mapping be used to effectively predict where crime is likely to occur? It is argued that with careful attention to the data that are used and hotspot analysis technique that is applied, accurate spatial predictions of crime can be generated using hotspot maps. By establishing a benchmark measure that documents the extent to which hotspot mapping output can predict spatial patterns of crime, other spatial analysis techniques that are designed for crime prediction can be more effectively compared. It is anticipated that a comprehensive examination of hotspot analysis techniques will also identify ways in which these techniques can be refined to improve their use for spatial crime prediction.

The thesis extends existing research in six main ways. First, the research examines the spatial prediction performance of a number of popular hotspot mapping techniques. The most commonly used techniques include thematic mapping of aggregations of crime to administrative geographic units (e.g., police beats), and kernel density estimation that produces a spatial surface representing the density distribution of crime. The main distinction between hotspot mapping techniques is in the calculations performed to determine hotspots. Each technique also requires the user to enter certain parameters as inputs to these calculations. In addition, hotspot maps can be produced using different periods of retrospective crime data. The influence that different techniques, input parameters and the retrospective period of crime data that is used to create hotspot mapping output is not known. To establish the extent to which hotspot mapping can produce accurate spatial predictions of crime, the current research examines whether hotspot mapping techniques differ in their ability to predict where crime may occur in the future, if these results are influenced by the input parameters used, and if different retrospective periods of crime data identify different hotspots.

Second, the research examines whether an approach that uses statistical significance mapping offers an improvement in the performance of predicting spatial patterns of crime in comparison to the predictions produced from the commonly used hotspot mapping techniques. At present, a weakness with the commonly used techniques is that there is no agreed upon, objective method for determining the areas that are identified as *hot*. This can result in different users identifying different hotspots based purely on their own subjective selection. The research establishes if the existing commonly used hotspot analysis techniques can be improved by using a statistical significance mapping approach that removes the ambiguity in determining areas that are *hot*. The research also establishes whether a statistical significance hotspot mapping approach results in improvements in predicting where crime is likely to occur.

The thesis extends existing research in a third way by examining whether the spatial predictions produced using hotspot mapping are relevant for only a short period into the future or if the predictions are reliable for longer periods. Identifying hotspots of crime using retrospective data assumes these hotspot patterns will persist into the future. That is, it assumes hotspots are stable, changing little from the point they are identified to the pattern that is predicted in the future. However, little consideration to date has been given to whether predictions produced using hotspot mapping and other new prediction techniques, are equally accurate for different periods of the future – for example, the immediate future (i.e., the next day), the near future (i.e., the next week) and for other periods beyond. The research examines the stability of hotspots and determines whether the accuracy of the spatial predictions they generate is more relevant for certain time periods than others.

Fourth, the research extends existing research by illustrating the importance in identifying the theoretical reasoning that can justify any spatial prediction of crime. If spatial crime prediction techniques are to be effectively used in practice, the predictions that are made need to be supported with clear reasoning on *why* it is likely that crime will be committed in the areas identified. This can be informed from existing theory on why crime tends to happen at certain locations and from the empirical observations of patterns in offending behaviour, victimisation and the spatial-temporal attributes of crime. Establishing why crime is likely to occur at certain locations can then inform the police tactics and prevention programmes that aim to counter the anticipated offending behaviour and high level of vulnerability in these areas. The current research draws together the principal spatial theories of crime to create a foundation for helping explain why crime is predicted to occur at certain locations. The drawing together of these theories will permit a critique of the application of spatial analysis techniques for predicting where crime is likely to occur. That is, if the prediction made by a spatial analytical technique cannot be explained in clear theoretical terms, it may suggest the prediction is weak, and as a result it could be difficult to identify the specific tactics and programmes to counter the predicted activity.

Fifth, the research introduces several new theoretical perspectives on why hotspots exist following new analytical observations that the current research makes in the spatial patterns of crime. To date, environmental criminology research has made little use of spatial regression techniques for helping explain relationships between crime and other variables. In addition to a theoretical critique, a statistical means for explaining why hotspots exist at certain locations could provide additional value to the existing theoretical foundation. The research examines the use of spatial regression techniques for explaining why hotspots exist and evaluates whether existing theory provides a sufficient enough foundation for explaining the range of predictable spatial patterns of crime it is anticipated that this research will uncover.

Finally, the research introduces a new template for measuring spatial crime prediction performance. In order for mapping techniques to be compared for their predictive accuracy, a standard set of measures are required that allow for direct comparisons to be made between techniques. To date, some measures have emerged as research into spatial crime prediction has gathered pace, but no focussed attention has been given to developing a single measure or set of standard measures that can be applied in all cases. This current research critiques the existing measures used for assessing spatial crime prediction performance and introduces a standard measurement template that allows for accurate and detailed comparisons between mapping outputs.

1.1. Thesis structure

The thesis comprises thirteen chapters, including the current one. Examined in the following chapter is the progression from spatial theories of crime to crime hotspot analysis and spatial crime prediction. The starting point involves describing the key theoretical principles that explain the spatial distribution of crime and in turn inform the theoretical foundation on the use of mapping techniques for predicting where crime is likely to occur. The technical processes to hotspot analysis are then examined, including a description of the commonly used hotspot analysis techniques and an introduction to spatial significance mapping. In recent years, several new techniques have been introduced that claim to provide direct spatial predictions of crime. These are reviewed in the context of the earlier theoretical discussion to qualify the theoretical logic on which each is based. The last section of this chapter introduces spatial regression analysis for helping to identify variables that correlate with crime, and how this knowledge could inform improvements in spatial crime prediction.

Chapter 3 describes in detail the separate research objectives and hypotheses. The hypotheses are used to direct the structure and content of a series of empirical studies. Eight hypotheses are posed:

1. Hotspots can be identified using retrospective data for a short period of time rather than requiring retrospective data for longer periods of time
2. Common hotspot mapping techniques differ on how accurately they predict spatial patterns of crime
3. The technical parameters used in hotspot analysis techniques have an influence on the techniques' spatial crime prediction performance
4. Spatial significance mapping methods provide an improved means of predicting where crime is likely to occur in comparison to common hotspot mapping techniques, and removes the ambiguity of defining areas that are hot
5. Areas that are identified as hotspots of crime are places where the concentration of crime has been endured consistently for at least one year, and where the concentration of crime is likely to continue to persist into the future
6. Recent incidents of crime provide an effective means of accurately predicting the immediate future, but the accuracy in these predictions reduces for longer periods of the future

7. Geographically Weighted Regression (GWR) provides an effective means of determining at the local level the reasons why hotspots exist, and that these explanatory variables vary between hotspots
8. GWR analysis can be used for supporting long-term predictions of crime by examining how a change in explanatory variables can influence a change in future crime levels.

The description of the research methods is presented in two ways. Described in chapter 4 are the methodological approaches that are generic to the whole research. This includes a description of the data and study areas, the software used, statistical measures for prediction performance and the processes for applying the range of spatial analysis techniques. The empirical section of the thesis is organised into a number of studies, each forming its own chapter. The specific detail on the method used in each study is described in each empirical study chapter in turn. Results are also presented in each empirical study chapter rather than in a single results chapter.

The empirical studies constitute chapters 5 to 11. Each study informs the next, leading to a comprehensive set of results, and a progression in the development of the arguments that form from the research findings. Each empirical study chapter begins by setting out the chapter's aims and structure. Examined in chapter 5 is how hotspots can be statistically determined in crime data. The empirical study described in chapter 5 also begins the assessment of whether different retrospective periods of crime data may result in different predictions. In chapter 6, the method and results are described of a detailed metric examination of commonly used hotspot techniques to determine if different techniques produce different spatial prediction results. The analysis also examines if spatial prediction performance differs by crime type. Examined in chapter 7 is the influence of input parameters on the spatial crime prediction performance of hotspot mapping output. The focus here is towards examining the influence of the input parameters for the techniques that were identified as consistently producing better spatial predictions of crime from chapter 6. Chapter 8 reports on a detailed examination of the Gi* statistic – a mapping technique that identifies statistically significant concentrations of crime. The output produced using the Gi* statistic is compared to the results generated from the empirical studies reported on in chapters 6 and 7 to determine if the Gi* statistic unambiguously identifies hotspots of crime and if these hotspot areas are more accurate than the best predictions determined using common hotspot mapping techniques. The technique determined to consistently produce hotspot maps with the best predictions is then used in the remaining three research studies.

Examined in chapter 9 is whether hotspots are temporally stable. That is, once hotspots are identified in retrospective crime data, are these the areas where crime will continue to persist at high levels in the future? In chapter 10, the temporal dimension of spatial crime prediction is developed a stage further. Prospective mapping is one of the new techniques specifically designed for predicting future incidents of crime. In chapter 10, prospective mapping is examined for its ability to predict where crime is likely to occur for different time frames of the future. These results are compared to the hotspot mapping technique identified as the best performer from chapter 8. Chapter 11 is the final empirical study, examining whether spatial regression modelling and the use of other data variables (such as land use and demography) can determine why hotspots exist. The analysis in this final research study also aims to identify if these explanatory variables can be used alongside or as a replacement for retrospective crime data to improve spatial predictions of crime.

Chapter 12 brings together the results from the empirical studies and is where the technical, methodological, practice, policy and theoretical implications of the current research are discussed. Discussion includes identifying how the results contribute to the existing research, presenting arguments that use the results to build on findings from previous research, and suggesting areas for future research. Chapter 12 also includes a summary of the main findings in relation to each hypothesis and the research conclusions. Chapter 13 lists the research references.

1.2. Dissemination of research findings

As this PhD research has developed there have been opportunities to discuss the spatial analysis techniques that have been used, discuss certain methodological approaches that have been applied, and share several of the preliminary results. There has also been the opportunity to discuss theoretical, practice and policy implications of the research results. These discussions have helped to test some of the technical processes and preliminary findings with a number of practitioner audiences. Two journal papers have already been published containing findings directly relating to this PhD research. The peer review of publishing these papers has placed confidence in the methods used and in several results that have been generated. It has also helped refine the direction of the research to ensure it offers a valuable contribution to the field of geographical crime analysis.

The following lists publications and presentations relating to this PhD research over the period it was completed.

1.2.1. Publications

Chainey, S.P., Tompson, L., Uhlig, S. (2008), "The utility of hotspot mapping for predicting spatial patterns of crime", *Security Journal* 21:1-2.

- This paper has received 89 citations (Source: Google Scholar 23 May 2014)

Chainey, S.P. (2014), "Examining the influence of cell size and bandwidth size on kernel density estimation crime hotspot maps for predicting spatial patterns of crime*", Bulletin of the Geographical Society of Liege* 60:7-19

1.2.2. Conference presentations

The Crime Prediction Framework: a spatial temporal framework for targeting patrols, crime prevention programmes and strategic policy. The International Symposium on Environmental Criminology and Crime Analysis, Kerkrade, Netherlands, 2014

Predictive Policing: a spatial temporal framework for targeting patrols, crime prevention programmes and strategic policy. International Association of Chiefs of Police, World Innovation Conference, Amsterdam, 2014

Understanding hotspots. Australian Crime Mapping and Analysis Conference, Melbourne, 2012

Advanced hotspot analysis: spatial significance mapping using nearest neighbour analysis and the Gi* statistic. International Crime and Intelligence Analysis Conference, Manchester 2012

Identifying hotspots: a review of common techniques. 2nd European GIS and Law Enforcement Conference, Munich, 2011

Understanding hotspots using Geographically Weighted Regression. International Crime Mapping Research Conference, Miami, 2011

Understanding hotspots. International Association of Crime Analysts, Arlington, Texas, 2010

1.2.3. Training courses developed for the UCL Jill Dando Institute of Security and Crime Science that further field tests some of the PhD findings

- Predictive mapping www.ucl.ac.uk/jdi/short-courses/Predictive-mapping
- Understanding hotspots www.ucl.ac.uk/jdi/short-courses/understanding-hotspots
- Advanced hotspot analysis www.ucl.ac.uk/jdi/short-courses/adv-hotspot

2. From spatial theories of crime to analytical techniques that predict spatial patterns of crime

2.1. Chapter aims and structure

The starting point for examining whether any spatial analytical technique offers potential for predicting where crime is likely to occur is to determine the theoretical basis on which it may operate. A number of spatial theories of crime have been developed over the years. This chapter brings these theories together to help explain why geographical patterns of crime are not random and why crime can be predicted to take place at certain locations. With this theoretical knowledge, it is proposed that through the careful selection of appropriate spatial analytical techniques it is possible to accurately identify where these predictable crime patterns will most likely occur.

In places throughout this chapter (and in other parts of the thesis), reference is made to disease mapping and epidemiology. This on the basis that the study of crime patterns can be compared to the study of disease patterns in two main ways. First, in a spatial unit sense, incidents of disease are typically represented as geographic points within a GIS and are, therefore, subjected to similar spatial analysis techniques that can be applied to crime (Elliot et al., 2000). For example, the use of kernel density estimation for identifying geographic clusters of incidents was applied to the analysis of disease before it was applied to the analysis of crime (see Bithell, 1990). As the spatial unit for representing crime is similar to that used for disease mapping, and the spatial analysis techniques that are applied to these data are similar, lessons can be learned from studies into disease mapping that would benefit geographical crime analysis. Secondly, theoretical explanations for crime have usefully drawn from other disciplines, including epidemiology, for helping to explain the spatial patterning of crime and the spatial behaviour of offenders. For example, certain crime types such as burglary have displayed patterns similar to those found in the study of disease contagion (see Johnson and Bowers, 2004a)

The following section begins by examining the theoretical principles that underpin the geography of crime. These theoretical principles include macro geographic level explanations for crime (i.e., at the neighbourhood level), meso geographic explanations of crime (i.e., at the street level) and micro geographic explanations of crime (i.e., at specific locations) (Brantingham and Brantingham, 1984). The aim in drawing together the principal spatial theories of crime is to create a foundation that will permit a continual critique across the thesis of the spatial analysis techniques for predicting where crime is likely to occur.

Hotspot analysis techniques are considered to be the most basic form of spatial crime prediction. A review of the study of hotspots follows the section on the theoretical principles of the geography of crime, with the aim to define clearly what is meant by a hotspot and identify the practical role that hotspot analysis plays in helping to tackle crime problems. An examination of a number of hotspot analysis techniques is then conducted, presented with examples that illustrate how these techniques have been applied in practice. The aim of this examination is to determine those hotspot analysis techniques that are most commonly used and to highlight the technical distinctions between them. At present, little is known about whether these commonly used hotspot mapping techniques differ in their spatial prediction performance. Analysis of the spatial prediction performance of these techniques is subject to one of the studies that is reported on in this thesis. The chapter then introduces spatial significance mapping with the purpose of illustrating the potential it has for improving the spatial predictions made by the existing common hotspot analysis techniques.

In recent years, the concept of predictive policing has emerged, with several mapping techniques being introduced to support these predictive objectives. These techniques are reviewed in this chapter, drawing from the spatial theoretical principles of crime to critique each technique. It is argued that many of these new techniques are based on weak theoretical foundations and their accuracy in predicting where crime will most likely occur is questioned. Separate to these new prediction techniques is spatial regression. Spatial regression analysis techniques offer a statistical means of identifying those variables that may explain why spatial patterns of crime vary. These techniques are introduced and reviewed to assess if they offer potential in helping to determine why hotspots exist, and whether the results from this type of analysis can help inform predictions of crime.

Following the review of theoretical principles and spatial analysis techniques, the final section identifies that a gap currently exists in the use of standard measures for comparing the prediction performance of mapping outputs.

2.2. The key spatial theories of crime

If spatial predictions of crime are to be made with confidence, the prediction must be based on clear theoretical principles. If it is not clear why crime is likely to occur at a location, it would suggest the prediction is weak. If the theoretical argument for the spatial prediction is clear, this in turn will help identify the specific tactics and programmes to counter the predicted activity.

A useful starting point for the spatial theoretical principles of crime is to recognise that crime has an inherently geographical quality. When a crime occurs, it happens at a place with a geographical location (Chainey and Ratcliffe, 2005). This can be further observed in Brantingham and Brantingham's (1981) description of the four dimensions to every crime;

- the legal dimension (a law must be broken),
- the victim dimension (someone or something has to be targeted),
- the offender dimension (someone has to do the crime), and
- the spatial dimension (it has to happen somewhere).

Crimes also do not occur randomly. If crimes were a random occurrence that had an equal chance of happening anywhere at any time, there would be little point in attempting to observe patterns and predict where crime may occur in the future.

The main theoretical area that underpins the geography of crime is the practical subset of mainstream criminology, *environmental criminology*. Environmental criminology involves the study of criminal activity and victimisation and how factors of space influence offenders and victims (Bottoms and Wiles, 2002). The next section in this chapter begins by exploring how the importance of this spatial influence on people came to be recognised, before examining the spatial dynamics of offenders and the interaction of the offender and victim in space. The progression in theoretical development also helps to illustrate the evolution from macro towards meso and micro level geographic explanations of crime. It is argued that macro, meso and micro level theories are each of value, particularly when there is the need to explain why crime is likely to occur at certain locations in the immediate future and explain why crime may occur at certain places at some further point into the future.

Researchers have long known that there is variation in the spatial arrangement of crime. Although there have been recorded spatial studies of crime for nearly 200 years, a number of key research periods have punctuated the history. While these periods have overlapped across time, for convenience they can be thought of as three distinct schools of thought – the Cartographic School, the Chicago School and the GIS School. The theoretical developments associated with the Cartographic and Chicago Schools provide a useful foundation for macro explanations of why crime may occur at certain locations. In the more modern era, the focus of the GIS School has been more towards meso and micro level geographic explanations of crime.

2.2.1. The Cartographic School

One of the earliest maps of crime originates from France and was published in 1833 by Andre-Michel Guerry, who showed, amongst other features, the distribution of violent and property crime in the various départements of France (Guerry, 1833). These maps indicated that not only was there spatial variation in crime but that the risk of property crime and violent crime was often different in the same areas (Brantingham and Brantingham, 1981). Analysing French data around the same time was Adolphe Quetelet, who supplemented his maps with statistics showing spatial variations across France, as well as between different social groups, including beggars and smugglers (Quetelet, 1842). These early pioneers are credited with founding what is termed the Cartographic School (Chainey and Ratcliffe, 2005).

Throughout the 19th century, studies into the spatial arrangement of crime and criminals continued, one of the most renowned being an examination of the infamous London rookeries (Mayhew, 1862). These thieves' quarters were areas of high offender concentration, situated on the boundary of the City of London, where offenders were seen to exploit any cross-border policing difficulties between the City of London Police and the larger surrounding Metropolitan Police jurisdictional area. These early studies set the foundation for illustrating that crime patterns were not random, and prompted discussions on the factors that influenced the spatial distribution of crime.

2.2.2. The Chicago School

I. The socio-cultural triggers of crime and models for urban development

In the twentieth century, more innovation followed with the research conducted by the Chicago School. This group of researchers included Clifford Shaw and Henry McKay who drew on the spatial and temporal ideas of social ecology, forged by their predecessors, notably Ernest Burgess (Burgess, 1916). Shaw and McKay mapped, by hand, the residences of juvenile delinquents across Chicago (Shaw and McKay, 1942). This pioneering mapping was used to explore the socio-cultural triggers of crime as Chicago expanded during a period of great economic growth. They drew on Burgess's work by comparing socio-economic factors and physical factors in different zones across the city.

Burgess introduced the zonal (or concentric) model in 1925. Burgess's idealised model, as shown in Figure 2.1, had concentric zones radiating outwards in bands from the city centre, with each band representing a different stage of the city's development. The innermost zone (zone I), termed 'the loop', contained the central business district and had little residential development. Adjoining this was the 'zone in transition' (zone II), an area taken over by business and light manufacturing industries, and which also included the factory zone. The third zone ('zone of workingmen's homes') was occupied by factory workers who had managed to escape the zone in transition, but who were still tied to the city due to the need to work in the factories. Travel cost and time was a factor for these workers, so they resided in zone III. The 'residential zone' (IV) comprised high-class apartments or single-family suburban dwellings where the occupants accepted the travel costs as a price for a higher quality of life. Beyond the city limits was the 'commuters zone' (zone V) where people lived in suburban areas or satellite towns with a commute of up to an hour (Burgess, 1925). To demonstrate the model, Burgess charted 1920's Chicago and overlaid his model onto this city's expanding and vibrant cityscape (Burgess, 1925: 55).

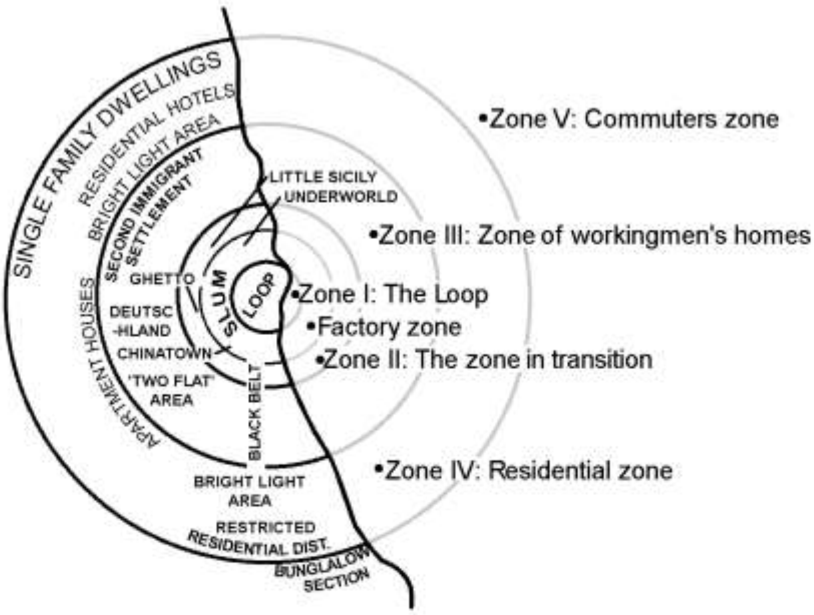

Figure 2.1. The concentric (zonal) circle model (Source: Chainey and Ratcliffe, 2005)

From a criminological perspective, the zone in transition was the area of most interest. Here mobility was greatest, the availability of stimulus peaked, and there was a concentration of 'juvenile delinquency, boys gangs, crime, poverty, wife desertion, divorce, abandoned infants, [and] vice' (Burgess, 1925: 59). Shaw and McKay (1942) went on to map the concentric zones using different bandwidths for different cities (e.g., in Chicago the bandwidths were two miles, while in Philadelphia they were one and a half miles). Their work spanned decades and formed the basis for much of American criminological thought coming into the latter half of the twentieth century, especially in establishing the longevity of delinquent areas discovered through longitudinal studies (Brantingham and Brantingham, 1981). In a macro sense, these studies led to the suggestion that crime concentrates in certain areas. The studies also suggested that these high crime areas persisted for some time, influenced by demographic and social mobility, and the availability and concentration of certain stimuli in these locations. From this, it is argued in the current research that if the concentration of these stimuli and mobility remain unchanged, crime will predictably continue to persist in these areas. In addition, if the stimuli and mobility factors that create favourable conditions for crime to occur can be identified, these may act as useful variables to predict where crime will likely occur.

The Burgess model worked well for American cities in the twentieth century, but was less applicable outside North America. The development of American cities took place over a fairly short temporal period, whereas European urban development occurred over a considerably longer time. Urban geographers have since developed models which help to better explain the mosaic pattern of development in urban areas outside North America. Knox (1994) describes an urban geography model of city expansion more fitting for a wider range of urban developments. These stages of neighbourhood change start with urbanisation before progressing to *in filling*, involving the population of vacant land blocks by multi-family units and rental properties. A stage of downgrading then follows when, as the housing stock ages, there is a slow decline of the area caused by degradation of the properties and general deterioration. Following this is a thinning out phase, characterised by a rapid population turnover and significant social and demographic change. The final stage is one of rehabilitation or gentrification.

From a crime perspective, the initial urbanisation brings in many young families, and while neighbourhoods may be harmonious for a few years, a suburb populated by a homogenous group, all arriving at the same time, will see many of the children reach their teenage (and highest crime-risk) years at the same time. Ten to fifteen years after the urbanisation stage (and often during the in filling and downgrading stages) it is possible that crime will therefore increase (Chainey and Ratcliffe, 2005). The in filling stage also has the capacity to increase crime by reducing the social cohesion of an area, due to the introduction of different socio-economic groups. This process continues to increase in the downgrading and thinning out stages due to social and structural neglect accelerating opportunities for property crime. These points offer further to the argument relevant in the current research that if the specific theoretical conditions that give rise to crime can be identified, determining where they geographically concentrate may offer value in helping predict where crime is likely to occur. The rehabilitation or gentrification stages may result in a reduction in crime, caused by the reintroduction of a more affluent population that can afford upgraded security features on cars and homes (Knox, 1994).

II. Social disorganisation and collective efficacy

A theory that grew from the Chicago School was social disorganisation. Social disorganisation theory posits the idea that increased levels of delinquency, especially juvenile delinquency, exist because of the lack of a local social fabric where the structure and culture of the community are strong enough to provide a concerted influence over local residents (Shaw and McKay, 1942). For example, social disorganisation theory suggests that if there is a high degree of cultural heterogeneity and a high turnover of residents, the community is unlikely to be able to agree to a common standard for behaviour on the street, and that few residents are likely to know the young people on the street who are causing trouble, or know their families. With no clear rules of acceptable behaviour and few sanctions available to curb adolescent exuberance (you cannot tell a child to stop misbehaving or you will tell his parents, if you and the child both know that you do not know his parents), juvenile delinquency increases.

The practical testing of social organisation theory has had its challenges because it can be difficult to construct variables that directly measure social disorganisation (Chainey and Ratcliffe, 2005). For example, it is unlikely that a household survey that asked the respondents to rate from one to ten the level of social disorganisation in their neighbourhood would be very successful, because few people would have a clear notion of what social disorganisation is. Additionally, even if social disorganisation could directly or indirectly be measured (using proxy variables), operationalising a policing impact on this issue in order to improve the crime situation would be quite a challenge. This is because policing strategies are often limited when faced with these more systemic causes of crime.

In response to the difficulties in measuring social disorganisation, some researchers have tried to measure its reverse, social or collective efficacy. Collective efficacy can be defined as the 'social cohesion among neighbours combined with their willingness to intervene on behalf of the common good' (Sampson et al., 1997: 918). Collective efficacy can be found in areas where neighbours cooperate on issues of mutual interest, share some areas of agreement with the people who live around them, and are prepared to intervene if local youths are behaving in a manner unacceptable to local norms. Such levels of cooperation require enough implicit or explicit communication between neighbours in order to define and agree the standard for local normative behaviour. It is, therefore, argued that areas high in collective efficacy are well suited to resisting crime at the local level by being able to influence local young people and exercise some control over a group in their peak offending years.

Collective efficacy has been directly measured using community-based surveys which have attempted to measure neighbourhood social and institutional processes. Collective efficacy is related to the notion of 'social capital', a feature that some researchers have operationalised as the number of interactions that take place with neighbours. Social capital is a measure of the skills and social position that a person possesses that provide them with the power to effect a positive social change on their local environment. Sampson and his colleagues have led the research in this area, actively seeking to measure collective efficacy. Their survey of over 8,000 Chicagoans included asking the respondents if they believed it was likely that neighbours could be counted on to intervene if children were spray-painting a local building, or if a fight broke out (Sampson et al., 1997). Their study also included census measures of race, poverty, and immigration, home ownership and residential stability (among others) and concluded that collective efficacy could be reliably measured and could act to control the level of violent crime. Sampson has built on these studies by further illustrating the relationships between crime and social conditions in Chicago (Sampson, 2012).

The review of macro-level (neighbourhood) spatial explanations for crime from the Cartographic and Chicago Schools suggests that crime can be predicted if the factors that create favourable conditions for crime to occur are identified. Although some difficulties may exist in measuring some of the variables that have emerged from the thinking in the Chicago School (and more latterly social disorganisation and collective efficacy), the identification of these factors may assist in informing the options for directing strategic policy that aim to achieve a long term, sustainable reduction in crime. The current research aims to examine this by identifying suitable spatial analytical techniques that determine the factors that create favourable conditions for crime to occur and whether, in turn, the empirical findings from this research inform spatial predictions of crime.

2.2.3. The GIS School

During the 1990s and 2000s there was an explosion of interest in environmental criminology, spatial crime analysis and the investigation of offender patterns using geographic tools. The catalyst for this enthusiasm has been attributed to the development of the ideas relating to defensible space, and the related principles associated with Crime Prevention Through Environmental Design (CPTED) (Brantingham and Brantingham, 1981, Jeffrey, 1971; Newman, 1972). CPTED in particular has grown into a significant discipline that addresses space management and architectural design, and urban planning (Crowe, 2000).

Building on the framework of CPTED and defensible space a number of important advances were made, most notably with:

- the theoretical developments of routine activity from Cohen and Felson (1979);
- the work of Rengert in his examination of residential burglary behaviour (Rengert and Wasilchick, 1985) and the spatial arrangement of drug markets (Rengert, 1996);
- the advances in crime pattern theory made by Brantingham and Brantingham (1981; 1984)
- the application of spatio-temporal analysis techniques to crime by LeBeau (1987; 1992); and
- the examination of crime across different spatial scales by Harries (1980).

These significant advancements helped to fortify the underlying theoretical principles of environmental criminology. However, it was the development of affordable geographical information systems (GIS) and the increasing technological developments within policing (such as the digitisation and geocoding of crime records) that have allowed crime researchers to exploit the wealth of data recorded by police agencies and map crime and calls for service[1]. The next section examines the theoretical developments of the new environmental criminologists, a group referred to as the GIS School (Chainey and Ratcliffe, 2005). These developments include the routine activity approach, the spatial arrangement of attractive targets, the rational choice perspective, crime pattern theory, the least effort principle, and the concepts of crime generators and attractors. These theoretical approaches do not just consider the crime offence, but also consider offender behaviour and the risk of victimisation. After all, if crime patterns are not random then offenders cannot be acting in a random way. The non-random nature to offending then suggests that the majority of offenders will act in a predictable manner, reacting in a similar way to the same opportunities to commit crime across space. In turn, this suggests that the patterns they individually or collectively create are predictable. Similarly, the risk of being a victim is unlikely to be random, with this non-random nature suggesting that certain individuals or groups of people will act in a predictable manner that makes them more vulnerable to crime. This suggests the patterns they individually or collectively create are also predictable.

[1] Affordable GIS and the digitisation of records has led to similar developments in the application of spatial analysis in many other disciplines, including epidemiology, demography, land use, transport, and other emergency services such as fire and rescue.

I. The routine activity approach

The routine activity approach originally started as a macro-level explanation of predatory crime (Cohen and Felson, 1979), but has progressed over the years to provide a worthwhile mechanism to examine criminal opportunity and crime prevention in a variety of settings. The original work examined changing patterns of employment, and the new criminal opportunities that are created when there are fewer people staying at home during the day. The routine activity approach is based on the simple idea that the behaviour of victims helps explain the occurrence of crime, and that for a crime to occur, three components are necessary. There must be the presence of a likely offender, the presence of a suitable target, and the absence of a capable guardian. The target does not necessarily have to be a person, but instead could be buildings, cars, or other property and objects. Similarly, a guardian may not just be a person, such as a police officer, security guard, neighbour or even passer-by, but could also include closed circuit television surveillance systems or a car alarm. The three components – offender, target and lack of a guardian – must meet in time and space to provide the necessary *chemistry* for crime (Felson, 1998). This meeting in time and space is not random but is dictated by the natural rhythm of daily life – people going about their *routine activities.*

The routine activity approach does not just discuss offenders, targets and guardians, but adds the important qualifiers. Not all offenders are *likely* offenders, as some will lack the technical knowledge and skill to attack certain types of premises. Similarly, not all targets are *suitable* targets, as they may be inaccessible (such as rooftop apartments) or too well defended. Many objects and people can be guardians, however at different times they may not be *capable* guardians. The routine activity approach can be summarised by the following simple equation:

Likely offender + suitable target – capable guardian = crime opportunity (1)

The combination of these three components and their qualifiers then dictates that the risk of crime changes over time with the movement of people throughout the daily routine activities of their lives. In the context of the current research, as these meetings in space and time are not random, where and when crime opportunities are most present can be argued to be predictable.

Since Cohen and Felson's original work, Eck recast the concepts of the routine activity approach into the crime triangle (Eck, 1995) (see Figure 2.2). The crime triangle introduced 'place' as the third side of the triangle, with offenders and targets/victims forming the other two sides. This recasting of the routine activity approach also led to the expansion of the original concept of guardianship by introducing the general term of *controller*. The concept of controllers is illustrated in Figure 2.2 by the positioning of guardians in relation to victims and targets. Controllers were then added to the offender and place sides of the crime triangle. For offenders, Felson (1995) introduced the concept of *handlers*. A handler is a person, a third party, who can influence the behaviour of the offender. For example, a parent may be a handler, as may a teacher or any other person who knows or who could determine the name of the person, and whose respect the offender might not wish to lose. For the place side of the crime triangle, the concept of *place managers* was introduced (Eck, 1995). A place manager is someone who is able to control a place even if they are not formally in charge of the area.

Figure 2.2. The crime triangle

In practice the crime triangle helps to focus analysis towards the causes of crime, from a routine activity perspective, and the mechanisms that can influence those causes (Chainey and Ratcliffe, 2005). The inclusion of controllers around the three sides of the crime triangle then helps to illustrate that if criminal opportunity can be predicted and theoretically explained, activities that involve one or more types of controllers could counter and minimise the crime opportunity.

The review of the routine activity approach, therefore, indicates that the crime opportunities that emerge from the coming together of offenders, targets and guardians, dictated by the rhythm of daily life, are predictable in space and time. Following this predictable behaviour, the occurrence of crime could also be prevented: as Felson states (1995: 55) 'Crime opportunity is the least when targets are directly supervised by guardians; offenders, by handlers; and places, by managers'. Therefore, if spatial analysis techniques can observe geographical patterns in crime that can be explained in routine activity terms, and are predicted to occur again because the chemistry of crime continues to persist, the routine activity approach can also be used to help identify effective means for removing these predictable criminal opportunities.

II. The spatial arrangement of attractive targets

Target suitability can change over space and time. While the routine activity approach provides a general model with which to consider the likelihood of crime occurrence, available targets can differ in their attractiveness to a criminal. Clarke's (1999) *CRAVED* summary of the basic characteristics of a suitable, or *hot* target provides a useful framework for determining the types of property that are at more risk of being stolen:

- *Concealable* – things that can easily be concealed after being stolen are at more risk of being stolen
- *Removable* – the easier a thing is to remove, the greater the chance it will be a hot target
- *Available* – objects that are plentiful, visible and not secured are most at risk
- *Valuable* – the more valuable an item is deemed to be, the higher the risk of it being a target for theft
- *Enjoyable* –the degree to which a product can be enjoyed marks its potential theft risk, an indication of both its value to the offender and possibly its resale value on the stolen goods market
- *Disposable* – many items are stolen so that they can be sold or traded to others, therefore disposability is an important characteristic for stolen goods.

Ratcliffe (2002) showed how the ideas of CRAVED products and the routine activity approach could help explain vehicle crime in Sydney, Australia. Ratcliffe's analysis identified that vehicle crime in the affluent suburbs near the beaches was highest during the overnight periods and that expensive cars were the main targets. The targeting of these cars fitted the CRAVED criteria, not only because the vehicles were valuable, but also because few residents in these areas had private garages (making them easier to remove). Vehicle crime was concentrated in the overnight period because the cars were suitable targets and more *available* to thieves at night (the cars were being used by their owners during the day, but were *available* on driveways or parked on the street at night), and the night time hours provided for few capable guardians.

The opportunity for crime and the spatial arrangement of targets using the routine activities and CRAVED principles suggests that if patterns from recorded crime data identify hot targets, it is likely these targets will continue to be high risk items for theft unless one or more features of their CRAVED profile are addressed. As the distribution of these attractive targets is not random, identifying where they are most prevalent would assist in predicting where crime is likely to occur. The challenge, therefore, that the current research aims to address is whether spatial analytical techniques can reveal where these CRAVED targets geographically concentrate. In addition, because the reason for their attractiveness to theft can be explained, the areas where these attractive targets are most prevalent is where prevention initiatives that counter their CRAVED appeal could be of most benefit. Thinking in terms of these theoretical principles again illustrates the important role these principles play in qualifying spatial predictions of crime and helping direct the analyst towards potential ways the predicted crime activity can be countered.

III. Rational choice

The rational choice perspective (Clarke and Felson, 1993; Cornish and Clarke, 1986) provides a framework to consider offender decision-making when a crime opportunity is presented. It can also be used to consider likely strategies that will influence the decision-making of the offender. Most offenders are known to make some sort of decision to commit a crime by weighing up some of the pros and cons (i.e., the rewards, against the chance of being caught) (Bernasco, 2010; Clarke and Felson, 1993; Cornish and Clarke, 1986). This suggests committing a crime is a (fairly) rational decision, and that an offender will commit an offence while trying to achieve some sort of desire or goal (Kennedy, 2009). The goal may be to derive personal gain, as in burglary or theft, or personal pleasure as in the crime of joyriding. For example, Rengert and Wasilchick noted after their interviews with 31 burglars, ‘the decision to commit burglaries was a purposeful, rational decision in almost every case’ (2000: 60). If the legitimate means of obtaining the offender’s goal are not available, then a decision may be made when a criminal opportunity becomes available. The decision may not be one that is fully calculated, as offenders may not weigh up all of the consequences.

It is argued that criminal decision-making is in two parts. There is a long-term, multi-stage decision to become generally involved in criminal activity (criminal involvement decision) and a shorter-term, more immediate decision (the criminal event decision) to grasp an opportunity that is presented (Cornish and Clarke, 1986). Factors such as drink, drugs, peer-pressure, or limited education do mean that not all offenders' decisions are purely rational, resulting in what has been termed *limited rationality*, also known as *bounded rationality* (Newman, 1997). These terms are an acceptance from researchers that some offences are committed with less-than-military planning behind them. Although the effects of drugs and alcohol can limit the rationality of offenders, the immediate decision-making of a burglar (for example) is primarily based on the environmental cues from the prospective target that can change from place to place: can the offender be seen breaking in, is there anyone home, and is there an easy way into the house (Cromwell et al., 1999)?

On its own, the theoretical principles of rational choice do not necessarily determine the spatial patterning of crime. The main exception to this is at the micro geographical level where the criminal event decision may lead to an offender choosing, for example, a specific house to burgle over a neighbouring property. However, armed with an understanding of the influence of routine activities, the chemistry for crime, the spatial arrangement of targets, and the decision-making of an offender, a model for the interaction between offenders and victims can now be described – crime pattern theory.

IV. Crime pattern theory

While the routine activity approach provides a model to predict if a crime has all of the chemistry to occur, and the rational choice perspective enables us to determine some of the thinking behind an offender's ultimate decision to commit a crime, crime pattern theory can help explain *where* and *when* crime is likely to occur. Crime pattern theory helps to bring together the two areas of offender spatial distribution and offence spatial distribution by examining the 'relationship of the offence to the offender's habitual use of space' (Bottoms and Wiles, 2002: 638). Crime pattern theory (sometimes also referred to as offender search theory) suggests that offenders are influenced by the daily activities and routines of their lives, so that even if they are searching for a criminal opportunity, they will tend to steer towards areas that are known to them (Brantingham and Brantingham, 1984). In an offender's day-to-day activities they will be watching for targets that have no guardians or place managers. Two important theoretical concepts associated with crime pattern theory are the least effort principle and the distinctions between how crime opportunities concentrate. These two theoretical concepts are considered after an explanation of the other main theoretical principles that are associated with crime pattern theory.

Like offenders, we all have various routine activities in our lives. Most of us have to go to work, college or school, and we usually go there from home. We may also go to shops and restaurants, bars and cinemas. These places from which we carry out the majority of activities are termed *nodes*, and connecting these nodes are *pathways* (the routes we take between our nodes). These repetitive activities around our nodes and the journeys along the pathways create within us a 'cognitive map' (Brantingham and Brantingham, 1984: 358) of places, routes and associations. Over time, these cognitive maps become a general list of well-known areas in which we feel comfortable. This environment consists of not just the physical things, such as buildings and train stations, but also the social and economic infrastructure through which we pass. Cities become an urban mosaic to us, places where we have no knowledge, interspersed with well-known places. We also become familiar with the routes between these known areas. These islands of knowledge, and the routes that link them, become our 'awareness space' (Brantingham and Brantingham, 1981: 35; Rengert and Wasilchick, 2000: 61).

Like us, offenders also have awareness spaces. They also move between places such as work, school, shops and home, and for many offenders, the search for criminal opportunities takes place around these areas. Opportunities are not spaced evenly throughout the landscape, and some offenders will only be able to take advantage of some offence opportunities. Also, some of the awareness areas will not be conducive to crime due to the presence of guardians or place managers. Therefore, for each offender we can generate a model of awareness space and criminal opportunity space (with the implicit absence of guardianship), and where they intersect we will find the areas of crime occurrence. Crime pattern theory is, therefore, strongly connected with the interactions of criminals and their physical and social environments.

The theoretical principles of crime pattern theory were modelled by Brantingham and Brantingham (1984) and an adaptation of the principles by Chainey and Ratcliffe (2005) can be seen in Figure 2.3. This shows that awareness space consists mainly of the places that are routinely frequented and the routes between those places. We may be more familiar with distant places that we frequent regularly than local places just around the corner that we never visit; therefore, proximity does not always mean the same thing as familiarity.

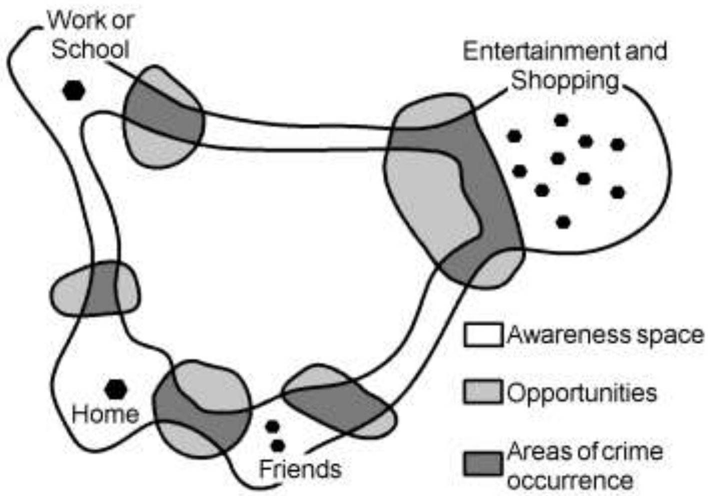

Figure 2.3. Hypothetical model of the creation of criminal occurrence space where offender awareness space and opportunities coincide (Source: Chainey and Ratcliffe, 2005)

The adaptation by Chainey and Ratcliffe (2005) of the original Brantinghams' model shown in Figure 2.3 includes the location of friends as an influential component on offender awareness space. Although for most people work or school play a significant part in daily life, this can be less so for offenders. Costello and Wiles (2001) found that many offenders in Sheffield who could have been in the workforce had never had full time employment, while Rengert and Wasilchick (2006) reported interviews with a number of offenders who had quit legitimate employment in order to pursue a professional criminal career. These examples focused on residential burglars only, however, it is likely that the attractiveness of both lifestyle and income of a life of crime are greater than poorly-paid legitimate employment for many burglars, drug dealers and robbers (Chainey and Ratcliffe, 2005). Having a place of work is not a necessary requisite of forming an awareness space. The Sheffield study found that as many offenders had never been employed, they were often transient and heavily influenced by the location of friends and criminal peers (Costello and Wiles 2001).

Figure 2.3 also shows that some opportunities are outside the awareness space of the offender, and are essentially unavailable, while those opportunities within the awareness space may have variable attractiveness. For example, for a burglar, all homes within their awareness space are theoretically potential targets, but some are more attractive than others. The areas with a greater proportion of targets that are potentially safer to burgle and more profitable will form part of the offender's *search space* – shown in Figure 2.3 as the opportunities area. Within this search space, certain targets will be more attractive than others, and the micro-search for a particular target occurs in the actual area targeted. Where an offender is most likely to commit crime based on their awareness space and the distribution of opportunities will, therefore, be in the area shown in Figure 2.3 as the areas of crime occurrence.

It should be noted that crime pattern theory is a general explanation of the crime patterns of offenders, and there will always be exceptions. Studies have found that a few offenders do not commit crime within their own awareness space, often as a result of the influence by other people who direct the offender to new opportunities (Rengert and Wasilchick, 2000), or by peers who introduce the offender to new areas (Wiles and Costello, 2000). When this happens, the spatial pattern of offences can be very different. However, for the majority of offenders, the tendency is to offend in their awareness spaces.

There are several other reasons why offenders might commit offences in familiar areas. It is helpful to know the layout of an area so if a quick getaway is required, the offender can avoid running straight up a dead-end street. Secondly, it has been suggested that offenders value feeling comfortable in an area and not feeling as if they stand out. This feeling of comfort has been suggested as a reason why offenders who live in poor neighbourhoods do not often commit offences in affluent areas (Rengert, 1989). However, if a more affluent area immediately neighbours a poor area where offenders reside, the familiarity between the borders of the two may result in higher crime rates (Hirschfield et al., 1995). In a similar vein to the *poor-rich* distinction is the *white-black* neighbourhood distinctions found in the burglary patterns of offenders in the United States. A number of studies have noted that black offenders avoid white suburbs and white offenders steer clear of black neighbourhoods, each group feeling that the avoided areas were unsafe (Rengert and Wasilchick, 2000; Wright and Decker, 1994).

An important concept that is central to crime pattern theory is the general theoretical principle of least effort (Zipf, 1949). In least effort terms, physical space can be likened to a friction surface in that it requires effort to cross it. The further the distance to travel, the greater the cost in time and possibly money, with people usually exerting the minimum effort possible to complete their tasks. The least effort principle, therefore, has an influence on the spatial behaviour of offenders by explaining why there are spatial limits to the geographical coverage of an offender's awareness space and why an offender chooses a particular location to commit crime. For example, the least effort principle explains why in Figure 2.3 there are boundaries to the extent of the offender's awareness space and their knowledge of criminal opportunities.

While increased distance to commit crime increases the effort, it also increases the risk, and increases the possibility that the offender will stray into an unknown area. This helps explain why an offender's desire for spatial exploration to extend criminal opportunities is rare (Rossmo, 2000). Offenders are often constrained by time and financial resources and lack the freedom to explore other opportunities or to search further afield for fresh prospects. The least effort principle can also be used to explain that if opportunities for an offender to commit crime are equally available in two areas, yet one area is much further away, the offender is more likely to choose the nearest offence opportunity in order to minimise their effort in crime commission (Chainey and Ratcliffe, 2005). This theoretical concept of minimising the distance travelled to commit crime is supported by a wealth of empirical research that has shown that the majority of offender journeys to crime tend to be short (Chainey and Ratcliffe, 2005; Rossmo, 2000). For example, Wiles and Costello (2001) showed that offenders of burglaries to residential properties in Sheffield on average travelled 1.9 miles from their home to their crimes. The least effort principle is, therefore, a useful mechanism for thinking about the geographical extent of an offender's spatial behaviour and the crime patterns that may result. In spatial crime prediction terms it can lead the analyst to hypothesising that the person who committed an offence is likely to be local and familiar with the area, or in aggregate form, that the many offences committed by individuals in the same area are most likely to involve offenders who live close to or have a good degree of familiarity with this area.

An offender's choice of location to commit a crime and the patterns that result can additionally be explained in terms of how opportunities to commit crime concentrate (i.e., how the opportunity space, shown in Figure 2.3, forms). The manner in which crime opportunities concentrate can be explained in two main ways - crime generators and crime attractors. A *crime generator* is a particular area where large numbers of people or other targets are drawn for reasons unrelated to criminal motivation (Brantingham and Brantingham, 1995). These places are crime generators because they provide times and places where there are many opportunities for offenders. For example, busy train stations and shopping areas may experience high levels of crime due principally to the large number of place users and targets at these locations (Clarke and Eck, 2003). By comparison, *crime attractors* are places which create criminal opportunities and in doing so, attract motivated offenders (Brantingham and Brantingham, 1995). The lure of a known criminal opportunity draws offenders to the area, enticing them with the knowledge that the area has a reputation for a particular type, or types of illicit opportunity. For example, areas close to schools provide an attractive location for street robbers, looking to prey on lone school children on their walk home and robbing them of their smartphone.

Crime generators and crime attractors, therefore, offer definitions to distinguish between areas where crime is likely to concentrate. They are typically used in a meso and micro geographic sense to explain that the high levels of crime are due to many targets being available or specific opportunities exist. However, they do not necessarily explain in an overarching theoretical sense the reasons why potential targets and opportunities can be found at certain locations. Additionally, the focus of crime pattern theory is towards explaining the criminal landscape from an offending viewpoint, rather than from a risk of victimisation perspective. To some degree, the spatial arrangement of attractive targets provides some explanation by considering where hot targets are concentrated, but again this does not necessarily explain why there are so many targets in crime generator areas, or why there are specific opportunities to target CRAVED items in crime attractor areas. In addition, a hot target explanation is only suitable for theft offences and not for offences against the person, such as assault. It is, therefore, suggested that a gap may exist in existing environmental criminology theory for explaining why certain areas experience high levels of crime.

A complementary principle to the existing environmental criminology theory that the current research considers is that of geographic areas exhibiting *favourable conditions* where many suitable or specific victims or targets concentrate. It is argued that this theoretical consideration of favourable conditions for victimisation could not only be used to help explain high levels of crime in meso (street level) and micro (specific locations) geographic settings, but also at the macro (neighbourhood) geographic level. The current research examines this possible theoretical gap further, and identifies whether an expansion to existing environmental criminology theory is required that may help explain why certain places possess favourable conditions for high levels of crime. In the first instance this requires a detailed examination of crime patterns to identify if existing theory is sufficient for explaining the spatial patterns of crime and explaining where crime is predicted to occur in the future.

Crime pattern theory provides a model to help explain why an offender chooses a particular location to commit a crime. It illustrates that this decision-making is informed by the offender's awareness of the area and the familiarity of opportunities, and is often influenced by their previous experience of crime commission. Crime pattern theory, therefore, provides a further theoretical explanation for why spatial patterns of crime are not random and can be predicted. While its theoretical focus is on explaining the geographical patterning of crime committed by individual offenders, the aggregation of many offenders being attracted to similar targets that are spatially concentrated means the theory can also help explain these aggregate geographical patterns of crime. The challenge for the current research is to ensure that spatial analysis techniques can accurately identify these aggregate geographical patterns. Crime pattern theory can then be used as a potential source for explaining why crime occurs in these locations and why crime is likely to continue to occur in these locations.

2.2.4. Initial conclusions and gaps in the existing research: the contribution of spatial theories of crime for spatial crime prediction

Crime pattern theory and the least effort principle provide a framework for helping to understand the criminal spatial landscape. The awareness spaces referred to in crime pattern theory are made up of different structures which provide a changing pattern of opportunities to commit offences depending on the environmental backcloth – the social, psychological, economic, physical and temporal mosaic of the offender passing through the landscape. These different structures change across space and affect the type of criminal opportunities that are available to the motivated offender. The motivations of the offender are, though, bounded by the rationality in their decision-making. Offenders prefer to commit offences in areas where they are comfortable, and by their preference to minimise effort. The theoretical explanations for the criminal spatial landscape, therefore, provide the source for helping to interpret the spatial patterns observed in crime data and predict why crime may occur in certain locations. The spatial theories of crime from the GIS School do, though, mostly refer to explaining the actions of individuals. However, the aggregation of this theorised individual behaviour can be used to inform the predictable criminal spatial behaviour of many offenders. Additionally, while the explanations for the non-random nature to the geographic distribution of crime that have developed from the Cartographic and Chicago Schools do not explain the individual actions of offenders and the characteristics of suitable targets, there is potential that these macro-level explanations can contribute to informing spatial crime prediction at the neighbourhood level.

Concepts such as the routine activity approach, rational choice, crime pattern theory, the least effort principle, crime attractors and crime generators are valuable in helping to conceptualise the underlying criminal landscape and the spatial decision-making processes of offenders. However, a theoretical gap may currently exist in environmental criminology for explaining why certain places experience favourable conditions for high levels of victimisation. This includes explaining these favourable conditions at the micro, meso and macro geographic levels. The current research examines whether this theoretical gap exists in environmental criminology and critiques how the contribution of a theoretical principle for explaining why places experience favourable conditions for crime can inform spatial crime prediction.

The inherent spatial quality to crime, and the knowledge that crime tends to occur at predictable locations rather than being randomly distributed, now directs the research attention towards identifying analytical techniques that can identify and help interpret these spatial patterns. These analytical techniques include hotspot analysis, new predictive mapping techniques such as prospective mapping, and spatial regression. The current research examines these spatial analysis techniques, qualifies their application and whether they provide an accurate means of predicting where crime is likely to occur. The following section reviews each technique and illustrates how, to date, they have been applied in policing and public safety.

2.3. Crime hotspots

This section defines what a crime hotspot is, describes the role of hotspot analysis in policing and public safety, and introduces the commonly used hotspot analysis techniques. The section also examines whether these hotspot analysis techniques have been subject to appropriate examination of how accurate they are for predicting spatial patterns of crime.

2.3.1. The definition of a hotspot

A hotspot is defined as an area of high concentration of crime relative to the distribution of crime across the entire study area (Chainey and Ratcliffe, 2005; Home Office, 2005; Sherman, 2009). In these terms, hotspots can exist at different geographic scales of interest, whether it is at the city level for examining localities where crime is highest, or at a local residential housing estate level, identifying particular streets or clusters of buildings where crime is seen to highly concentrate.

2.3.2. The role of hotspot analysis for helping to tackle crime problems

The mapping of hotspots of crime has become common practice in police agencies across the world and has been applied to many forms of crime. For example, hotspot mapping has been used in the analysis of residential burglary, street robbery and vehicle crime in London (Eck et al., 2005), gang-related murders in Belo Horizonte, Brazil (Beato, 2008), violent crime in Philadelphia (Ratcliffe et al., 2011), and street assaults in Melbourne, Australia (Mashford, 2008).

The primary purpose of hotspot analysis is to visually identify where crime tends to be highest, and from this aid decision-making on where to target and deploy resources to tackle the crime problem. The ability to be able to identify where crime concentrates, and direct law enforcement and crime prevention activity to these areas, is used routinely in policing and public safety for helping to address crime problems. Examples include the following:

- Hotspot analysis is routinely used in policing to support the operational briefing of police patrols (Goldsmith et al., 2000; Harries, 1999; Home Office, 2005; LaVigne and Wartell, 1998, 1999; Osborne and Wernicke, 2003). Hough and Tilley (1998) illustrate an example of hotspot maps that were used in police patrol briefings by the Metropolitan Police in the London Borough of Brent. These hotspot maps identified recent events (i.e., in the last four days) against a backdrop (hotspot map) of the distribution of crime over the last month.

- A more focused use of hotspot analysis for informing the targeting of police patrols is hotspot policing. This involves the dedicated targeting of police patrols to high crime locations (Sherman et al., 1989; Weisburd and Braga, 2006) where the patrols typically remain stationary for 15 to 20 minutes (and returning each hour for the duration when crime is determined to be high). This patrolling approach aims to create a visible police presence that emphasises a higher certainty of detection if an offender decides to attempt an act of crime (Kennedy, 2009). In the first documented example of hotspot policing in Minneapolis, USA, doubling the number of patrols at hotspots reduced crime by 6-13%, and disorder by 50% (Braga, 2007). In a more recent study of hotspot policing tactics in Philadelphia, targeted patrols reduced violent crime by 23% (Ratcliffe et al., 2011).
- The use of *intelligence products* has become a key component in modern policing (Chainey, 2012; Chainey and Chapman, 2013). The production of intelligence, requiring the gathering of information and its interpretation, is fundamental to the intelligence-led policing approach, with its use at the core of informing police business and decision-making (Ratcliffe, 2008). Intelligence products include tactical assessments for informing the regular briefing, tasking and coordination of front line resources; problem profiles that aim to help better understand particular crime problems; and strategic intelligence assessments that identify the key issues and threats in a police command area to determine the strategic priorities that require specific attention. Hotspot analysis is a common feature within these products, informing operational tactics, helping to identify persistent problem areas and support a strategic targeted approach for crime reduction (Chainey and Ratcliffe, 2005; Clarke and Eck, 2003; Home Office, 2005). An example of hotspot analysis that was used to assist the strategic targeting of crime reduction initiatives was illustrated by Chainey and Chapman (2013) for Newcastle-upon-Tyne. This example showed how hotspot analysis of youth-related anti-social behaviour assisted in identifying where to focus strategic interventions for tackling this issue in certain persistently problematic areas.
- Hotspot analysis has also been a common feature in Compstat-style performance meetings (Chainey and Ratcliffe, 2005; Harries, 1999; McDonald, 2002; Schick, 2004; Walsh, 2001). Compstat is a policing management meeting process, with meetings held regularly (from every week to every month depending on adoption by

police agency) to inform on recent events and trends, determine future activity, and build in a process of accountability to these activities. The use of hotspot analysis is a regular feature in these meetings for helping to review performance and determine future actions. Figure 2.4 shows an example of the use of hotspot mapping by Thames Valley Police in their Compstat meeting. The meeting is chaired by the Chief Constable and is attended by his deputies and police commanders who are each responsible for policing and crime reduction in one of Thames Valley's ten command areas (Home Office, 2005).

Hotspot analysis has therefore become a regular application in policing and public safety, helping determine where crime may happen next by using data from the past to inform future actions. In this sense, it acts as a basic technique for predicting where crime may occur, using the premise that retrospective patterns of crime are a useful indicator of future patterns. Similar applications of hotspot analysis have been used in epidemiology, using data on where incidents of disease have occurred in order to target (and hence predict) where other incidents of disease are likely to occur in the future (e.g., Atkinson and Molesworth, 2000).

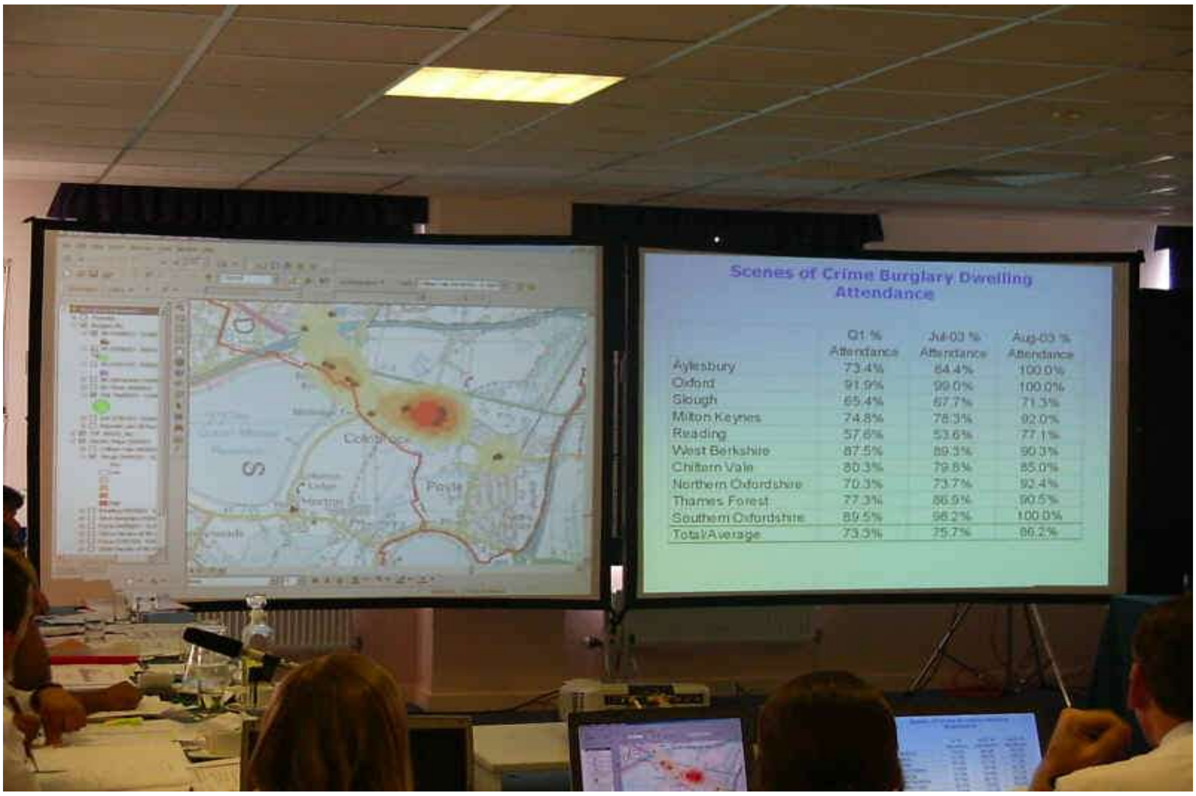

Figure 2.4. Hotspot analysis used in Thames Valley Police's Compstat meeting

2.3.3. Common approaches to crime hotspot analysis

There are many mapping techniques that can be used for identifying patterns of crime. These techniques could be as straightforward as representing each crime event as a point and observing the geographic distribution of these points; utilising functions within a GIS for thematically shading administrative areas (e.g., Census zones or police beats); or representing the distribution of crime as a continuous surface that relates to the volumetric densities of the geographic distribution of crime.

Identifying hotspots using point mapping has become outdated since the proliferation of GIS software and the increasing sophistication of mapping techniques. There are four techniques for hotspot mapping that have been regularly used since the widespread adoption of desktop mapping software tools and GIS (Chainey and Ratcliffe, 2005). These are the following:

- Spatial ellipses
- Thematic mapping of administrative areas
- Thematic mapping of grids, and
- Kernel density estimation

In the next section, each of these four commonly used hotspot mapping techniques are reviewed to illustrate their application and examine whether they have been subject to research that has measured their accuracy for predicting spatial patterns of crime.

I. Spatial ellipses

One of the earliest crime mapping software applications that became widely available to practitioners for hotspot analysis was Spatial and Temporal Analysis of Crime (STAC) (Illinois Criminal Justice Information Authority, 1996). STAC is not a GIS, but instead acts as an aid to researchers who already have a GIS or desktop mapping capability. STAC is a spatial tool to find and examine hotspot areas within the study area. STAC works by first finding the densest concentration of points on the map (hot clusters), and then fitting a standard deviational ellipse to each one. The ellipses themselves indicate through their size and alignment the nature of the underlying crime clusters.

Examples demonstrating the use of STAC include Martin et al. (1998) in their study of crime patterns in Detroit; Bowers and Hirschfield (1999) who explored links between crime and disadvantage in North West England; Block and Block (2000) who analysed hotspots around rapid transit stations in Chicago; and Langworthy and Jefferis (2000) who examined the influence of the school holiday period on the spatial distribution of crime hotspots in the Bronx, New York. Baltimore County Police Department has also used STAC extensively to analyse a range of crime types (Block and Perry, 1993).

Reported benefits of using STAC include that it derives hotspots without relying on defined boundaries such as census units or police administrative boundaries; that it requires few parameters; and that it is compatible with most GIS applications (Martin et al., 1998). However, STAC has attracted criticism for several reasons. Firstly, it is preferable for the user to be well versed in the routines at work within the software. For the novice, there is little counsel on appropriate parameter values and this leads to the introduction of ambiguity and increasing variability in the results (Eck et al., 2005). Secondly, crime hotspots do not naturally form into convenient ellipses, and thus STAC hotspots do not represent the actual spatial distribution of crime and can often mislead (Eck et al., 2005; Ratcliffe and McCullagh, 2001). Finally, the visualisation of the STAC-produced results negates any comparison with events that do not fall into the spatial ellipses (Eck et al., 2005). An example of STAC-produced spatial ellipses is shown in Figure 2.5b. To date, the STAC technique of representing hotspots as spatial ellipses has not undergone any analysis that determines how accurate it is for predicting spatial patterns of crime[2].

II. Thematic mapping of geographic boundary areas

A widely used approach for representing spatial distributions of crime events is geographic boundary thematic mapping, or choropleth mapping (Home Office, 2001). Boundary areas that are used for this type of thematic mapping are usually arbitrarily defined for administrative or political use. For example, geographic boundary areas can be police beats, census output areas, wards or districts. Offences as points on a map can be aggregated to these geographic unit areas and then shaded in accordance with the number of crimes that fall within them. Williamson et al. (2001) assert that maps created in this way are quick to produce and require little technical expertise to interpret. Furthermore, this technique enables the user to quickly determine which areas have a high

[2] Results from the current research that examine STAC and other commonly used hotspot mapping techniques have already been published in the Security Journal: Chainey, S.P., Tompson, L., and Uhlig, S. (2008), "The utility of hotspot mapping for predicting spatial patterns of crime", *Security Journal* 21:1-2..

incidence of crime, and allows further diagnosis of the problem by *zeroing in* on these areas. In addition, census areas can easily be linked with other data sources, such as population, to calculate a crime rate, so increasing the hotspot map's versatility for analysis.

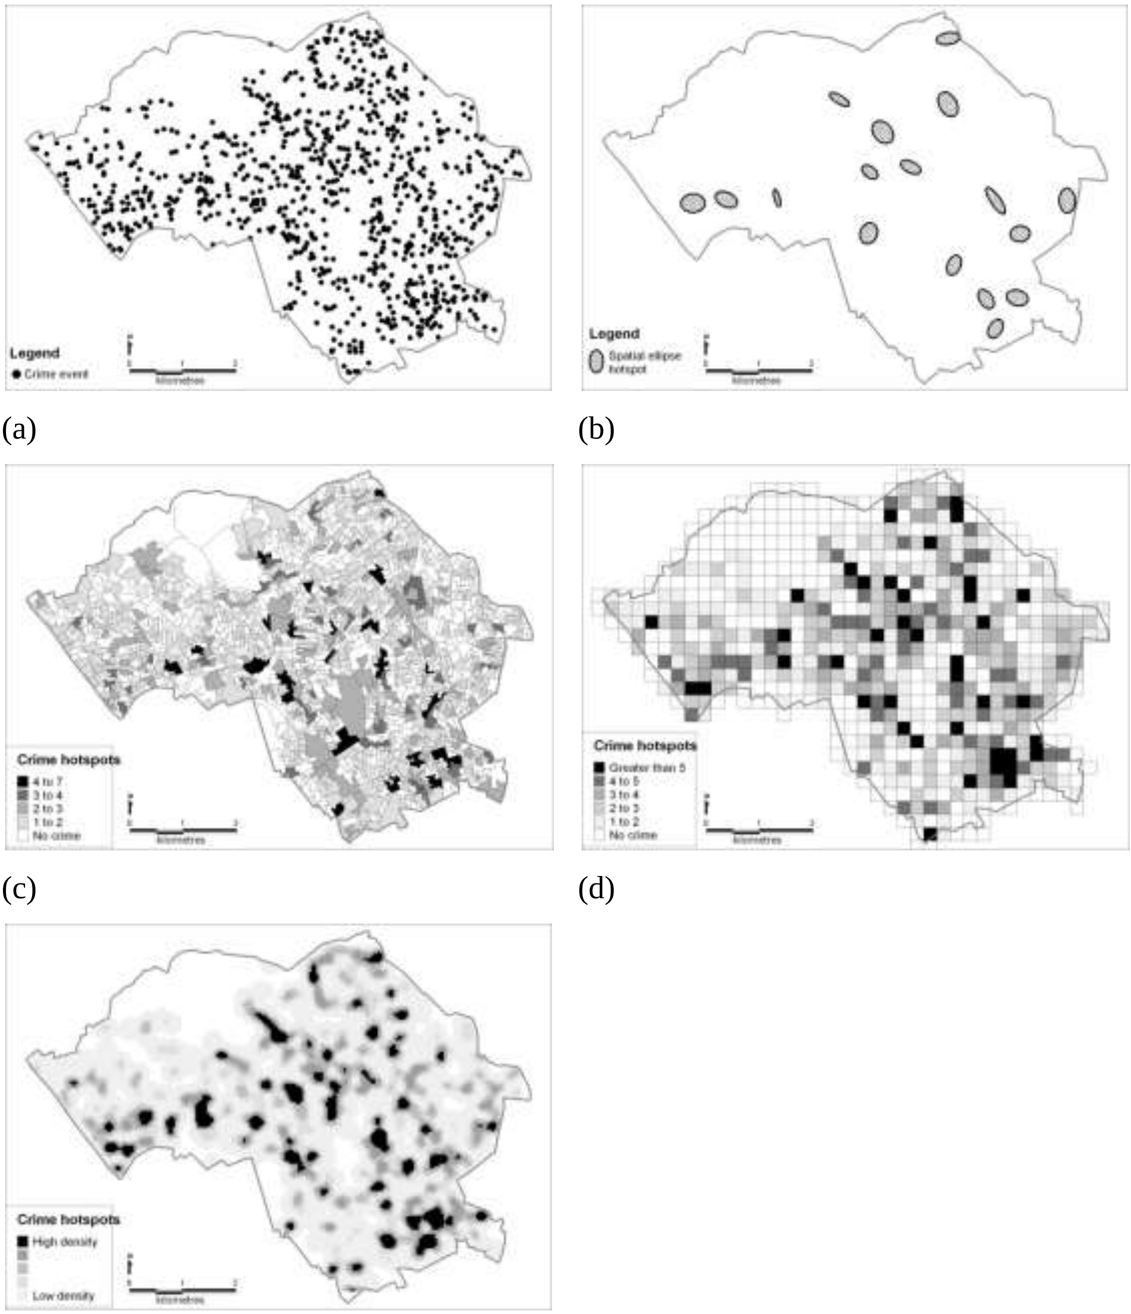

(a) (b) (c) (d) (e)

Figure 2.5. Common hotspot mapping techniques. (a) Point mapping, (b) standard deviational spatial ellipses, (c) thematic mapping of administrative geographic areas, (d) grid thematic mapping, and (e) kernel density estimation

Due to the varying size and shape of most geographical boundaries, thematic shading can beguile the map reader into identifying the existence of the highest crime concentrations (Eck et al., 2005). Hence, thematic mapping of geographic boundary areas can fail to reveal patterns across and within the geographical division of the boundary areas (Chainey and Ratcliffe, 2005). Also, as with all mapping reliant on defined geographical boundaries, the problem of the Modifiable Areal Unit Problem (MAUP - Openshaw, 1984) produces further complications. This is where changes in the boundaries themselves can directly affect the patterns shown on the map.

Thematic mapping of boundary areas continues to see widespread application, such as comparing the different volumes of unique and repeat burglaries across a study area's census zones (Ratcliffe and McCullagh, 2001), comparing vehicle theft in relation to land use in Overland Park, Kansas (Harries, 1999), and analysing and presenting crime patterns across partnership administrative zones (Chainey, 2001; Home Office, 2005). An example of a thematic map, generated for census output areas, is shown in Figure 2.5c. To date, the thematic mapping of administrative boundary areas technique for identifying hotspots has not undergone any analysis that determines how accurate it is for predicting spatial patterns of crime.

III. Grid thematic mapping

In order to address the problems associated with different sizes and shapes of geographical regions, uniform grids (or quadrats) can be drawn in a GIS as a layer over the study area and thematically shaded. Therefore, all areas used for thematic shading are of consistent dimensions and are comparable, assisting the quick and easy identification of hotspots. Bowers et al. (2001) used this method as a component of a GIS-based database application set up to identify vulnerable residences where target hardening was then implemented. LeBeau (2001) also found this technique useful when mapping the volume of emergency calls and violent offences per square mile in North Carolina.

The grid thematic mapping approach does have some limitations. The use of grids still restricts how the hotspots can be displayed. Spatial detail within and across each quadrat is correspondingly lost because the crime events have to conform to one specific quadrat, and this can then lead to inaccurate interpretation by the map user. Additionally, many comments have been made about the *blocky* appearance of this technique (Chainey and

Ratcliffe, 2005; Eck et al., 2005; Home Office, 2001), which is affected by grid cell size. The solution, reducing the size of each cell, can destroy the visual interpretation of the thematic map by making it look *specklely* and can fail to provide any useful information about where crime clusters (Chainey and Ratcliffe, 2005). Finally, grid thematic mapping suffers from the same MAUP problems outlined above (Bailey and Gatrell, 1995). An example of a grid thematic map is shown in Figure 2.5d. To date, the thematic mapping of grid cells technique for identifying hotspots has not undergone any analysis that determines how accurate it is for predicting spatial patterns of crime.

IV. Kernel density estimation

Kernel density estimation (KDE) is regarded as the most suitable of the common hotspot mapping techniques for visualising crime data (Chainey and Ratcliffe, 2005; Chainey et al., 2002; Eck et al., 2005; McGuire and Williamson, 1999; Williamson et al., 1999; Williamson et al., 2001;). KDE is an increasingly popular method due to its growing availability in many GIS software[3], its perceived good accuracy of hotspot identification, and the aesthetic look of the resulting map in comparison to other techniques (Chainey and Ratcliffe, 2005; Eck et al., 2005; Jefferis, 1999). Point data (i.e., recorded crime offences) are aggregated within a user-specified search radius and a continuous surface that represents the density of the points across the study area is calculated. A smooth surface map is produced, showing the variation of crime density across the study area, with no need to conform to geometric shapes such as ellipses or other geographic units. There is flexibility when setting parameters such as the grid cell size and bandwidth size (search radius); however despite many useful recommendations (see Chainey and Ratcliffe, 2005; Eck et al., 2005, Ratcliffe and McCullagh, 1999), there is no universal doctrine on how to set these parameters and in what circumstances. Examples of the use of KDE are now widespread, with popular crime mapping texts showing many examples of its use (see Chainey and Ratcliffe, 2005; Eck et al., 2005; Goldsmith, et al; 2000, Harries, 1999).

KDE is not without its faults. Eck et al. (2005) highlight that the choice of thematic range to use still presents itself as a problem as agencies fail to question the validity or statistical

[3] Kernel density estimation is available in MapInfo via the add-ons Hotspot Detective and Hotspot Gridder, and in ArcGIS via the ESRI extensions Crime Analyst, Spatial Analysis and the free extension Hawth's Tools. KDE is also available in free spatial analysis tools such as GeoDa.

robustness of the produced map, instead being caught in its *visual lure*. This largely affects how hotspots are identified and increases the variation of maps fashioned from the same data. There are also concerns that small amounts of data can misinform the map reader (Home Office, 2001). Nevertheless, the KDE technique has been in vogue for several years, not only because it is visually appealing but also because it can be applied by using a series of easy to follow systematic guiding principles (Chainey and Ratcliffe, 2005; Chainey et al., 2002; Eck et al., 2005; Williamson et al., 1999). An example of a KDE crime hotspot map is shown in Figure 2.5e.

Similar to spatial ellipses, thematic mapping of administrative areas and the thematic mapping of grid cells, the KDE technique for mapping hotspots has not been subject to any detailed research that measures its performance in predicting spatial patterns of crime. Johnson et al. (2008b) did, however, compare the spatial prediction performance of KDE to a prospective mapping approach (the prospective mapping approach is discussed in section 2.5.4) but did not consider how the cell size or bandwidth size, and the retrospective period of crime data used to create KDE mapping outputs influenced their results.

2.3.4. Initial conclusions and gaps in the existing research: the accuracy of hotspot mapping techniques for predicting spatial patterns of crime

Theoretical principles indicate that crime concentrates at certain places, with the review in this chapter of spatial analytical techniques for identifying these hotspots of crime illustrating examples of these non-random spatial patterns. The application of hotspot mapping has been used to support a range of practices in policing and crime prevention – from supporting the targeting of police patrols, to being a stage in the analytical process of developing intelligence products that can help determine a range of responses for these areas. These examples illustrate that hotspot analysis forms a basic method of crime prediction – using data on retrospective crime events to determine where crime may take place in the future. However, to date, very little research has been conducted that measures how accurate these hotspot mapping techniques are for predicting spatial patterns of crime. In addition, each technique requires the analyst to enter certain input parameters that have an influence on the mapping output that is generated. The current research aims to fill this gap by completing a comprehensive examination of these common hotspot mapping techniques to determine how accurate they are in predicting spatial patterns of crime.

2.4. Spatial significance mapping

Maps generated using KDE and the other common hotspot mapping methods are useful for showing where crime concentrates but may fail to unambiguously determine what is *hot* (in hotspot analysis terms) from what is *not hot.* That is, they may fail in separating the significant spatial concentrations of crime from those spatial distributions of less interest or from random variation. This weakness in failing to specifically determine the areas that are *hot* on a hotspot map has led in recent years to practitioner interest in the use of techniques that provide a more robust statistical assessment of the spatial distribution of crime by determining if the patterns observed are statistically significant. These spatial analysis techniques include local indicators of spatial association, often referred to as LISA statistics. The process of testing whether the spatial concentration of crime is unusual (i.e., is statistically significant) offers promise for building upon the current commonly used hotspot analysis techniques by statistically determining the hotspots and improving the prediction of where crime is likely to occur.

LISA statistics include Local Moran's I, Local Geary's C and the Gi* statistic. These techniques identify the association between a single point and its neighbours up to a specified distance from the point (Anselin, 1995). These statistics, therefore, provide an indication of the extent to which the value of an observation is similar or different to its neighbouring observations. Each technique requires data to be aggregated to some form of geographic unit (e.g., census output areas or grids) and a spatial neighbourhood for each of these geographic units to be defined. This could be units that are adjacent, units within a specified radius, or all other geographic units that are negatively weighted by the distance from the observation zone. Each of these statistics can be used to assess the local association in crime data, but they do so in different ways.

The following section reviews the use of Local Moran's I, Local Geary's C and the Gi* statistic to crime data and examines if each have been subject to research that identifies how they could improve on the common hotspot analysis techniques for predicting spatial patterns of crime.

2.4.1. Local Moran's I and Local Geary's C

Local Moran's I and Local Geary's C are variants of their global spatial autocorrelation equivalents. The global Moran's I statistic compares the similarity in values between each location and its near neighbours (Eck et al., 2005; Levine, 2010; O'Sullivan and Unwin, 2003). Where values in a geographic unit are high and are surrounded by other geographic units with similarly high values, positive spatial autocorrelation exists. If, however, areas with low values are surrounded by high crime areas and high crime areas are surrounded by spatial units with little crime, the series would display negative spatial autocorrelation. Geary's C statistic is a measure of the deviations in intensity values of each point with one another. Similar to Moran's I, the results from a Geary's C test determine if there is evidence of positive spatial autocorrelation or negative spatial autocorrelation. The local variants of Moran's I and Geary's C effectively apply each measure to each local spatial unit in comparison to its neighbours. The local variants can be used to determine for each observation the extent to which there is spatial association (i.e., clustering of similar values around that observation). Both Local Moran's I and Local Geary's C provide a means for testing the significance of the observed concentration. Although an exact test of significance is not possible, high positive or high negative standardized scores of I and C are taken as indicators of similarity or dissimilarity (Levine, 2010).

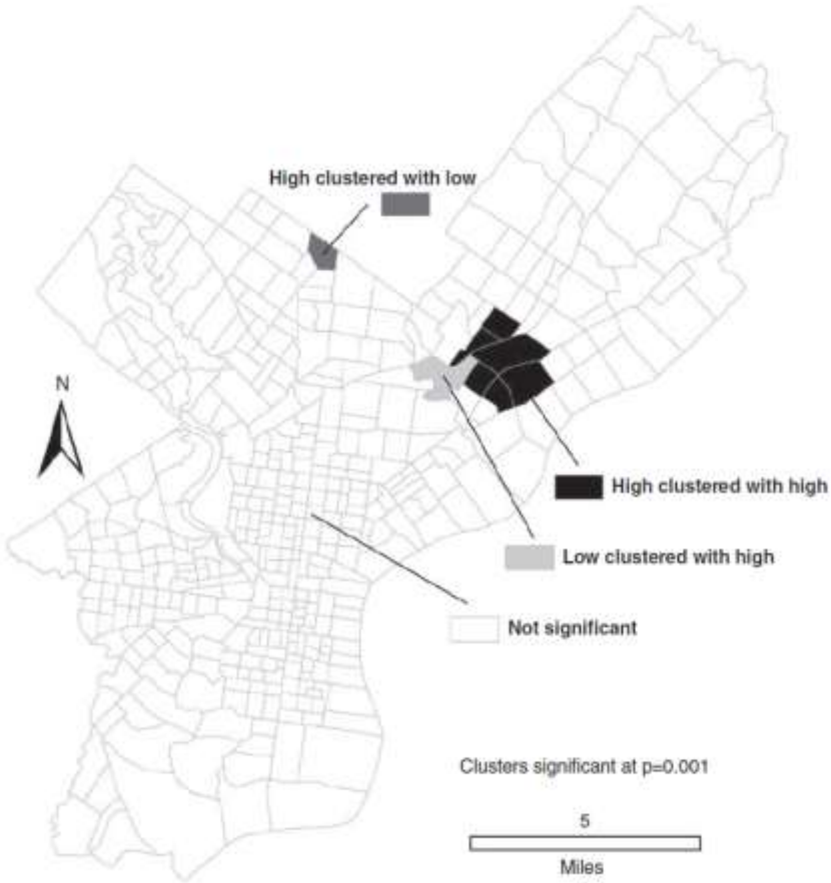

Figure 2.6. Local Moran's I significant spatial clusters in robbery frequency change from 2005 to 2006 by police sector in Philadelphia (Source: Ratcliffe, 2010)

The use of Local Moran's I for crime analysis has been demonstrated by Ratcliffe (2010) to analyse the spatial dispersion of robbery in Philadelphia (see Figure 2.6). This application showed those areas where the emerging pattern of robberies was similar between neighbourhoods, both in terms of where robberies had increased (high areas surrounded by high areas) and where they had reduced (low areas surrounded by low areas). Apart from this example, Local Moran's I and Local Geary's C have not been applied to identifying hotspots of crime. In addition, although Ratcliffe's example showed where robberies in Philadelphia had recently emerged, it did not go further by investigating whether these areas were where robberies were predicted to occur in the future. This suggests there would be value in examining further whether Local Moran's I and Local Geary's C are accurate in identifying hotspots of crime and predicting where crime may occur. The current research examines the technical application of each test in a subsequent chapter of this thesis.

2.4.2. Gi* statistic

The Gi* (pronounced G-i-star) statistic is a measure that compares local averages to global averages in order to identify those areas that are significantly different in comparison to what is generally observed across the whole study area (Ord and Getis, 1995). Using data that have been aggregated into geographic units (e.g., grid cells or census areas), the sum of values for the unit of interest and its neighbours within a user defined radius is compared to all the values in the geographic units for the entire study area. If local spatial association exists, it will be exhibited by a spatial clustering of high or low values. When there is a clustering of high values, the Gi* values will be positive. Low values will yield a negative Gi* value (Anselin, 1995; Chainey and Ratcliffe, 2005; Getis and Ord, 1996; Ord and Getis, 1995). The output generated from the Gi* statistic comprises standardised Z scores which can be used to determine whether the Gi* value for each cell of interest is statistically significant.

The use of the Gi* statistic was demonstrated by Eck et al. (2005) in their study identifying hotspots of robbery in the London Borough of Hackney (see Figure 2.7). This example showed that the Gi* mapping output identified very similar areas to a KDE hotspot analysis of the same robbery data, but showed potential in the use of the Gi* statistic for defining hotspot areas in spatial significance terms. Eck et al.'s (2005) research on the Gi* statistic did not go further by investigating whether the robbery hotspots that were identified were where robberies also occurred in the future. This suggests there would be value in examining further whether the Gi* statistic is accurate in identifying hotspots of crime and predicting where crime may occur. The current research examines the technical application of the Gi* statistic in subsequent sections of this thesis.

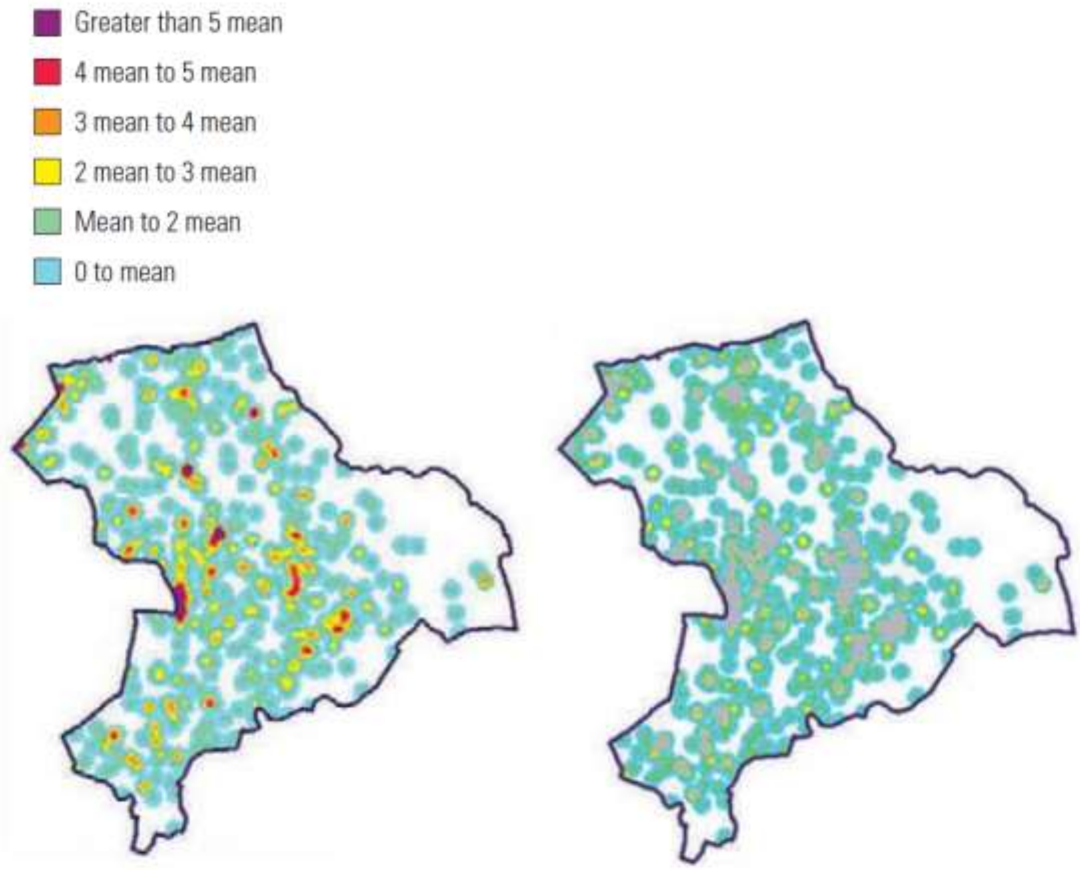

Figure 2.7. A comparison between KDE hotspots (shown on the left) and Gi* hotspots (shown in grey overlaid on the KDE hotspot output on the right) of robbery in the London Borough of Hackney (Source: Eck et al., 2005)

2.4.3. Initial conclusions and gaps in the existing research: the use of spatial significance mapping for predicting spatial patterns of crime

Local Moran's I, Local Geary's C and the Gi* statistic offer the potential to identify more robustly (in a statistical sense) where hotspots occur, to define the spatial extent of *hot* areas, and in turn to predict where crime may occur in the future. That is, Local Moran's I, Local Geary's C and the Gi* statistic may be better than many of the commonly used hotspot mapping methods used in policing and public safety for determining where to target resources. The current research examines Local Moran's I, Local Geary's C and the Gi* statistic to determine whether they offer improvements to common hotspot mapping techniques for identifying hotspots and predicting where crime may occur in the future.

2.5. New developments in risk and predictive mapping

In recent years, a number of new predictive crime mapping techniques (also referred to as predictive policing techniques) have been introduced. These include software products developed by IBM (IBM, 2010a) and PredPol (PredPol, 2013a), prospective mapping developed by Johnson and Bowers at University College London (Bowers et al., 2004; Johnson and Bowers, 2004a), and Risk Terrain Modelling (RTM) developed by researchers at Rutgers University in the USA (Kennedy et al., 2011). The typical marketing message that supports these predictive mapping techniques is that they identify where crime is most likely to occur in the future, and criticise hotspot analysis as only looking back at the past. However, these new techniques have been subject to very little critical assessment that examines how accurate they are in predicting spatial patterns of crime. The current research does not aim to assess metrically the prediction performance of all the new predictive mapping techniques, but by conducting an analysis of the prediction performance of hotspot mapping techniques it will establish a benchmark against which new predictive mapping techniques can be compared. In the current research, only a short review of RTM is included in the section that follows because its application did not emerge until the latter stages of this PhD and, to date, there is little published research available on how it can be practically applied to support spatial crime prediction. A detailed metric examination of the spatial prediction accuracy of RTM is though a recommended area for future research.

A short review of the main predictive software solutions follow to help illustrate and critique the foundations on which they are based. The use of repeat and near repeat patterns of crime for predicting where crime is likely to occur receives more detailed treatment. This detailed treatment is because the analysis of repeats and near repeats does not require expensive software solutions, theoretical developments explain the patterns of repeats and near repeats, and patterns of repeats and near repeats have been shown to be good predictors of crime (Bowers et al., 2004; Chainey, 2012b; Haberman and Ratcliffe, 2012; Johnson et al., 2008b).

2.5.1. IBM predictive analytics

IBM's predictive analytics software analyses multiple data sources to monitor, measure and predict when and where crime is most likely to occur. The software is based on combining IBM's SPSS Statistics software with ESRI ArcGIS. It primarily uses recorded crime data and calls for service data, but can also include any other recorded information that the user has the desire to include in the model. The first stage in adopting this IBM solution is for data to be passed through the IBM SPSS Modeler software that uses 'sophisticated statistical, data exploration and machine-learning techniques … to help agencies uncover hidden patterns and trends' (IBM, 2010a: 1) and determine the best model that fits the data. The model is then used with IBM SPSS Statistics and ArcGIS to produce 'crime forecast maps … that show the likely crime rate in different areas' (IBM, 2010a: 3). IBM claim their software has helped reduce crime in Memphis, Tennessee, by more than 30%, and by 35% in Lancaster, Los Angeles. The Lancaster IBM application cost $46,000 in the first year, $6,000 for each subsequent year, plus an additional $24,000 in training and consultancy (IBM, 2010b). Other than IBMs promotional material about their software and its apparent impact, little else has been published on the data it uses, the modelling techniques it applies, and the outputs it generates. No independent evaluation of the software's prediction accuracy has been conducted, nor are there any results that illustrate its prediction accuracy in comparison to other mapping techniques.

2.5.2. PredPol

Perhaps the most aggressively marketed predictive mapping tool is PredPol. The PredPol tool was developed over the course of six years by a team of mathematicians and social scientists at the University of California, Los Angeles, Santa Clara University, and University of California, Irvine in collaboration with crime analysts and police officers at the Los Angeles and Santa Cruz Police Departments. The principle aim of PredPol is to 'place officers at the right time and location to give them the best chance of preventing crime' (PredPol, 2013a). Based on models for predicting aftershocks from earthquakes, the predictions the PredPol software generates use a combination of historical patterns of crime with more recent information on crime incidents. The inclusion of recent incidents in the PredPol prediction procedure is based on the empirical finding that earthquake aftershocks often occur close in space and time to previous earthquake events (PredPol, 2013b). This concept is similar to that of repeat and near repeat victimisation that is discussed in the following section (section 2.5.3). 'In contrast to technology that simply maps past crime data, PredPol applies advanced mathematics and adaptive computer learning. It has resulted in predictions twice as accurate as those made through existing best practices by building on the knowledge and experience that already exists.' (PredPol, 2013b).

PredPol predicts where crimes are likely to occur in *prediction boxes* that are 500 feet by 500 feet in size. Officers are briefed before they go out on patrol about the highest-probability hotspots for that day and are told to devote extra attention to those areas. The cost for PredPol is not clear, however, the author has learned that for a medium sized UK police force the cost of the PredPol solution would be approximately £200,000.

Since implementing PredPol in Santa Clara, there has been a decrease in crime, including a reduction in burglaries of 27%. In the Los Angeles' Foothill Division, crime reduced by 13% in the four months following the rollout of PredPol compared to an increase of 0.4% in the rest of the city where PredPol had not been implemented (PredPol, 2013c). These assessments of changes in crime have been conducted by the police departments and PredPol developers themselves rather than being subject to independent scrutiny. Similarly, the claim that PredPol predictions are twice as accurate as other forms of hotspot analysis has been subject to some criticism by the International Association of Crime Analysts (IACA, 2013), which has stated these claims as being over exaggerated, non-representative of the hotspot mapping output generated by police analysts, and misleading.

2.5.3. Risk Terrain Modelling

Another developing area in predicting where crime is likely to happen in the future has been with Risk Terrain Modelling (RTM) (Kennedy et al., 2011). This approach uses a number of variables to generate a composite risk map as an aid to hotspot analysis. However, this method leaves it to researchers to choose variables they believe are suitable for determining crime risk, rather than being directed by a statistical procedure. In addition, the RTM approach assumes that each variable selected by the researcher has equal weight across space, rather than recognising that the explanatory (and causal) influence that each variable has is likely to vary across space. To date, despite its growing popularity, RTM has not been subject to any metric examination of how it performs in predicting where crime is likely to occur in comparison to hotspot analysis techniques.

2.5.4. Predicting crime using the patterning of repeat and near repeat victimisation

Perhaps the most utilised concept that has aimed to improve the spatial predictions of crime is repeat victimisation and near repeat victimisation. Repeat victimisation is the concept of a person or building being subject to victimisation a number of times. Research into repeat victimisation has shown that, overall, risk doubles following a victimisation, and that repeats occur swiftly after the initial incident (Farrell and Pease, 1993; Polivi et al., 1991). Interviews with offenders have also supported evident patterns of repeat victimisation, with Ericsson (1995), for example, finding that 76% of offenders interviewed returned to a number of houses to burgle them 2-5 times.

Near repeat victimisation is the observed finding that targets near to a recent incident are also at a heightened risk of being victimised. While this concept was initially developed from analysis of residential burglary (Johnson and Bowers, 2004a), near repeat patterns have also been found for vehicle crime, violent crime, pedal cycle theft and shootings (Haberman and Ratcliffe, 2012). The level of risk to neighbouring targets is lower than the risk of victimisation to the recent victimised target, and decays with distance from this original target. Similar to repeat victimisation, this heightened risk to neighbouring targets decays over time.

The reasons why repeats and near repeats occur cannot be simply explained using just the spatial theoretical principles for crime that have been previously described in this thesis. Subsequent to the key spatial theories for crime has been the development of the boost account, flag account and optimal foraging theories for explaining why offenders commit crime repeatedly at the same locations, and at locations nearby. The boost account refers to an offender deciding to return to the same location or nearby locations, boosted by the success of previous crime commission. Optimal foraging theory provides an explanation for why the boost account occurs by suggesting that offenders commit offences in spates, seeking to benefit from the rich supply of targets in an area, but moving on once the supply of targets have been exhausted or there is risk of capture. The flag account suggests there is some enduring quality about the target that highlights its high level of vulnerability to would-be offenders. Each theory will now be explained in more detail.

Repeats and near repeats are primarily believed to relate to the boost account theory (Bowers and Johnson, 2004; Pease, 1998). This theory states that future victimisation is boosted by the initial incident. That is, the offender got away with it, so why not do it again? This is best explained by considering residential burglary where the offender is *boosted* in his decision-making to target the same property because he now knows how to get in, the layout of the property and what he left behind previously. This boost theory also applies to neighbouring properties in that the chances of a neighbouring property being burgled are boosted by the initial burglary due to the following reasons:

- the neighbouring properties are more likely to be similar in design (in comparison to properties further away),
- the offender is already familiar with the area (having chosen to burgle here previously),
- the means of breaking in and the layout of the property are likely to be similar (so he knows his way around), and
- the neighbours are likely to have possessions worth stealing, similar to those stolen in the initial burglary.

Optimal foraging theory provides a means of explaining why the boost account occurs (Johnson and Bowers, 2004b; Johnson et al., 2009). This approach likens offenders to foraging animals. As a forager, an animal makes a trade-off between the energy value of the food that is immediately available and the effort that will be expended in reaching a better food source. The better food has to be good enough to offset the energy required to travel and attain it. The quality of the food in over-grazed areas diminishes until it re-grows. This is akin to a repeatedly burgled property and the burglary of properties nearby, where the value of the items taken from these properties declines until these items have been replaced. Once an area has been grazed out (i.e., skimmed of the best theft opportunities), the forager moves on.

The flag account theory provides an additional explanation for why repeats and near repeats occur (Pease, 1998). The flag account suggests there is some enduring characteristic about the property that *flags* it as being vulnerable. For example, in terms of burglary, it could be a property at the end of a terrace, which has an alley running along the back, and appears to have poor door and window security – all of which are signals that identify an easy target to a would be offender. In practice, it is likely that both boost and flag theories are at play in explaining why repeats and near repeats occur. For instance, the flag characteristics of a property may initially attract an offender because it is seen as an easier target, with the risk of future burglary being boosted following an initial incident. A distinction between the boost and flag accounts is though concerned with who is likely to commit the repeats or near repeats. In terms of the boost account, any repeats or near repeats will be committed by the offender who committed the initial offence (Bowers and Johnson, 2004). In terms of the flag account, any repeats or near repeats are likely to be committed by other offenders who have each identified the vulnerable characteristics of the target, or a returning offender, but after the initial boost period has diminished (i.e., an offender returns to commit another offence at the same house some months or years after an initial incident) (Bowers and Johnson, 2004).

The boost and flag principles, and foraging behaviour are consistent with the findings from interviews conducted with offenders. For example, as one offender reported in Ashton et al. (1998: 276), 'once you've been in a place it's easier to burgle because you are familiar with the layout, and you can get out much quicker'. The empirical findings from research into repeats and near repeats, and the theoretical principles that underpin these patterns, has led some commentators to suggest that recent incidents, particularly those where other incidents have previously occurred, provide the best variable for predicting where crime is likely to take place in the future (Bowers et al., 2004).

These principles of repeats and near repeats have been programmed into both software and operational policing activities. For example, the Vigilance Modeller[4] is based on an algorithm that incorporates the patterning principles of repeats and near repeats to generate mapping output that identifies places at the highest risk of crime over the next seven days. This knowledge of repeat and near repeat patterns has also begun to be embedded into how police conduct patrols. For example, in Trafford, Greater Manchester Police conducted an analysis of burglary to domestic properties and identified patterns of repeats and near repeats. From this, a series of operational tactics were designed to respond with crime prevention advice to properties that had been burgled within the last 24 hours. In addition, visits were also made to neighbouring properties to raise awareness about the recent burglary, reassure residents, and ask them to conduct their own practical crime prevention to minimise their risk of burglary and to immediately report any suspicious activity. Targeting uniformed police patrols to areas that had recently been burgled was also considered to act as a deterrent to future offending to any returning *forager*. The result was a 42% reduction in burglaries in those areas that were targeted (Chainey, 2012b; Fielding and Jones, 2012).

[4] The Vigilance Modeller is an online tool for generating risk areas, developed by the Home Office in partnership with Lincolnshire Community Safety Partnership, Astun Technology and UCL.

The results from analysing patterns of repeat and near repeat victimisation and designing tactics to counter these in Greater Manchester show promise that this approach can predict spatial patterns of crime. A set of experiments that has further illustrated the potential predictive power of repeat and near repeat patterns was conducted by Johnson et al. (2008b) that used a prospective mapping tool (that incorporates the same mathematical algorithm as the Vigilance Modeller) to predict where crime was likely to occur in the future. These experiments compared this prospective mapping approach to KDE and found the former technique generated the better predictions. However, these experiments were only conducted using residential burglary data, did not examine whether the KDE parameter settings (i.e., cell size and bandwidth size) influenced the outcome of the results, and only made an assessment based on the prediction of crime in the week that followed.

While there is sound empirical evidence for the patterns of repeats and near repeats, very little research has been conducted that metrically assesses whether the spatial patterns it predicts are relevant for all temporal periods of the future. For example, Johnson and Bowers (2004b) showed that burglary patterns tend to be *slippery*, in that while recent events offer an effective means of predicting where crimes are likely in the immediate future (i.e., over the next few days), the spatial distribution of events beyond this immediate future tends not to match as accurately as those from the recent past. However, following from the many other longitudinal studies of the spatial distribution of crime, these studies have shown the persistent longevity of crime at certain places (Brantingham and Brantingham 1981; Groff et al., 2008; Shaw and McKay, 1942; Weisburd et al., 2004).

In practice, the findings relating to both the *slippery* nature of short-term spatial patterns of crime, and the stability of crime concentration over longer periods suggest that a predictive mapping approach based on the principles of repeats and near repeats can support the daily tasking of police resources, but can overlook some of the necessary strategic focus that crime prevention targeted on persistently problematic areas also requires. From this it is argued that one single method for predicting spatial patterns of crime is unlikely to be accurate for predicting where crime will occur for different periods of the future. While recent incidents may be accurate for the purpose of predicting where other individual incidents of crime are likely to occur in the immediate future (i.e., the next few days), the locations where multiple incidents of crime are likely to concentrate beyond this immediate future could be in different areas.

The accurate prediction of spatial patterns of crime may, therefore, require a range of spatial analysis methods, with each suited for different prediction time periods. If a range of spatial analysis methods are required for predicting spatial patterns of crime, each must also be underpinned with appropriate theory that explains why this spatial patterning is likely to occur. In this sense, boost and foraging principles appear to offer the theoretical basis for explaining where incidents are likely to occur in the immediate future and the observed *slippery* nature to some of this spatial patterning as the forager moves on to another area. The flag account theory then provides some of the theoretical justification to explain why hotspots of crime tend to be stable (i.e., there is some intrinsic desirability about the targets in an area that explains why they persistently experience crime). However, similar to how foraging explains the boost account and the *slippery* nature to spatial patterns of crime over short time periods, there appears to be a gap in the existing environmental criminology literature that provides a supporting explanation to the flag account and the stability of crime concentration over longer periods. As introduced in section 2.2.3, crime is likely to concentrate in areas where there are *favourable conditions* for crime to occur. While the theoretical concepts of crime generators and crime attractors (Brantingham and Brantingham, 1995) provide some explanation for crime concentration, the inclusion of a new theoretical principle relating to favourable conditions may be of value for helping to explain the flag account and the stability of crime concentration over longer periods. A new theoretical principle relating to favourable conditions for crime to occur is considered further in the discussion chapter (chapter 12). In the first instance a detailed examination of crime patterns is required to identify if existing theory is sufficient for explaining the spatial patterns of crime and explaining where crime is predicted to occur in the future.

2.5.5. Initial conclusions and gaps in the existing research: developments in risk and predictive mapping

Since 2009 a number of new techniques have emerged that are designed more specifically for predicting where crime may occur. These include commercial software solutions such as PredPol and IBM’s Predictive Analytics, and tools developed from academic research such as prospective mapping and Risk Terrain Modelling. All of these methods aim to advance hotspot analysis, yet to date, none of them have been tested with much rigour to identify if and how they improve on the predictive ability of hotspot analysis. The approach the commercial software solutions take is that more variables than just retrospective crime data should be used to predict crime, and that these data need to undergo sophisticated data mining and computer algorithmic processes to generate predictions that will be more reliable and accurate. To apply this higher level of processing and sophistication comes at a fairly high financial cost. In addition, (in some cases) there appears to be a lack of theoretical grounding on which these commercial software solutions are based, with the companies producing the software failing to specifically describe the environmental criminology theoretical principles on which they are based.

In comparison, the use of repeat and near repeat patterns provide a sound theoretical basis for determining where an offender may offend next, and hence quite logically can support a crime prediction approach. However, the periods into the future for which predictions are made using repeat and near repeat patterning principles needs further investigation to determine whether these predictions are just as accurate for the immediate future as they are for predicting where crime is likely to concentrate in the next few weeks and beyond. To date, no suggestion has been made by other researchers or commercial companies about the importance of using different spatial analysis techniques for predicting different time periods of the future. This current research has introduced this potential issue by suggesting there may be value in using different spatial analysis techniques for predicting spatial patterns of crime for different periods of the future, but with the requisite that each is also supported with theory that underpins why these crime patterns are likely to occur. The current research aims to help fill this gap by comparing the accuracy of the spatial predictions made using prospective mapping for different periods of the future. These prospective mapping predictions will be compared with those generated using the benchmark analysis results from the examination of hotspot mapping techniques.

Risk Terrain Modelling is an interesting addition to the recent developments in predictive analysis techniques, again developed in the notion that variables other than retrospective patterns of crime should be used for determining where crime is likely to occur in the future. However, to date there has been very little research that shows that the addition of these extra variables results in improving the predictions over and above the predictions that can be made from using retrospective crime data alone. In addition, the selection of additional non-crime variables for RTM is typically not subject to an analysis of the influence they have on crime levels, and fails to recognise the influence that each variable has on crime is likely to vary spatially.

2.6. Exploring geographical relationships as a means to improve hotspot analysis and crime prediction

While hotspot analysis is useful for identifying where high levels of crime concentrate, it offers little for helping to explain *why* crime concentrates at these locations. Most practitioners recognise that hotspot analysis is one of the techniques that can be used at the start of a process for examining why a crime problem exists. This diagnostic approach is very much a key principle to the problem-oriented policing method (Goldstein, 1979, 1990) where the focus is to deal with the underlying causes of problems. Understanding these causes requires knowledge of appropriate theory, accurate data and an iterative interchange between analysis and theory to help make judgements on why the crime problem may exist (Clarke and Goldstein, 2003; Townsley et al., 2011). However, to date, little use has been made in crime analysis of spatial regression techniques to statistically inform the reasons for spatial variations in crime. Spatial regression analysis is the process of identifying variables that explain the spatial variation in the phenomenon under study. It is argued that if spatial regression can be used to help identify the variables that explain why hotspots are present, these explanatory variables could then be used alongside or as a replacement for retrospective data on crime to predict where crime is likely to occur. The current research examines the use of spatial regression techniques for improving hotspot analysis. In particular, attention is focused on geographically weighted regression (GWR) because the GWR approach addresses many of the statistical challenges associated with spatial regression and because the software used is freely available to many practitioners. A starting point for explaining spatial regression is to review the use of Ordinary Least Squares regression as a means for exploring possible causal factors.

2.6.1. Regression analysis

Regression analysis involves testing the relationship between one phenomenon (e.g., a crime problem) and factors that are considered to cause the phenomenon. The result of the regression analysis determines which factors can be used to explain the phenomenon and the strength in the relationship, and identifies if there are other factors that were not included in the model that need to also be examined for their causal influence. All regression models seek to examine the relationship between a dependent variable and explanatory variables to determine the presence, direction and strength of any relationships. The most basic regression model is Ordinary Least Squares (OLS). OLS assumes that these relationships are consistent geographically. That is, the reasons that explain the presence of a phenomenon in one location are exactly the same as those in another. For example, if hotspots of burglary are present in two locations, an OLS regression would assume that the reasons for these hotspots are exactly the same, influenced equally by the factors that are causing burglary. In reality, this is unlikely to be the case.

In addition, OLS regression assumes there to be zero or negligible errors in the explanatory variable and that the error terms are independent of one another. When dealing with spatial data, assumption that error terms are independent of each other is often violated due to spatial autocorrelations in data; that is, nearby observations are similar to each other. This can lead to a biased estimation of the standard errors of the model's variables, and consequently result in misleading significance results (Anselin and Griffith, 1988; Fox et al., 2001).

2.6.2. Spatial regression analysis

Analysis techniques for spatial regression examine the spatial relationship between variables. If the spatial relationships between these variables are significant, the spatial autocorrelation between these variables needs to be taken into account. If not, this can lead to statistical problems with the spatial regression model, including unreliable significance results and bias in the estimation of the standard errors of the model's variables.

A number of approaches are used for capturing spatial effects in regression analysis. These include incorporating spatial effects through an error term in the regression analysis (spatial error models), or where the spatial effects are incorporated by including a spatially lagged variable as an additional predictor (spatial lag models). The spatial lag model adds to the standard OLS regression model equation by considering the impact of the dependent variable in neighbouring areas through the inclusion of a weighted mean value of the local dependent variable (Anselin, 1988a; Anselin, 1988b). By doing so, the spatial lag model implies the level of crime, for example, in one area is influenced by the level of crime in another (Chainey and Ratcliffe, 2005). A spatial error model provides an indication that the relationship between the dependent and explanatory variable is influenced by unmeasured variables (Messner and Anselin, 2004). For example, a spatial error model suggests that the amount of variance in the level of crime (the dependent variable) that is not predicted by the explanatory variables is due to spatially autocorrelated missing variables (Chainey and Ratcliffe, 2005).

Other spatial regression approaches that have been developed and tested to address the violation of the OLS assumption of the independence of observations when treated with spatial data include: spatial dependency models (Can, 1990 and Dubin, 1992; 1998); the generalised linear model (Littell et al., 1996; Schabenberger and Pierce, 2002) and its non-parametric extension, the generalised additive model (Guisan et al., 2002; Hastie and Tibshirani, 1990); artificial neural networks (Fischer, 1997; Shin and Goel, 2000); and Bayesian spatial regression (LaSage, 1997; Wheeler and Calder, 2007). Each of the models has been applied to a wide range of spatial applications including forestry (Austin and Meyers, 1996; Foody, 2003; Haiganoush et al., 1997; Wang et al., 2005; Zhang et al., 2004; Zhang and Shi, 2004; Zhang et al., 2009), house and land prices (Gao et al., 2006), child poverty (Voss et al., 2006), epidemiology (Mohebbi et al., 2011), demography (Loftin and Ward, 1983) and crime (Cahill and Mulligan, 2007; Wheeler and Waller, 2009; Waller et al., 2007).

Geographically Weighted Regression (Fotheringham et al., 2002) is an alternative approach to addressing the limitations of OLS regression when applied to spatial data. GWR is a local version of spatial regression that actively seeks to measure the variation in relationships between variables by performing localised regression equations all over the study area (Chainey and Ratcliffe, 2005). The current research focuses on the use of GWR rather than a broad range of spatial regression techniques because in a number of tests GWR has outperformed several of the others mentioned in its modelling performance (Waller et al., 2007; Wang et al., 2005; Zhang et al., 2005) and because it is more accessible in software to practitioners than many of the other spatial regression techniques[5]. However, it is recognised that GWR is not without its critics, including some commentators who question its use for determining inference (Griffith, 2008; Wheeler and Tiefelsdorf, 2005; Wheeler and Walker, 2009).

[5]GWR is available in ESRI's ArcGIS, and as a free download as a standalone application from the University of St Andrews

The process of applying GWR (as well as many other spatial regression models) involves the calculation of statistical diagnostic tests to help determine if the model is suitable. These tests include the calculation of measures to determine model performance, model significance, and model bias. For instance, these tests help to determine which explanatory variables contribute to the model, the statistical significance of the model, the degree of stationarity (i.e., whether the explanatory variables in the model have a consistent relationship with the dependent variable), if the residuals are normally distributed (i.e., exploring the presence of bias in the model by identifying whether the model performs well for low values but does not perform well for high values), and if the residuals are spatially clustered. If the residuals are spatially clustered, it suggests that a key explanatory variable is missing from the model.

GWR and other spatial regression models generate a number of outputs that can be mapped. These include the parameter estimates (i.e., the coefficient describing the relationship between the explanatory variable and the dependent variable for each observation on the map), t-values (indicating the statistical significance of the parameter estimates), and residuals (the level of error or unexplained component of the relationship between the explanatory variable and the dependent variable in the model). An example of GWR applied to crime data is shown in Figure 2.8. This figure identifies local spatial trends and processes between religious institutions, social disorganisation variables and their combined effect on the local homicide rates for the city of Philadelphia. The mapping of independent variable t-values revealed trends across the city that pointed to the presence of non-stationary spatial processes, something the global tests could not detect or depict. Global tests indicated that the density of religious institutions did not have a significant effect on homicide across the entire city. However, a spatial analysis that was sensitive to the possibility that some variables were more influential in various parts of the city suggested that religious institutions had a significant influence on homicide rates in neighbourhoods in North and North West Philadelphia (Chainey and Ratcliffe, 2005).

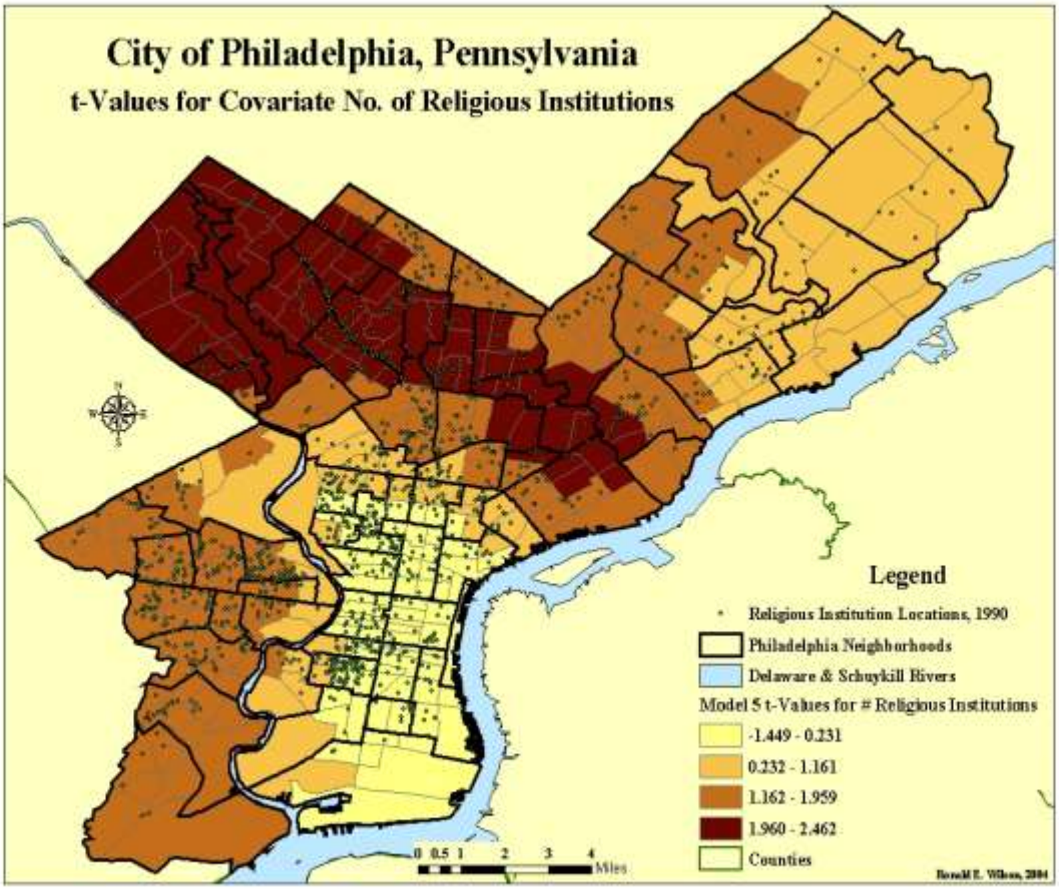

Figure 2.8. GWR generated t-values of the spatial distribution of religious institutions and their relationship with he incidence of homicide across the city of Philadelphia. This example suggested that the influence of religious institutions was stronger in some parts (the darker shaded areas) than others (Source: Chainey and Ratcliffe, 2005).

The additional challenge that crime data presents is that they typically follow a Poisson distribution (that is, many areas have a low count of crime, and few have a high count) and require different, specific treatment in regression models (most examples of spatial regression are based on using variables that follow a Gaussian distribution). While Poisson-based GWR models have been developed (with examples of its use including disease mapping – see Nakaya et al., 2005), to date very little research has been conducted that identifies the specific type of treatment, if any, that is required for the spatial regression analysis of crime data. The current research examines the application of GWR using crime data and determines if the GWR modelling process results in identifying variables that explain why hotspots exist. In turn, these results from an examination of GWR could identify how hotspot analysis could be improved by using variables that are statistically correlated with crime, alongside or in replacement of retrospective crime data for predicting where crime is likely to occur.

2.6.3. Initial conclusions and gaps in the existing research: identifying geographical relationships as a means to improve hotspot analysis and crime prediction

GWR provides the potential to explore the spatially varying relationships between levels of crime and what causes this variation in crime. In turn, this could help identify why hotspots exist and if there are differences in what is causing these high concentrations of crime. Armed with this information, potential then exists to associate these explanatory variables with the theoretical concepts on what causes crime to concentrate (e.g., crime attractors, routine activities and crime pattern theory) and construct these into a hotspot modelling process alongside retrospective data on crime, or in replacement of this retrospective information. This hotspot modelling approach that uses a combination of explanatory variables and crime data may lead to improvements in spatial crime prediction. In addition, GWR parameter estimates may potentially be used to inform how crime levels may change based on the change in one or a number of related variables. That is, offering strategic forecasts on how crime would be expected to change if a particular policy direction was chosen (e.g., the regeneration of what is currently a high crime housing estate).

2.7. Summary: from spatial theories of crime to analytical techniques that predict spatial patterns of crime

Examined in this chapter was the theoretical basis for explaining why spatial patterns of crime can be predicted. A number of spatial theories of crime exist, providing macro, meso and micro geographic explanations for crime. In the macro neighbourhood sense these include explaining how specific socio-economic conditions are likely to exist in certain places that lead to higher levels of crime and the factors that may attract many suitable or specific targets to concentrate in certain places. In a more meso and micro level sense, theoretical principles such as routine activities, rational choice, crime pattern theory, and least effort are valuable in helping to conceptualise the underlying criminal landscape and the spatial decision-making processes of offenders. The boost account, flag account and optimal foraging behavioural concepts add to these meso and micro explanations by explaining further the likely behaviour of offenders in their crime commission decision-making. Collectively, these theories explain why geographical patterns of crime are not random and why crime can be predicted to take place at certain locations. The review of these theories may have also identified a gap in existing environmental criminology theory that explains why certain places experience favourable conditions for high levels of victimisation. This PhD research investigates this potential gap further and identifies whether an expansion to existing environmental criminology theory is required that may help explain why certain places possess favourable conditions for victimisation. In the first instance, this requires a detailed examination of crime patterns to identify if existing theory is sufficient for explaining the spatial patterns of crime and why crime may occur in certain places in the future.

Hotspot mapping is a popular method used by police and public safety agencies to determine where to target resources. Hotspot mapping has also been used in many other study disciplines, in particular disease mapping. However, to date, the accuracy of hotspot mapping for predicting where crime is likely to occur has not been rigorously examined. A number of hotspot mapping techniques exist, including the thematic mapping of administrative areas and kernel density estimation, but the actual prediction accuracy of these techniques has yet to be fully evaluated. The research will examine these techniques by determining if they differ in the spatial predictions they produce, and if their mapping output is influenced by differences in their input parameters. The review of hotspot mapping techniques in this chapter has also identified the potential of spatial significance mapping for producing better spatial predictions of crime than those generated using commonly used hotspot analysis techniques, particularly in terms of defining hotspots in less ambiguous terms.

As recent developments in predictive policing have emerged, several new mapping techniques have been introduced. The review of these techniques in this chapter has resulted in questioning the IBM and PredPol software solutions for their theoretical foundations and the claims they are better than hotspot mapping. Risk Terrain Modelling is a new technique of interest but is subject for further research beyond the scope of this PhD. Prospective mapping, modelled on the theoretical principles of the boost account and optimal foraging theory is, though, of specific interest and will be subject to analysis that compares the output this technique generates in relation to hotspot mapping output. In particular, this will include examining if the spatial crime predictions produced using prospective mapping are accurate for multiple periods of the future – the immediate, near and more distant future.

Spatial regression analysis techniques offer a statistical means of identifying those variables that may explain why spatial patterns of crime vary. Geographically weighted regression is a particular technique of interest due its availability to practitioners and researchers, and its favourable reviews in comparison to other spatial regression techniques for addressing many of the challenges in analysing spatial relationships between data variables. The current research will examine GWR to determine if it can be used for statistically explaining why hotspots exist, and whether the results from this type of analysis can help inform predictions of crime.

To date, because no detailed examination of the spatial crime prediction performance of mapping techniques has been conducted, it is likely that sufficient measures for calculating the prediction differences between mapping outputs have not been developed. This research will review the literature further to identify methods that have either been used previously in spatial crime analysis studies for comparing differences between mapping output, or have been used in other fields of science that are appropriate for measuring the prediction performance of crime mapping output. If required, this research will develop new measures for calculating spatial crime prediction performance, or adapt those from other disciplines to make them suitable for these measurement requirements.

3. Objectives, research questions and hypotheses

3.1. Principal research question and objectives

The primary question this research aims to answer is: to what extent can hotspot mapping be used to effectively predict where crime is likely to occur? To answer this question requires identifying hotspot analysis techniques that accurately predict where crime occurs, and for the spatial predictions the techniques generate to be explained in clear theoretical terms. If the spatial predictions are accurate and the theoretical reasons for the predictions are clear, this in turn helps identify where policing and public safety activity should be targeted and the specific police tactics and crime prevention programmes to counter the predicted activity.

A number of spatial theories of crime exist, providing macro, meso and micro geographic explanations for crime. Collectively, these theories help to explain why geographical patterns of crime are not random and why crime can be predicted to take place at certain locations. Several common hotspot mapping techniques exist, but the actual prediction accuracy of these techniques has yet to be evaluated in detail. Each of these techniques is also influenced by several input parameters, such as the data they use and technical parameters that influence the calculations for determining where hotspots exist. The influence of these input parameters on the mapping outputs they generate has also yet to be evaluated. Potential exists in spatial significance mapping techniques for producing accurate predictions of crime, but these have also not been evaluated for their spatial crime prediction performance. A detailed examination of the commonly used hotspot mapping and spatial significance mapping techniques is therefore required in order to establish a benchmark analysis against which new predictive mapping techniques could be compared.

One of the new predictive mapping techniques is prospective mapping. Tests that have compared prospective mapping to common hotspot analysis techniques such as KDE have suggested that the prospective mapping approach is more accurate in making spatial predictions of crime. However, these tests only involved examining where crime was likely to occur in the next seven days, rather than examining where crime will occur for any longer periods into the future. In addition, previous research has shown that high levels of crime endure in places for long periods. This, therefore, suggests that prospective mapping (that uses very recent incidents) may be accurate in predicting where crimes occur in the immediate future (i.e., within the next few days), but may not be as accurate for predicting where crime is likely to persist for longer periods into the future. Instead, it is argued that where crime has previously persisted at high levels for some time (i.e., hotspots) is more likely to be where high levels of crime will occur for longer periods into the future. To determine if differences exist in the ability to predict spatial patterns of crime for different temporal periods requires a detailed examination of the temporal stability of hotspots and whether the prediction accuracy of prospective mapping is consistently better than hotspot analysis techniques for multiple periods of the future – the immediate, near and other time frames beyond.

Additional to the new prediction techniques is spatial regression analysis. Spatial regression analysis techniques, such as GWR, offer a statistical means of identifying those variables that may explain why spatial patterns of crime vary. However, techniques such as GWR have yet to be examined as to whether they can be used for statistically explaining why hotspots exist, and whether this type of analysis can help improve predictions of crime.

The gaps in the existing research into spatial crime prediction and the arguments that have been set forward for establishing the extent to which hotspot mapping can produce accurate spatial predictions of crime suggest the following key questions:

- Does the spatial crime prediction accuracy of common hotspot mapping techniques vary?
- Can these common hotspot mapping techniques be improved through attention to the influence that technical input parameters have on mapping output?
- Can the prediction accuracy of hotspot analysis be improved using statistical significance mapping techniques?
- Are hotspots of crime stable over time?
- What influence does the retrospective period of crime data that is used for producing mapping output have on spatial crime predictions? And are the accuracy of these predictions consistent for multiple periods of the future – the immediate, near and other time frames beyond.
- Does the use of techniques that explore spatially varying relationships help identify why hotspots exist?
- Can the analysis of spatially varying relationships be used for supporting long-term crime predictions by examining how a change in explanatory variables can influence a change in future crime levels?

The following section addresses these research questions in turn. Each research question is reframed as a hypothesis, with these hypotheses directing the empirical studies that then follow.

3.2. Testing for the presence of crime hotspots

Hypothesis 1: Hotspots can be identified using retrospective data for a short period of time rather than requiring retrospective data for longer periods of time

A preliminary stage to hotspot analysis involves determining if there is evidence that hotspots exist in the data under examination. This also requires an assessment of the volume of data that are required for statistical evidence of hotspots to be present. The first part of the research will therefore test the position in retrospective crime data when hotspot patterns first appear.

3.3. A metric assessment of the prediction performance of common hotspot analysis techniques

Hypothesis 2: Common hotspot mapping techniques (i.e., spatial ellipses, thematic mapping of administrative areas, thematic mapping of grids, and kernel density estimation) differ on how accurately they predict spatial patterns of crime

To date, a metric comparison of the prediction performance of the commonly used hotspot mapping techniques has not been completed. This stage of the research will involve a series of experiments that compares the spatial crime prediction performance of spatial ellipses, thematic mapping of geographic units, thematic mapping of grid cells and kernel density estimation.

3.4. A metric comparison of the influence that technical parameters used in hotspot analysis can have on spatial crime prediction performance

Hypothesis 3: The technical parameters used in hotspot analysis techniques have an influence on the techniques' spatial crime prediction performance.

As a result of testing hypothesis 2 and determining which of the common hotspot analysis techniques consistently performs well at spatial crime prediction, the technical parameters of this technique will be examined to see if they influence the technique's prediction performance. For example, if kernel density estimation is determined to be the best of the common hotspot mapping techniques, the influence of cell size and bandwidth size on the prediction performance of the mapping output will be tested.

3.5. Improving hotspot analysis using spatial significance mapping

Hypothesis 4: Spatial significance mapping methods provide an improved means of predicting where crime is likely to occur in comparison to common hotspot mapping techniques, and removes the ambiguity of defining areas that are *hot.*

Common hotspot analysis techniques such as grid thematic mapping and kernel density estimation require the analyst to determine from the spatially mapped values the crime intensity level that represents *hot*. This means the selection of hotspots is subjective and can lead to variation between analysts in the areas that are determined to be hotspots. Spatial significance mapping offers potential in removing this ambiguity in defining hotspots by using the principles of statistical significance testing. By determining, in a statistical sense, the areas that are hotspots, it is also possible that these identified areas offer a more accurate means of determining where crime is likely to occur in the future.

3.6. Examining the temporal stability of hotspots

Hypothesis 5: Areas that are identified as hotspots of crime are places where the concentration of crime has been endured consistently for at least one year, and where the concentration of crime is likely to continue to persist into the future.

Policing and crime prevention resources that are designed to tackle hotspots assume these hotspots have been endured for some time and are likely to continue to be the areas where high levels of crime will persist in the future. This is based on the findings from many longitudinal studies of crime and places. However, other research has suggested that hotspots tend to move around. The research on the *slippery* nature to hotspots was, though, based on findings referring to how crime changes on a daily and weekly basis, rather than how crime changes over longer periods (weeks and months). Hotspot analysis is more naturally suited to identifying areas where the concentration of crime is based on a longer temporal retrospective period than just the previous few days. If hotspots are found to possess stable, enduring levels of high crime, it suggests that in practice they provide an effective means for targeting crime prevention resources to the places where high levels of crime are likely to persist (unless action is carried out to address the crime problems in these areas).

3.7. Examining the influence that recent incidents of crime have on predicting different future periods of crime

Hypothesis 6: Recent incidents of crime provide an effective means of accurately predicting the immediate future, but the accuracy in these predictions reduces for longer periods of the future.

Empirical findings into the patterns of repeat and near repeat victimisation have shown previous incidents to be very good predictors of where crime is likely to occur in the future. However, these predictions may have a short temporal horizon, in that they are good at predicting the immediate future (i.e., where crime may occur in the next few days) but are weaker at predicting where crime is likely to concentrate and persist in the longer future (i.e., the next few weeks and months). If predictions of where crime hotspots are likely to exist in the future are better informed from hotspot analysis rather than from the prediction principles of repeat and near repeat victimisation, it may suggest that a multi-method approach to crime prediction should be considered.

3.8. Examining the use of geographically weighted regression for helping to identify why hotspots exist, and for informing spatial predictions of crime

Hypothesis 7: GWR provides an effective means of determining at the local level the reasons why hotspots exist, and why these explanatory variables vary between hotspots.
Hypothesis 8: GWR analysis can be effectively used for supporting long-term predictions of crime by examining how a change in explanatory variables can influence a change in future crime levels.

Hotspot analysis and other mapping methods identify where crime is likely to take place in the future, but do not determine why crime concentrates in these areas. GWR provides an analytical framework for helping to determine the variables that explain why crime concentrates in certain locations, and if these explanations vary spatially. For example, the reasons for the presence of one hotspot in a study area may be different for the reasons for another hotspot in the same study area. If GWR provides an effective analytical framework for identifying these explanatory variables, this provides promise for improving prediction mapping techniques through the inclusion of these variables alongside, or in replacement of, retrospective patterns of crime. In addition, if the inference between explanatory variables and levels of crime can be effectively deduced using GWR, this could provide a means for supporting longer-term predictions by determining how the change in one variable would influence future changes in crime. That is, the results from a GWR analysis can be used to predict how crime would likely change based on the change in a variable that is related to the cause of crime.

4. Method

4.1. Introduction

This PhD research is conducted as a progressive sequence of studies. Findings from each study feed into the next study. This means that rather than providing a description of the full research method in a single section, it was more appropriate to describe the method used for each research part in turn, but begin here by describing the methodological framework that was generic across the entire research.

This chapter includes a description of the data and software used, statistical approaches for measuring the prediction performance of mapping output, and the processes for applying the range of spatial analysis techniques. Subsequent chapters on each empirical study begin by describing the methodological details relevant to that study. The research was conducted as seven studies, with each study relating to at least one of the research hypotheses posed in chapter 3:

- Study 1: Testing for the presence of hotspots
- Study 2: A metric assessment of the prediction performance of common hotspot analysis techniques
- Study 3: A metric comparison of the influence that technical parameters used in hotspot analysis can have on spatial crime prediction performance
- Study 4: Improving hotspot analysis using spatial significance mapping
- Study 5: Examining the temporal stability of hotspots
- Study 6: Examining the influence that recent incidents of crime have on predicting different future periods of crime
- Study 7: Examining the use of geographically weighted regression for helping to identify why hotspots exist, and for informing spatial predictions of crime

4.2. Software

A key methodological consideration for this research involved examining in detail the features and parameter requirements of each technique in order for processes to be repeated by practitioners. This meant the choice of software was oriented towards GIS products commonly used and freeware spatial statistical tools that were easy for practitioners to access to help promote replication.

ESRI ArcGIS version 10.1 and Pitney Bowes MapInfo Professional version 10.5 were used in this research. Both products were used in order to test the feasibility of the research that was conducted, and to harness the respective strengths of each to facilitate the analytical experiments and the production of visual outputs. In addition to ArcGIS and MapInfo, the extensions ESRI Crime Analyst (ESRI(UK), 2013) and Hotspot Detective for MapInfo (Ratcliffe, 2004) were used for hotspot analysis in studies 2 and 3 of the research. CrimeStat version 3 (Levine, 2010), a freeware software product developed by the United States National Institute of Justice, was used for producing spatial ellipses in study 2. Study 6 of the research used the online Vigilance Modeller tool (Astun Technology, 2013) for producing prospective mapping output for import into ArcGIS. ArcGIS GWR and GWR version 4 were used for modelling in study 7[6].

Listed below are the research parts and the main functions that were used from each of the GIS and spatial analysis software products.

- Data preparation including geocoding accuracy tests: MapInfo SQL for selecting crime type data sets from an all crime data set for each study area, MapInfo for running geocoding tests on the sample recorded crime data, and MapInfo Universal Translator for the conversion of crime data from MapInfo format to ArcGIS shapefile format
- Study 1: Testing for the presence of hotspots: ArcGIS nearest neighbour index function
- Study 2: A metric assessment of the prediction performance of common hotspot analysis techniques: CrimeStat (generation of spatial ellipses) and MapInfo (thematic mapping of Census Output Areas, thematic mapping of grid cells) and Hotspot Detective for MapInfo (kernel density estimation)
- Study 3: A metric comparison of the influence that technical parameters used in hotspot analysis can have on spatial crime prediction performance: ArcGIS Crime Analyst extension (kernel density estimation), ArcGIS raster to point conversion, ArcGIS point to Theissen polygon conversion and ArcGIS SQL (Structured Query Language) for determining cumulative counts of the area searched relative to the number of offences committed in the prediction period for the calculation of prediction measures

[6] R code for GWR has also been developed (see: http://www.st-andrews.ac.uk/geoinformatics/gwr/gwr-software/), from which the results can be displayed and analysed in the researcher's preferred GIS.

- Study 4: Improving hotspot analysis using spatial significance mapping: ArcGIS Getis-Ord Gi* statistic and ArcGIS SQL for determining cumulative counts of the area searched relative to the number of offences committed in the prediction period for the calculation of prediction measures
- Study 5: Examining the temporal stability of hotspots: MapInfo for using imported ArcGIS Getis-Ord Gi* hotspots and MapInfo SQL for calculating values for measuring temporal stability
- Study 6: Examining the influence that recent incidents of crime have on predicting different future periods of crime: Vigilance Modeller for generating prospective mapping outputs, ArcGIS for mapping Vigilance Modeller output and calculating the volume of crime in prospective mapping risk areas
- Study 7: Examining the use of geographically weighted regression for helping to identify why hotspots exist, and for informing spatial predictions of crime: ArcGIS GWR for running OLS and Gaussian Models, GWR version 4 for running Poisson Models, ArcGIS for displaying GWR mapping outputs.

4.3. Study areas and crime data

4.3.1. Study areas

The research used data from two areas – Newcastle-upon-Tyne, and the combined boroughs of Camden and Islington in London. These two areas were selected due to good contacts with the police in both areas (in order to arrange the supply of crime data) and because the two areas are quite different in their geography - while both are city areas, the landscape of Camden and Islington is dominated by its urban geography, while Newcastle has a vibrant city centre, but with suburbs to the East, North and West, and which extend towards rural parts of Northumberland County. Crime levels between the two areas also differ, with on average 64,000 recorded offences per year in Camden and Islington and 24,000 per year in Newcastle. These contrasts between the study areas provided the opportunity for the empirical studies to examine how different levels of crime, and consequently different levels of the spatial concentration of crime, influenced the results. Further information on each of the study areas now follows.

The London Metropolitan Police provided recorded crime data for the period of the 1st January 2009 to 31st December 2010 for the London Borough of Camden and the London Borough of Islington. The Camden/Islington study area is located in Central/North London. This area encompasses a wide range of urban geography, including three mainline train stations (Euston, Kings Cross and St. Pancras), a Premier League football stadium (Arsenal FC), popular shopping areas such as Tottenham Court Road, Holborn, Angel and Camden Market and large open parks such as Hampstead Heath and Regents Park. The area contains a synthesis of different land uses (commercial, retail and residential), and a resident population that reflects London's cosmopolitan diversity. The area experiences a high influx of visitors who commute to the area for work, education (UCL is based in the Bloomsbury area of the London Borough of Camden), for shopping, tourism, or visit the area for its theatres, cinemas, restaurants, bars, music and other attractions. Figure 4.1 shows a map of the study area with some supporting general statistics about the area from the 2011 Census of England and Wales.

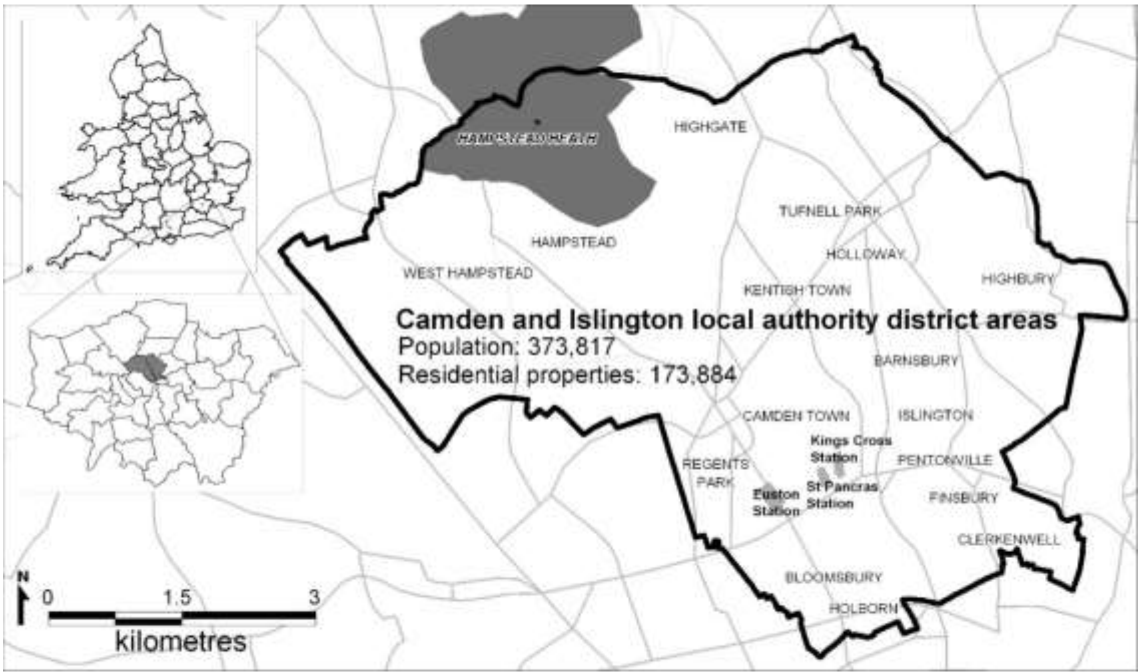

Figure 4.1. The Camden/Islington study area in Central/North London

Northumbria Police provided recorded crime data for the period 1st October 2009 to 30th September 2010 for the district of Newcastle-upon-Tyne in the North-East of England. Newcastle is one of England's largest ten cities and therefore includes many of the urban geographical features and amenities that one would expect in a typical city. This includes a vibrant shopping and entertainment area in the centre of the city, a large number of economic and commerce functions, a mainline train station, a metro system, a Premier League football stadium (Newcastle FC), and two large universities. The district also includes rural areas towards the north (see Figure 4.2). The district population was

292,400 and the number of households 123,242 at the time of the 2011 Census of England and Wales.

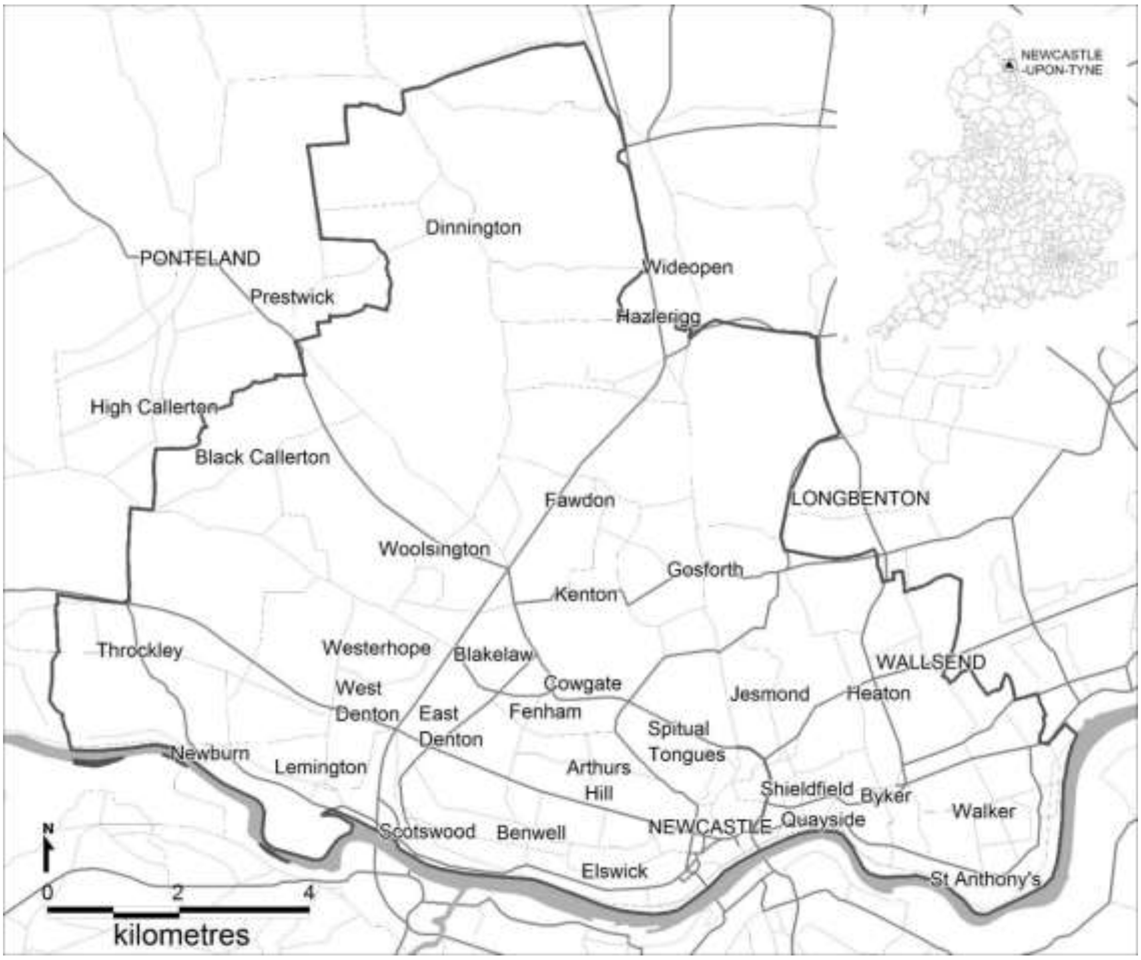

Figure 4.2. The Newcastle-upon-Tyne study area in North East England

4.3.2. Crime classifications, definitions and crime volume in the two study areas

The recorded crime data from both areas were for burglary to a dwelling (hereafter referred to as burglary dwelling), theft from the person, theft from vehicles and theft of vehicles. Northumbria Police were also able to provide data on assaults with injury for Newcastle. The Metropolitan Police were not able to provide data on assaults with injury due to issues they had experienced in the recording of these data for the period that crime data was made available. The following provides definitions of each offence:

- Burglary dwelling is defined as the breaking and entering into a residential dwelling and stealing property. It includes aggravated burglary and distraction burglary (e.g., entry gained by the offender posing as a utility engineer).
- Theft from the person involves theft without the use or threat of force such as when the goods stolen were being worn by the victim, or the goods stolen were physically attached in some way to the victim, or carried by the victim, or the goods stolen were

contained in an article of clothing being worn by the victim. For example, if a person was walking along a street holding their mobile phone, and the phone was snatched out of their hand, this would be recorded as a theft from the person.

- Theft from vehicles is defined as removing and stealing property from within and on the vehicle. This includes the theft of items attached in some way to the interior of the vehicle such as a radio/CD player, the theft of loose items within the vehicle such as a bag, electronic items (e.g., mobile phone, satellite navigation system), or money, the theft of items from the exterior of the vehicle such as a licence plate, exhaust system, wing mirror and badges, and the theft of petrol from siphoning from the petrol tank.
- Theft of vehicles is defined as taking without having the consent of the owner or other lawful authority, a vehicle for the offender's use or use by another, or knowing that any vehicle has been taken without such authority, drives it or allows himself to be carried in or on it. For example, a car is stolen from a car park, with the offender's intent for personal use, selling the vehicle, giving to another person or breaking the vehicle into parts for sale or use. Theft of a vehicle does not include the break-in to a house where the vehicle keys are taken and the vehicle in stolen. This is recorded as a burglary dwelling.
- Assault with injury involves the malicious wounding or inflicting of grievous bodily harm, assault occasioning actual bodily harm, a driver injuring persons through furious driving, assault with intent to resist apprehension, or owner or person in charge allowing a dog to be dangerously out of control in a public or non-public place and injuring a person. Over 90% of assaults with injuries are assaults occasioning actual bodily harm.

These crime types were chosen because they are groupings that are regularly analysed by police and public safety practitioners, and they are crime types of high volume that are of particular concern to the public and consume a great deal of police resources; therefore, the implications of the research would be accessible, of particular interest, and could be more readily translated into policing and public safety practice. Table 4.1 lists the number of offences for each crime type in the two study areas for the time period that data were provided. The differences in the periods for which crime data were provided were due to limitations imposed by the two police forces. However, it was judged that these would not affect any of the studies because both sets of crime data were sufficient for the analytical experiments that were to be conducted.

Table 4.1. Number of offences by crime type for Camden and Islington (1 January 2008 to 31 December 2010) and Newcastle (1 October 2009 to 30 September 2010).

Crime type	Camden and Islington (1 Jan 2009 - 31 Dec 2010)	Newcastle (1 Oct 2009 - 30 Sept 2010)
Burglary dwelling	11971	1304
Theft from the person	10,160	781
Theft from vehicle	22,328	1740
Theft of vehicle	8,385	196
Assault with injury	-	1838

4.3.3. Geocoding and geocoding accuracy

Geocoding is the process of assigning geographic coordinates to a record using the recorded address or some other geographic or location description. The geographic coordinates are most typically derived from a gazetteer - a look-up table that lists addressable and non-addressable locations and references these by their relevant geographic coordinates. These coordinates are usually precise to one metre. For example, geocoding a crime record for a burglary dwelling would involve using the recorded details on the house number, street name and locality to determine the dwelling's geographic coordinates by matching this address to its gazetteer entry. The geocoding process then attaches these geographic coordinates to the crime record. For a theft from the person that has occurred in a public place such as a café, the address details of the café, such as its name and street name, would also be recorded in the gazetteer and used in the geocoding matching process. In many cases the offence cannot be attached to an addressable property, such as an offence in a car park. Car parks are also included in the gazetteer, but usually a central geographic coordinate to the car park rather than the exact car parking space is used in the geocoding process. In some situations the address details entered on the crime record may not be complete. For example, it may be recorded that an assault occurred on the street. In such cases the name of a venue, shop, place of interest or street junction is also recorded to assist in locating as precisely as possible where the offence took place. These location details are then used in the geocoding process to determine as precisely as possible the geographic coordinates relevant to the location of this offence. In those cases where only the street name is recorded in the crime record, the offence is positioned to some central location, but reviewed to ensure that a large number of offences are not subject to this type of geographic placement.

The two sets of geocoded crime data were validated using a method for geocoding accuracy analysis as reported in Chainey and Ratcliffe (2005, 61-63). This geocoding accuracy approach tests a sample of 10% of the offence records and analyses whether the geographic placement of the offences on a detailed Ordnance Survey MasterMap displayed within a GIS matches with the address information on the crime record. For example, the researcher may test whether a burglary recorded at 5 Acacia Avenue has been correctly geocoded and is geographically positioned on the OS MasterMap to this address. Analysis of both sets of recorded crime data revealed them to be more than 95% accurate to the street address level and fit for purpose for this research. Three-hundred-and twenty-four of the crime records for Camden and Islington were not geocoded because of incomplete or inaccurate address information on the crime record. This represented less than 1% of records for each of the four Camden and Islington crime categories. Seventy-eight crime records for Newcastle could not be geocoded for similar reasons. Again, for no crime type was more than 1% of its original total affected. The records that were not geocoded were deleted from the crime data from both study areas. Table 4.1 lists only the number of offences that were successfully geocoded.

4.3.4. Temporal input periods, measurement date and temporal output periods

A suitable date had to be chosen from each crime dataset as the day on which retrospective data were selected to generate hotspot maps against which *future* events could be compared (this date is referred to as the *measurement date*). In the first instance, for simplicity, the mid-point in the Camden/Islington data (the 1st January 2010) and the mid-point in the Newcastle data (1st April 2010) were chosen in order to maximise the use of 12 months of retrospective data for generating hotspot maps in Camden/Islington and 6 months of retrospective data for Newcastle. This also meant that 12 months of data for Camden/Islington and six months of data for Newcastle after each measurement date could be used for measuring the accuracy of the different mapping techniques for predicting future events. The 1st of January (New Year's Day), however, is a day when the normal routines of people's day-to-day lives could be quite different to most other days in the year. The three-week period before the 1st of January could also be considered as unusual as it is the period before and during Christmas, when again the routine activities of peoples' lives could be quite different to other dates in the year. For this purpose another more *normal* measurement date was selected for the Camden/Islington dataset. The second measurement date chosen was the 11th March 2010: this date, being a Thursday, was during school term time and was considered as a date on which most people would go about their routine activities in a manner that was similar to many other dates in the year. Choosing an alternative measurement date of the 11th March 2010 meant that the results generated for this measurement date and the results generated using a measurement date of the 1st January 2010 could be compared.

Table 4.2. The temporal periods of input data used in the research studies for (a) a measurement date in Camden/Islington of the 1st January 2010, (b) a measurement date in Camden/Islington of the 11th March 2010[7], and (c) a measurement date in Newcastle of the 1st April 2010

(a)

Input data time periods used in experiments for a Camden/Islington measurement date of 01/01/10									
12 months	6 months	3 months	2 months	1 month	2 weeks	1 week	3 days	2 days	1 day
01Jan09-31Dec09	01Jul09 - 31Dec09	01Oct09-31Dec09	01Nov09-31Dec09	01Dec09-31Dec09	18Dec09-31Dec09	25Dec09-31Dec09	29Dec09-31Dec09	30Dec09-31Dec09	31 Dec09

(b)

Input data time periods used in experiments for a Camden/Islington measurement date of 11/03/10									
12 months	6 months	3 months	2 months	1 month	2 weeks	1 week	3 days	2 days	1 day
13Mar02-12Mar03	13Sep03-12Mar03	13Dec02-12Mar03	13Jan03 - 12Mar03	13Feb03-12Mar03	27Feb03-12Mar03	06Mar03-12Mar03	10Mar03-12Mar03	11Mar03-12Mar03	12 Mar03

(c)

Input data time periods used in experiments for a Newcastle measurement date of 01/04/10								
6 months	3 months	2 months	1 month	2 weeks	1 week	3 days	2 days	1 day
01Oct09 - 31Mar10	01Jan10-31Mar10	01Feb10-31Mar10	01Mar10-31Mar10	18Mar10-31Mar10	25Mar10-31Mar10	29Mar10-31Mar10	30Mar10-31Mar10	31Mar10

Each study areas' retrospective crime data were arranged into a number of time periods used as *input data* to generate hotspot maps. This meant that rather than using just one retrospective time period for the research (e.g., the three months prior to the measurement date), the use of a number of retrospective time periods could be considered together in

[7] Note that the maximum time range of output data for analysis in Camden and Islington based on the 11th March 2010 measurement date was nine and a half months and not twelve months.

order to provide a more reliable basis on which to draw findings. The input data were organised into the time periods shown in Table 4.2. In most cases, these input data periods were used for each research study. However, for certain experiments, only some of the data input periods were used or different ones were used due to the requirements of certain experiments. Confirmation on the input data that were used for each set of experiments is provided in the method sections for each research study.

Table 4.3. The temporal periods of output data used in the research experiments for (a) a measurement date in Camden/Islington of the 1st January 2010, (b) a measurement date in Camden/Islington of the 11th March 2010[8], and (c) a measurement date in Newcastle of the 1st April 2010.

(a)

2010 time periods used in experiments for a Camden and Islington measurement date of 01/01/10									
1 day	2 days	3 days	1 week	2 weeks	1 month	2 months	3 months	6 months	12 months
01Jan10	01Jan10 - 02Jan10	01Jan10 - 03Jan10	01Jan10 - 07Jan10	01Jan10- 14Jan10	01Jan10 - 31Jan10	01Jan10 - 29Feb10	01Jan10 - 31Mar10	01Jan10 - 30Jun10	01Jan10 - 31Dec10

(b)

2010 time periods used in experiments for a Camden and Islington measurement date of 11/03/10									
1 day	2 days	3 days	1 week	2 weeks	1 month	2 months	3 months	6 months	All of 2010
11Mar10	11Mar10 - 12Mar10	11Mar10 - 13Mar10	11Mar10 - 17Mar10	11Mar10 - 24Mar10	11Mar10 - 10Apr10	11Mar10 - 10May10	11Mar10- 10Jun10	11Mar10 - 10Sep10	11Mar10 - 31Dec10

(c)

2010 time periods used in experiments for a Newcastle measurement date of 01/04/10								
1 day	2 days	3 days	1 week	2 weeks	1 month	2 months	3 months	6 months
01Apr10	01Apr10 – 02Apr10	01Apr10 – 03Apr10	01Apr10 – 07Apr10	01Apr10 - 14Mar10	01Apr10 - 30Apr10	01Apr10- 30May10	01Apr10– 31Jun10	01Apr10- 30Sep10

If hotspot maps are generated to help determine where crimes may occur in the future, the definition of *the future* also needs consideration. For the purposes of this research, the analysis was limited to predicting crime patterns for up to a year from the measurement date for Camden/Islington and up to 6 months from the measurement date for Newcastle. For example, using the data from London meant that twelve months of

[8] Note that the maximum time range of output data for analysis in Camden and Islington based on the 11th March 2010 measurement date was nine and a half months and not twelve months.

crime data could be used when the measurement date was the 1st January 2010, and nine and a half months when the measurement date was the 11th March 2010. It was anticipated that the different time frames would not have implications on the empirical studies that were to be conducted because at least six months of data were considered to be sufficient for measuring differences in prediction levels.

In following a similar argument to the temporal arrangement of input data, data after the measurement date (referred to from this point as output data) were organised into temporal periods. This meant that rather than using just one output data period for the research (e.g., the three months after the measurement date), the use of a number of output data time periods would allow for comparisons in the results. While most experiments used the same output data periods, the requirements of some experiments resulted in certain other temporal output periods being used. Table 4.3 lists the output periods that were used in most experiments. Hotspot maps could then be generated for each period of input data and measured for their ability to predict spatial patterns of crime, when the prediction period was the next day, the next two days, the next week, and to the next 12 months. If the output periods used were different to those listed in Table 4.3, the details on the output periods used for a particular empirical study is described in that study's method section.

4.3.5. The possible impact of changes in crime patterns from police and crime prevention activity

During the data time periods that were used in the experiments there could have been police operations and crime reduction initiatives in both study areas that had an impact on crime. For the purposes of this research, the focus was on comparing analytical techniques against the same data. This would mean that any changes in crime patterns would be similarly applied to each of the techniques and would not affect the ability of making comparisons and drawing conclusions on the analyses that were conducted.

4.3.6. Summary of study area and crime data

Crime data for Camden/Islington and Newcastle covering the period of 2009 to 2010 provided an adequate data resource for the research studies that were proposed. The crime data were categorised into burglary dwelling, theft from the person, theft from vehicles, and theft of vehicles to offer a range of crime scenarios for the research. Assault with injury data were also provided for Newcastle, offering a valuable addition by permitting the comparison of violent crime against these other theft offences. The geography of the two study areas also provided an interesting contrast – both are city areas, however the landscape of Camden and Islington is dominated by its urban geography, while Newcastle has a vibrant city centre, with suburbs to the East, North and West, and rural areas that extend towards parts of Northumberland County. Crime levels between the two areas are also quite contrasting, with on average 6,000 burglary dwelling offences per year in Camden/Islington and 1,300 per year in Newcastle. The contrast between the two study areas provided the opportunity for the experiments to examine how different levels of crime and the spatial distribution of crime influenced the results.

The crime data were geocoded to a level that was fit for the purpose of the experiments proposed, and measurement dates were selected to maximise the full use of the crime data that was provided, but also to explore if results were consistent for different dates. Consideration towards the range of input and output data temporal periods provided a structure for testing the prediction performance of hotspot analysis techniques under different data conditions. The testing of data for different time periods also made any results less susceptible to the dangers of interpreting anomalies as representative results.

4.4. Measuring the prediction performance of mapping output

To date, no standard method has been proposed for determining how to measure the prediction performance of mapping output. In the next section, a number of methods that have been used for measuring prediction performance are described, new ones are introduced, and a framework of prediction performance metrics are proposed that can set the standard for how the prediction accuracy of mapping outputs can be measured.

4.4.1. Hit rate and search efficiency rate

One obvious measure for determining the effectiveness of mapping techniques for predicting where crime may occur would be a hit rate. Hit rate is the percentage of *new* crimes that occur within the areas where crimes are predicted to occur. While useful and easy to understand, this measure does not take into account the size of the areas where crimes are predicted to occur. For example, a hit rate could be 100%, but the area where crimes are predicted to occur could cover the entire study area – a result of little use to practitioners who have the need to identify where to target resources. As an alternative to the hit rate, Bowers et al. (2004) proposed the search efficiency rate. The search efficiency rate is the number of events per square kilometre in the areas where crimes are predicted to occur. This measure works well when considering just a single study area, but does not easily allow for comparisons between study areas of different sizes because the size of the entire study area should be considered in relation to the size of the areas where crimes are predicted to occur. For example, hotspots have been identified in a study area that is 10km² in area. From this, a search efficiency rate of 20 crimes per km² is calculated. A study area that is 50km² in size may have experienced the same volume of crime as the smaller study area and also have a search efficiency rate of 20 crimes per km². Yet in the larger study area there is more area where no crime has been predicted to occur (i.e., the areas between the hotspots), meaning that the areas where crime has been predicted to occur covers a smaller relative area than the predicted areas determined in the smaller study area. As a result, the hotspots identified for the larger study area provide a more useful basis from which to target resources, but with this greater utility not being captured in the search efficiency rate.

4.4.2. The prediction accuracy index

The current research introduces the prediction accuracy index (PAI)[9] devised by the author with colleagues as a measure of mapping output prediction performance. This index has been devised to consider the hit rate against the areas where crimes are predicted to occur with respect to the size of the study area. The PAI is calculated by dividing the hit rate percentage (the percentage of crime events that were committed in the areas where crimes were predicted to occur, i.e., the crime hotspots) by the area percentage (the proportion of the area where crime is predicted to occur i.e., the area of the hotspots, in relation to the whole study area: see Equation 2).

$$\frac{\left(\frac{n}{N}\right)*100}{\left(\frac{a}{A}\right)*100} = \frac{HitRate}{Area\,Percentage} = \text{Prediction Accuracy Index} \qquad (2)$$

In this formula, *n* refers to the number of crimes in areas where crimes are predicted to occur (i.e., the number of *future* crimes in the hotspots), *N* is the number of *future* crimes in the study area, *a* refers to the area (e.g., km²) of areas where crimes are predicted to

[9] Since the commencement of this PhD research a paper by the author that introduces and uses the PAI has been published: Chainey, S.P., Tompson, L., Uhlig, S. (2008a), "The utility of hotspot mapping for predicting spatial patterns of crime", *Security Journal* 21:1-2.

occur (i.e., the area of the hotspots), and *A* is the area (e.g., km^2) of the study area. *Future* crimes means the crimes that were committed during the output period for which the mapping techniques are measured.

Finding 100% of future events in 100% of the area would give a PAI value of 1. If the hit rate and the area percentage fall by an equal measure, the value would also be computed as 1. Finding 25% of future crime events in 50% of the study area would return a PAI value of 0.5; and finding 80% of future crime events in 40% of the area would return a PAI value of 2. Thus, the greater the number of *new* crime events in a hotspot area that is smaller in areal size to the whole study area, the higher the PAI value.

The PAI is easy to calculate and considers the number of *future* crimes that fall into the areas determined as hotspots against both the size of the hotspots and the size of the study area. The PAI is also a measure that is applicable to any study area, to any crime point data, and to any analysis technique that aims to predict spatial patterns of crime. Many practitioners still, though, find great use in the hit rate as a measure to predict how many crimes they may be able to impact upon by targeting resources to just the hotspot areas. As the hit rate is calculated as part of the PAI, the hit rate can be considered alongside a PAI calculation. The use of the PAI has subsequently been discussed further by Pezzuchi (2008), Levine (2008) and Chainey et al. (2008b; 2008c), with researchers concluding it to be a useful measure for comparing multiple hotspot mapping outputs. These discussions have included illustrating how chance expectation can be minimised by using the mean PAI results from a large number of experiments across different temporal input and temporal output data periods, and by observing the variation in the standard deviation of the PAI generated from these many experiments.

4.4.3. Accuracy concentration curves

Since the PAI was introduced, other approaches for measuring the predictive performance of mapping output have been developed. Perhaps the most rigorous of these is described by Johnson et al. (2008b; 2012). The problem with a single measure such as the PAI is that it only offers a comparison between one hit rate and one defined hotspot area, and no comparison against chance expectation or across the full range of prediction from 1% of all offences to the prediction of 100% offences. Johnson et al. (2008b; 2012) proposed the use of an accuracy concentration curve. An accuracy concentration curve is generated by plotting the percentage of crimes that have been accurately predicted against the incremental risk ordered percentage of the study area. This accuracy concentration curve process involves plotting the number of future crimes in 1% of the study area (the 1% of the study area containing the highest hotspot values); plotting the number of future crimes in the areas containing the top 2% of hotspot values in the study area; plotting the number of future crimes in the areas containing the top 3% of hotspot values in the study area; and so on, until the number of future crimes in the areas containing 100% of the study area are plotted. The plots of these individual readings are then connected with a line that represents the accuracy concentration curve for these data. In practice, the process may also involve determining readings for the incremental proportion of future crimes rather than for just single percentage increments of the study area. This is because readings for the proportion of future crimes may be high in comparison to the proportion of the study area. For example, 10% of future crimes may fall within 1% of the study area, therefore, it would also be useful to determine and plot the proportion of the study area relating to 1%, 2%, 3% and so on, to 10% of future crimes. Incremental readings for the proportion of future crimes, where this proportion is high in comparison to very small proportions of the study area, are of particular interest in the measurement of the spatial prediction performance of mapping output. These measurements for very small proportions of the study area are useful because they would indicate the proportion of crime that is predicted, and could be prevented by targeting policing and public safety activity towards these very small areas.

Generating accuracy concentration curves means that mapping output for different crime types and under different conditions (e.g., the use of different temporal input and output periods) can be compared across the complete range of predicted offences against the size of the area that would need to be searched before the full range of predicted offences were identified. Of most interest in comparing these curves is how vertical the curve is, particularly for small proportions of the study area size - the more vertical the observed curve, the better the prediction performance of the mapping output from the technique being tested. Figure 4.3 shows an example of an accuracy concentration curve for kernel density estimation from Johnson et al.'s (2008b) study. The accuracy concentration curve shows the proportion of the study area that needed to be searched to identify the proportion of offences represented on the *y* axis. Johnson et al. (2008b) also generated upper and lower threshold curves against chance expectation from 99 runs of a Monte Carlo simulation. The lower threshold was based on the mean expected accuracy and the upper threshold determined the 95th percentile against chance expectation.

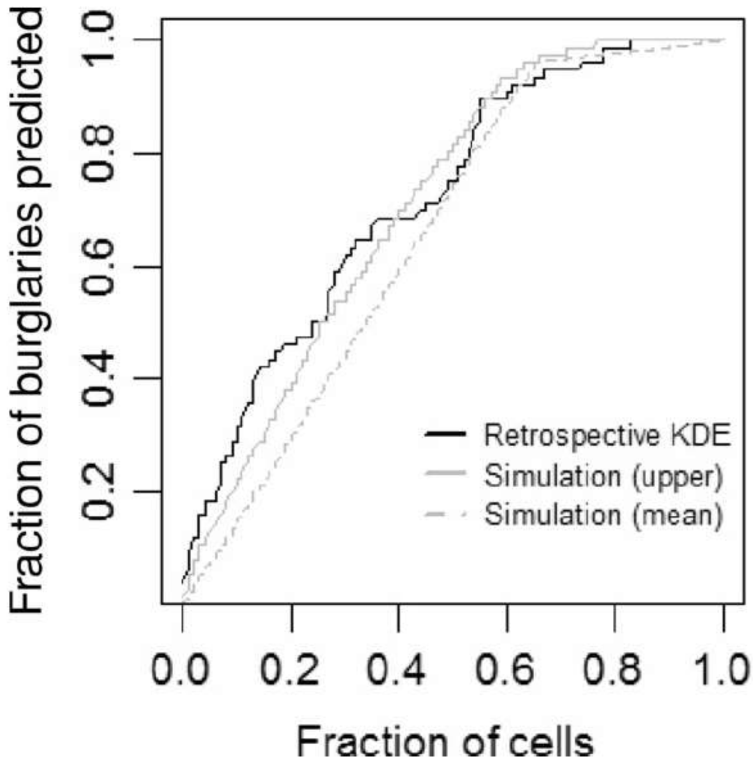

Figure 4.3. An accuracy concentration curve for KDE hotspot analysis, compared to chance expectation. The mean threshold indicates the observation to be no better than chance. The upper threshold indicates whether the observed results were greater than 95% significant. The more vertical the observed curve, the better the prediction (Source: Johnson et al., 2008b).

4.4.4. Area under the curve

Since the 1970s, the medical profession has made use of the *area under the curve* approach to help determine the accuracy of medical trials and for clinical predictions. The area under the curve approach has also been used in other fields of science and engineering, including psychology to predict behaviour, such as the chances of a person offending (Dolan, 2000; Mossman, 1992). The approach (originally referred to as the area under the ROC curve) was initially developed from signal detection theory during World War II for measuring the ability of radar receiver operators to detect whether a blip on a radar image was an enemy target, a friendly ship, or noise. The radar receiver operator's ability to do so was called the receiver operating characteristics (ROC). A ROC curve chart is created by plotting the fraction of true positives out of the total actual positives against the fraction of false positives out of the total actual negatives. The true positive rate is also known as sensitivity, and the false positive rate is one minus the specificity or true negative rate (Hanley and McNeil, 1982).

Figure 4.4 shows three ROC curves representing results that relate to the trial of a medical drug: (A) perfect impact of the trial drug i.e., everyone got better, (B) good impact of the trial drug i.e., lots of people got better, and (C) no impact of the trial drug i.e., no better than chance expectation. The effectiveness of the trial drug relates to how well it separates the group of patients being tested into those with or without the disease in question.

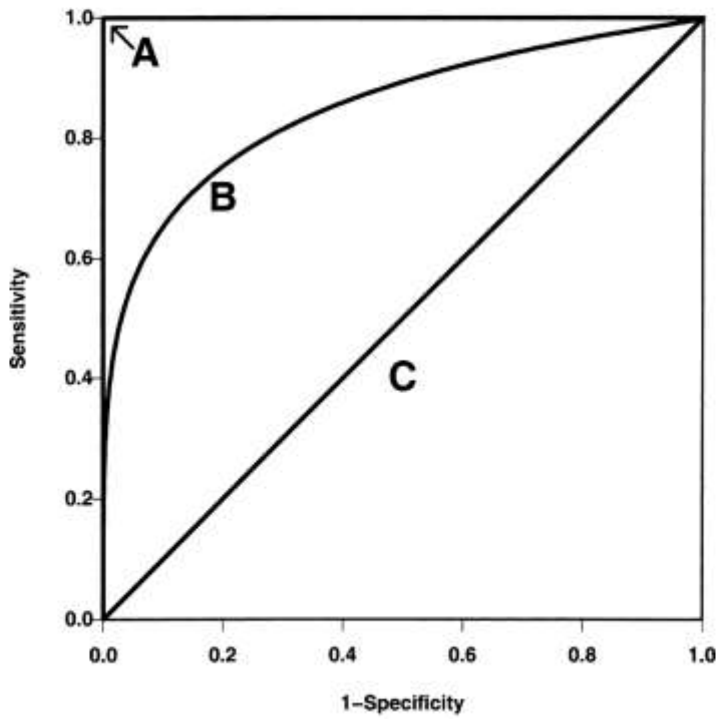

Figure 4.4. An example of three ROC curves, with A representing a perfect impact of trial drug or perfect prediction, B good results and C a result that is no better than chance expectation (Source: Zou et al., 2007).

The accuracy of the medical trial is measured by the area under the ROC curve. An area of 1 represents a perfect test and an area of 0.5 represents a worthless test. A rough guide that has been developed for classifying the accuracy of a diagnostic test utilises the following point system:

- 0.90-1 = excellent (A)
- 0.80-0.90 = good (B)
- 0.70-0.80 = fair (C)
- 0.60-0.70 = poor (D)
- 0.50-0.60 = fail (F)

An example of the use of the area under the curve approach is shown in Wigton at al.'s (1986) clinical study into predicting strep throat. In this study they found the presence of fever and two other medical conditions, and the absence of a cough, all predicted strep. Figure 4.5 shows a comparison between the findings from two study areas, showing that prediction rules performed more accurately for patients from Virginia (the area under the curve = 0.78) compared to Nebraska (the area under the curve = 0.73).

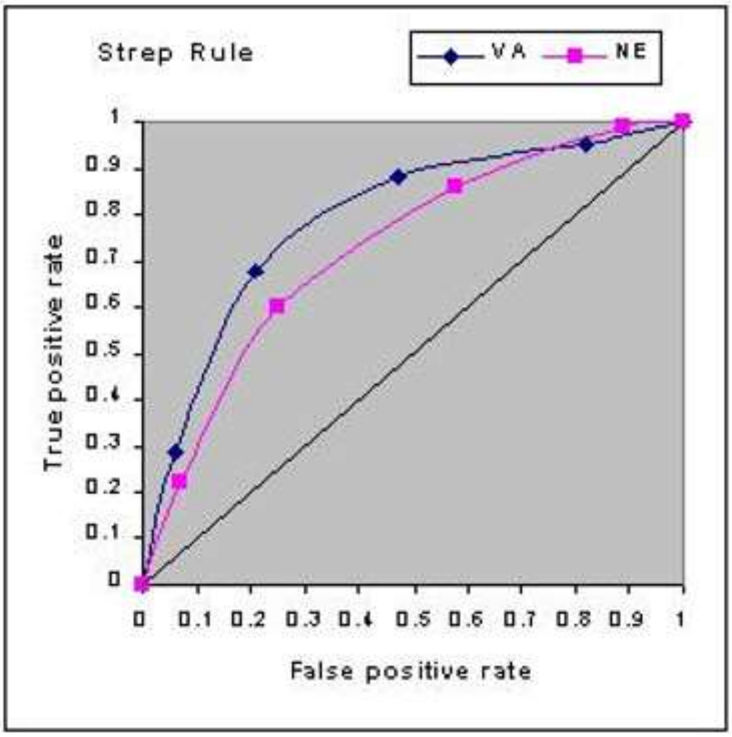

Figure 4.5. A comparison of ROC curves for two areas (Virginia and Nevada) from a medical trial that tested the variables that predicted strep throat (Source: Wigton at al., 1986).

There are a number of methods that can be used for measuring the area under the curve. The two most common are a non-parametric method that is based on constructing trapezoids under the curve as an approximation of the area, and a parametric method using a maximum likelihood estimator to fit a smooth curve to the data points.

A weakness in the predictive crime mapping research to date is the absence of a universal measure that is easy to compute and that allows for easy direct comparisons between the predictive accuracy of mapping output that is generated from different spatial analysis techniques. This current thesis introduces the application of the area under the curve approach to accuracy concentration curves for measuring the predictive accuracy of hotspot analysis output (and any predictive mapping technique's output). The trapezoid method for measuring the area under the curve is used as it is the easier of the two common types of methods to produce, and can be calculated using a simple formula in Microsoft Excel, allowing for easy replication by practitioners. The trapezoid method works by dividing the area under the curve into a series of trapezoids, calculating the area for each, and then summing these to determine the total area. The area of each trapezoid is calculated by determining the average height of the curve between two points and multiplying this by the base. In Excel, this requires pairs of *x* axis coordinates (the proportion of the study area) to be identified relative to pairs of *y* coordinates (the proportion of offences) for the trapezoid input calculations. The area under the curve can be calculated using the following formula in Excel:

SUMPRODUCT(A2:A100-A1:A99,(B2:B100+B1:B99)/2 (3)

where A is the *x* array of values (i.e., proportion of study area), B is the *y* array of values (i.e., proportion of offences), representing an example coverage of 100 values.

A distinction between measuring the predictive accuracy of mapping output compared to the approach used for clinical prediction is that more value is placed in a mapping technique that does well at predicting a high proportion of crime in small geographic areas, rather than a general measure for the mapping technique for the whole study area. That is, measuring the proportion of *future* crime that takes place in the relatively small areas that were identified as hotspots (using retrospective crime data), rather than the proportion of crime that took place in all areas for which a hotspot mapping technique calculated values for. Therefore, rather than generating a single area under the curve measure, values for this metric should be calculated across a range of area and offence proportion sub-sections of the accuracy concentration curve.

Table 4.4 lists the proposed combinations of the study area and offence proportions for measuring the area under an accuracy concentration curve for the purpose of determining the spatial crime prediction performance of mapping output. An area under the curve measure from the section of the accuracy concentration curve covering 0.5% of the study area and 5% of offences will indicate the predictive ability of the mapping output for very small areas (i.e., areas for focused police or crime prevention activity). This then extends to 1% and 5% of the study area's coverage area, relative to 10% and 25% of offences respectively. These metrics continue in order to give a sense of any improvements or degradation in the mapping outputs prediction performance for larger coverages of the study area (including 80% of the offences relative to 20% of the coverage area to utilise the albeit arbitrary but indicative practice of the 80:20 rule (also referred to as the Pareto Principle) that is commonly applied in crime analysis (Clarke and Eck, 2003)). A final measure relating to the area under the entire curve (i.e., 100% of the coverage area and 100% of offences) is also calculated.

Table 4.3 also lists the maximum area under the curve for each of these metrics. For example, for a perfect test covering the chart axis range of 0%-10% of the study area and 0%-50% of all offences, the maximum area would be 0.1*0.5 = 0.05. The closer the observed area under the curve is to the maximum value, the better the prediction performance of the mapping output.

Table 4.4. Proposed sub-sections for area under the curve measurements to test for mapping output prediction performance

Area under curve metrics	0.5% area X 5% offences	1% area X 10%offences	5% area X 25% offences	10% area X 50% offences	20% area X 80% offences	100% area X 100% offences
Maximum area under section of curve	0.000250	0.001	0.0125	0.05	0.16	1

4.4.5. The crime prediction index

While the area under the curve provides a more complete measure for assessing the prediction performance of mapping output, the values generated across each range of study area and offence proportions do not allow for straight-forward comparisons. The current research introduces the crime prediction index (CPI) as a simple way to overcome this. The CPI is calculated by dividing the observed area under the curve for a given proportion of the study area and offence range, by the maximum area under the curve for this range.

$$\frac{\text{Observed Area Under Curve}}{\text{Maximum Area Under Curve}} = \text{Crime Prediction Index} \qquad (3)$$

For example, if an observed area under the accuracy concentration curve for the 0.5% coverage area and 5% offences range was 0.000184, and as the maximum area under this section of the curve would be 0.00025 (i.e., 0.005*0.05), the CPI would be 0.736 (1 is a perfect prediction). If an observed area under the curve for the 5% coverage area and 25% offences range was 0.0081, and as the maximum area under this section of the curve would be 0.0125 (i.e., 0.05*0.25), the CPI would be 0.648. The results using these examples would show that the mapping method performs better at predicting crime for smaller coverage areas than for larger coverage areas. The CPI therefore chimes with the medical approach of using a classification point system for the accuracy of a diagnostic test (e.g., 0.90 to 1 is an excellent prediction). The CPI also makes it easier for mapping techniques to be compared for their prediction performance – the closer the observation is to a value of 1, the better the prediction.

4.4.6. Summary of measuring the prediction performance of mapping output

Since 2008 a number of techniques have been developed to measure the prediction performance of crime mapping output. These have included the hit rate, the search efficiency rate and the prediction accuracy index. Of these, the PAI has become most used by crime researchers. However, it is a simple method that does not allow for a full comparison of the prediction performance of the mapping output across the complete range of offence proportions in comparison to the size of the area that would need to be searched to predict this level of offences. The accuracy concentration curve provides a useful means for this more complete analysis of prediction performance. Curves for different types of mapping techniques can be plotted on the same chart to compare prediction performance. However, rather than just comparing curves, the area under the curve and the crime prediction index each provide single, robust measures that allow for a more straightforward comparison of the prediction performance between mapping outputs. These approaches borrow the area under the curve approach used in several other areas of science including medical trials and clinical predictions, but are customised for policing and crime prevention by applying these metrics to sub-sections of the accuracy concentration curve. These adaptations are in order to identify how spatial crime prediction methods perform when only a very small coverage area is searched for the offences that are predicted in comparison to how it performs as the area that needs to be searched increases. That is, if the accuracy concentration curve is very vertical within the 0%-50% range of offence proportion values and where the proportion of the study area is very small, and the CPI for this sub-section of the accuracy concentration curve is close to a value of one, this would suggest the mapping technique is an excellent predictor of crime. If this curve then starts to flatten after the 50% offence proportion level, it could be argued that the hotspot mapping technique has *done its job* by identifying small geographic areas (in proportion to the study area) where it predicted a high proportion of crime to take place.

The experiments carried out in the current research use a combination of the PAI, hit rates, accuracy concentration curves, the area under the curve and the CPI to measure the prediction performance of mapping outputs. This is in order to conduct the first assessment of how these different prediction performance measures compare.

4.5. Spatial analysis processes: testing for clustering, hotspot mapping techniques and input parameters, temporal stability, temporal predictions, and geographically weighted regression

4.5.1. Testing for clustering

The first research study involves testing each recorded crime data set for statistical evidence of clustering. The application of these tests will determine the point in the temporal input periods at which clustering is evident, and if this clustering continues across all the temporal data input periods. Establishing if clustering is evident will identify whether the data for this input period is suitable for further hotspot analysis. The technical methodological processes for testing for statistical evidence of clustering are explained in research study 1 (chapter 5).

4.5.2. Hotspot mapping techniques and input parameters

Research study 2 involves a metric comparison of the prediction accuracy of common hotspot analysis techniques. Each of these techniques requires certain parameters to be set for them to operate or for them to generate a visual map output. The research will experiment with different parameters settings in order to identify the impact these have on the prediction accuracy of mapping outputs. A technical description of each hotspot analysis technique and the parameters each technique uses is provided in research study 2 (chapter 6).

If a common hotspot analysis technique is consistently identified to perform better than the others for predicting spatial patterns of crime, this technique will be further examined for the impact its input parameter settings have on the prediction performance of the output it generates. For example, if kernel density estimation is identified as the best of the common hotspot analysis techniques, the two parameters it requires the researcher to enter – the cell size and the bandwidth size – will be tested to identify the impact that different sized parameters have on the prediction performance of the mapping output. The technical methodological processes for this hotspot analysis technique's parameter settings will be described in detail in research study 3 (chapter 7).

Study 4 (chapter 8) of the research involves testing whether spatial significance mapping using Local Moran's I, Local Geary's C and the Gi* statistic improves on the common hotspot analysis techniques. The focus of these tests are to determine whether spatial significance mapping output can unambiguously define the areas that are hotspots, and if these outputs perform better at predicting where crime occurs in the future. The technical features of these techniques, including the parameters they use, will be explained in research study 4. The influence these technical parameters have on the outputs and in the prediction performance of these spatial significance mapping techniques will also be tested.

4.5.3. Temporal stability of hotspots

Study 5 (chapter 9) of the research involves examining the temporal stability of hotspots. The study will involve using the hotspot analysis technique that consistently performed the best in predicting spatial patterns of crime from research studies 2, 3 and 4. The analysis will determine if the hotspots it identifies are based on high crime levels that have persisted over the input data period, if high crime levels then endured in the hotspots it identifies over the temporal output periods, and if these high levels of crime differ between hotspots that were generated using different temporal data input periods. For example, input data for 1 month and 3 months will be used for generating hotspots, from which the data for the period before and during the input data period will be analysed to identify if crime levels have remained stable and high in the hotspots, and if crime levels have continued to remain stable and high across the data output periods in the hotspots. This part of the research introduces the temporal stability index (Haberman and Ratcliffe, 2012) and will be explained in full in the method section for this research part.

4.5.4. Temporal predictions of mapping output

Study 6 of the research examines the influence that different retrospective periods of crime data have on spatial crime predictions for different time periods of the future. This will involve testing the influence of repeat and near repeat patterning principles that have been coded into the Vigilance Modeller prospective mapping tool. This process involves using data for the seven days prior to a measurement date as the inputs to the Vigilance Modeller. The mapping outputs from this technique will be tested to identify its prediction performance for the near future (within the next seven days), and for longer temporal output periods (e.g., the next two weeks, the next month, the next six months). A full description of this process is provided in the method section for research study 6 (chapter 10). Research study 6 also involves comparing the research findings from this set of experiments to the hotspot analysis results from parts 2 to 4. This will include identifying whether prospective mapping performs better than hotspot analysis for predicting spatial patterns of crime, and if this performance differs for the temporal output periods. That is, the analysis will determine if prospective mapping performs better than hotspot analysis in predicting where crime is likely to occur in the immediate future (i.e., the next few days), and for predicting where crime is likely to occur for longer periods of the future (i.e., the next week, the next two weeks and for other periods beyond).

The prediction performance of each mapping output for research studies 2 to 6 will use a combination of the prediction performance metrics described in section 4.4. GIS functions, as described in section 4.2, will be used for processing hotspot analysis techniques. The results of these functions will be calibrated in Microsoft Excel for further analysis and for producing tabular and charting outputs.

4.5.5. Geographically weighted regression

The detailed methodological process for applying geographically weighted regression is described in research study 7 (chapter 11) where this method is used. This includes a description of the types of models that are applied, and the diagnostic statistical processes that are tested that examine variable influence, model bias, model significance and model performance. The use of the GWR outputs is also explained in the method section in research study 7.

Two approaches are proposed for examining spatially varying relationships – a hypothesis testing approach and an exploratory approach. The hypothesis testing approach requires the researcher to consider, based on sound theoretical grounds, the reason why a phenomenon may be present. For example, violent crime hotspots may cluster in a city centre and around entertainment facilities. From this, it would be plausible to suggest that the reason why these violent crime hotspots are present are because these are the areas where bars, pubs and nightclubs are located. The logic is that the high concentration of violent crime in these hotspots is related to the location of premises where people gather to consume alcohol. It would also be plausible to suggest that the violent crime hotspots are related to altercations between pupils from rival schools, with these altercations occurring where children from the schools come together at the main transport hubs whilst returning home from school. These are two plausible hypotheses that are based on sound logical and environmental criminology theoretical grounds. These two explanatory variables (the location of bars, pubs and nightclubs; and transport hubs) can then be included in the regression model to explore, statistically, if they are related to violent crime, and if these relationships vary geographically. The hypothesis driven approach, therefore, involves determining a plausible explanation at the outset, determining the data that can be used to test the hypothesis, applying a spatial regression model, and then using the results from the model to identify if there is evidence that supports this theoretical reasoning.

The exploratory approach involves choosing a range of variables that are then tested in the regression model, and where correlations are found, attempting to interpret the results based on plausible logic or some theoretical basis. That is, it involves fitting a theory to the results to explain the relationship between the dependent and each significant explanatory variable. The exploratory approach can also be used when the researcher is not initially confident on the theoretical grounds to determine plausible hypotheses, or is limited with their access to data. For example, data may not exist that is perfectly associated with a hypothesis, requiring the researcher to instead rely on a proxy measure for exploring if a relationship exists.

The current research explores both approaches, seeking to identify strengths and weaknesses that lead to identifying and validating the inclusion of explanatory variables in a hotspot analysis process. To facilitate this part of the research, a number of datasets for small areas were sourced. These are described in research study 7.

Recall that the main objective in examining the use of GWR is to help identify why hotspots exist and for informing spatial predictions of crime. GWR will, therefore, be tested to examine if it is suitable for helping to understand local variations in why hotspots exist. Hotspots that are identified are likely to be small in area; therefore, GWR will be tested to see if the results it generates are geographically precise enough to facilitate this type of hotspot interpretation. If these results prove promising, a hotspot modelling process will be explored that will aim to identify a practical means of combining hotspot mapping outputs with explanatory variables to generate a map that is more effective (than hotspot mapping methods) for predicting spatial patterns of crime. At present, it is anticipated that a raster summation (map algebra) process (de Smith et al., 2007) will be used. This would involve combining the results from a hotspot mapping process (i.e., the presence, or not, of the cell forming part of a crime hotspot) with results from a spatial regression model (i.e., the presence, or not, of an explanatory variable in each cell). If appropriate, the research will also consider the practicality of assigning spatially varying weights to each of these variables.

4.5.6. Summary of spatial analysis processes

The first research study will identify the point in the temporal data input periods at which clustering is evident, for each of the crime types, and for both of the study areas. From this, assuming that there is some point at which clustering is evident, the common hotspot analysis techniques (research study 2) will be tested for their prediction accuracy. Study 2 will involve some testing of each technique's parameter settings that are used for generating hotspot mapping output, however, these parameters will be examined in more detail for one of the common hotspot mapping techniques in study 3. In study 3, only the hotspot mapping technique that was identified to consistently produce spatial crime predictions that were better than the others would be examined to determine the influence its parameter values have on the prediction accuracy of the mapping output it generates.

Study 4 involves examining whether spatial significance mapping removes the ambiguity of defining the areal coverage of hotspot areas, and if the areas that are statistically defined as *hot* perform better in predicting where crime occurs in comparison to the common hotspot analysis techniques. Study 5 analyses whether hotspots that are identified are temporally stable so that practitioners can be confident that these are the areas where crime is likely to continue to persist in the future, and therefore any resource targeting should have some impact. These temporal features of hotspots are then examined further in study 6 when the principles of repeat and near repeat patterns are tested for producing accurate predictions. An examination of these temporal features will involve using a prospective mapping tool to determine if good predictions are not only generated for the near future but that these predictions are also valuable for predicting the longer-term future, or whether hotspot analysis techniques perform better at predicting where crime is likely to occur for certain periods of the future.

Study 7 involves using GWR to help interpret why hotspots exist. This will involve examining whether GWR is sensitive enough for exploring relationships between explanatory variables and spatially varying levels of crime at the geographic scale at which hotspots are identified. Consideration will need to be given to the type of GWR model to apply (Gaussian or Poisson) alongside a full assessment of the diagnostic statistical tests that are required for GWR modelling. The different approaches for applying GWR (or any other form of regression analysis) will also be examined for their merits – hypothesis testing and exploratory analysis. If the GWR results prove promising, the technique will then be examined to identify if it can support an improvement in spatial crime predictions by constructing a modelling process that includes explanatory variables alongside, or in replacement of, retrospective data on crime. In addition, the potential use of GWR mapping outputs will also be examined for their value in informing how crime levels may change based on the change in one or a number of related variables.

5. Research study 1: Testing for the presence of crime hotspots

5.1. Introduction

A preliminary stage in any form of hotspot analysis is to determine if there is evidence that hotspots exist in the data under examination. Hotspots are spatial clusters of events, for which a number of tests exist for examining the statistical presence of clusters. The most common statistical tests for identifying clusters in spatial data are Moran's I, Geary's C and the Nearest Neighbour Index. Research study 1 tests the hypothesis (hypothesis 1) that hotspots can be identified using retrospective data for a short period of time rather than requiring retrospective data for longer periods of time. In so doing, the research will identify the position in retrospective crime data when crime hotspot patterns first appear and if this statistical evidence for clustering continues across the range of input data periods.

5.2. Chapter aims and structure

Examined in this chapter is the use of spatial statistical tests for determining if hotspots exist in the crime data that are examined. This includes determining if the presence of hotspots (clusters of crimes) is statistically significant. The following method section describes the technical features of Moran's I, Geary's C and the Nearest Neighbour Index, and their suitability for testing for the presence of hotspots. The experiments that were conducted are then described, including a description of the temporal periods that were used for the input data. In these experiments, input data were structured differently to that described in the general method chapter (chapter 4) in order to suit the requirements of these statistical tests.

Section 5.4 presents the results from the statistical clustering experiments. These are then summarised to inform how these results influence subsequent studies in this research.

5.3. Method

5.3.1. Statistical tests for clustering

This research study examines the use of Moran's I, Geary's C and the Nearest Neighbour Index for testing for statistical evidence of clustering (i.e., hotspots) in the crime data under examination. The following sections will describe each of these statistical measures in turn to determine their suitability for testing for the presence of clustering in crime data.

I. Moran’s I

The Moran’s I statistic is a spatial autocorrelation test that compares the similarity in values between each location and its near neighbours (Eck et al., 2005; Levine, 2010; O’Sullivan and Unwin, 2003). Moran’s I requires an intensity value, and the position (i.e., geographic coordinates) of each event that is under study in order to be calculated. In practice, this requires crime data to be aggregated to some form of geographic unit where its position is represented by the centroid of the geographic unit (e.g., a grid cell, or output area). This centroid point is then assigned an intensity value that represents the number of crimes in the geographic unit (Eck et al., 2005).

The Moran’s I result varies between –1.0 and +1.0. Where values in a geographic unit are high and are surrounded by other geographic units with similarly high values, positive spatial autocorrelation exists; that is, a spatial cluster of high crime is identified. When areas with low values are surrounded by high crime areas and high crime areas are surrounded by geographic units with low levels of crime, the series would display negative spatial autocorrelation; that is, a spatial cluster of high crime areas is not identified. For these reasons, Moran’s I has been used by practitioners to identify if hotspots are present in spatially distributed crime data (Eck et al., 2005; Levine, 2010). The significance of the Moran’s I result can be tested against a theoretical distribution (one that is normally distributed) by dividing by its theoretical standard deviation (Levine, 2010). This determines, for example, if the clustering of the data (indicated by positive spatial autocorrelation) is statistically significant to the 95th percentile, the 99th percentile or the 99.9th percentile. However, as well as positive spatial autocorrelation indicating the presence of clusters of crime, positive spatial autocorrelation also indicates areas where low values are surrounded by low values. That is, evidence of positive spatial autocorrelation using Moran’s I relates to areas that are surrounded by similar values, rather than exclusively relating to areas where high crime levels cluster.

II. Geary’s C

Geary’s C statistic is a spatial autocorrelation measure of the deviations in intensity values of each point with one another. Similar to Moran’s I, it requires data to be aggregated to some form of geographic unit (e.g., a grid cell, or output area) to determine an intensity value. The location of the geographic unit is represented by the units’ centroid geographic coordinates. The values of C typically vary between 0 and 2, where values less than 1 indicate evidence of positive spatial autocorrelation and values greater than 1 indicate evidence of negative spatial autocorrelation. The Geary's C statistic can be tested for significance against a theoretical distribution (one that is normally distributed) by dividing by its theoretical standard deviation (Levine, 2010). Similar to Moran's I, this can be used to determine the level of statistical significance in the Geary's C result. Also, similar to Moran's I, evidence of positive spatial autocorrelation in data that are examined using Geary's C relates to areas that are surrounded by similar values, rather than exclusively referring to areas where high (crime) levels cluster.

III. Nearest Neighbour Index

The Nearest Neighbour Index (NNI) test compares the observed point-based spatial distribution of the data under examination against random variation (Clark and Evans, 1954). That is, it does not require aggregation to a geographic unit to be performed prior to applying the test. If the result generated from the NNI test is 1 then the data are randomly distributed. If the NNI result is less than 1 then the data show evidence of clustering. A NNI result that is greater than 1 indicates evidence of a uniform pattern in the data. The statistical significance of the result can be derived (at least approximately) by comparing the theoretical distribution of the nearest neighbour distances between points under complete spatial randomness, with the observed nearest neighbour distances (Bailey and Gatrell, 1995).

It was decided that the NNI would provide the better test for clustering of crime data. This was for two reasons:

- Positive spatial autocorrelation under Moran's I and Geary's C requires further investigation to identify if the clustering identified refers to areas of high crime surrounded by high crime, rather than areas of low crime surrounded by low crime
- The data used in this research were point-based. Therefore, rather than losing some spatial detail in the original data by aggregating the point data to geographic units in order to perform Moran's I or Geary's C, a test of spatial clustering using the NNI preserved a point-based analytical approach.

Two sets of experiments were carried out. The first set was based on input data for a set of temporal periods prior to the measurement date. Typically, practitioners would use data for a certain number of previous days or weeks for producing a hotspot map. The first set of experiments replicated this process. The second approach used input data for a set of uniform categories of n events prior to the measurement date (e.g., multiples of 5 events). This approach would better determine the typical number of events that were required for generating hotspot maps. The tests were repeated for each crime type for the two study areas to examine consistencies in the findings.

5.3.2. Input data based on a set of temporal periods prior to the measurement date

Measurement dates for the Camden/Islington and Newcastle study areas were selected as the 1st January 2010 and the 1st April 2010 respectively. Input data for both the study areas and for each crime type (burglary dwelling, theft from the person, theft from vehicle, theft of vehicle) were prepared into the temporal periods shown in Table 5.1. For example, one day of burglary dwelling data (i.e., the day prior to the measurement date) was selected and the NNI test was applied. The result of the test was noted, including whether it was statistically significant to 95%. This approach meant that the typical temporal period prior to the measurement date for which the crime data showed statistical evidence of clustering could be determined. This resulted in 112 tests (14 temporal periods x 4 crime types x 2 study areas). An additional set of fourteen NNI tests were performed on the Newcastle assaults with injury data. Where there were no crime events or only one crime event for any temporal input period, the nearest neighbour statistic could not be applied.

Table 5.1. Temporal periods of input data used for NII tests, for both study areas

Temporal input data periods	1 day	2 days	3 days	4 days	5 days	6 days	1 week
	2 weeks	3 weeks	4 weeks	5 weeks	6 weeks	7 weeks	8 weeks

5.3.3. Input data based on a set of uniform categories of n events prior to the measurement date

NNI tests were also performed on input data for each of the study areas and for each crime type for a set of uniform categories of n events. Table 5.2 lists these categories of input data. For example, the ten burglary dwelling events that were committed immediately prior to the study area's measurement date were selected and the NNI test was applied. The result of the test was noted, including whether it was statistically significant to 95% confidence level. Using these results, an additional set of NNI tests was carried out to determine the exact number of events in each data set when clustering was statistically significant. This resulted in 184 tests (23 temporal periods x 4 crime types x 2 study areas), plus up to 64 further tests (up to 8 NNI tests to establish the exact point of clustering x 4 crime types x 2 study areas). An additional set of 23 NNI tests were also performed on the Newcastle assaults with injury data, including further tests to determine the exact volume of data when clustering was significant to 95%.

Table 5.2. Uniform categories of input data used for NNI tests

Uniform input data periods												
	5	10	15	20	25	30	35	40	45	50	55	60
	65	70	75	80	90	100	110	120	130	140	150	

5.4. Results

The first set of experiments was to determine if clusters (i.e., hotspots) were evident in the sample data, and how much data were required before clustering was statistically significant to 95%. The experiments were performed using two approaches:

- Using input data for a set of temporal periods prior to the measurement date (i.e., 1 day, 2 days, 3 days ... 8 weeks). This approach would determine the typical temporal period of input data that was required to generate hotspot maps.
- Using input data for a set of uniform categories of n events prior to the measurement date (i.e., 5, 10, 15 ... 150). This approach would determine the typical number of events that were required for generating hotspot maps.

Figure 5.1 and Table 5.3 show that for Camden/Islington only 3 days of data on thefts from vehicles were required before clustering was evident. Burglary dwelling and theft from the person showed evidence of clustering using just 4 and 5 days of data respectively, whereas theft of vehicles required 1 week and 3 days of data before clustering was evident. Before these respective thresholds were reached, the crime data displayed patterns of random variation or were significantly uniformly distributed ($p = 0.05$). That is, 1 week and 3 days of theft from vehicles data would typically be required in order to be able to generate a meaningful map identifying where hotspots have occurred. Figure 5.1 and Table 5.3 show the exact number of crime events that were required in Camden/Islington before clustering was significant to 95%. This ranged from 45 offences for burglary dwelling and 66 offences for theft from vehicles. All data input periods prior to the point at which clustering was found to be statistically significant also showed evidence of clustering. For example, all longer data input periods for theft from the person, prior to 5 days before the measurement date (representing 52 offences) showed significant evidence of clustering.

Table 5.3. The amount of crime data, by volume and retrospective temporal period, required in Camden/Islington and Newcastle for clustering to be statistically significant to 95%

	Camden/Islington		**Newcastle**	
	Days/weeks of data for evidence of clustering	**n of cases for evidence of clustering**	**Days/weeks of data for evidence of clustering**	**n of cases for evidence of clustering**
Burglary dwelling	4 days	45	2 weeks 4 days	65
Theft from the person	5 days	52	2 weeks 5 days	34
Theft of vehicle	1 week 3 days	64	16 weeks 5 days	63
Theft from vehicle	3 days	66	1 week 1 day	37
Assaults with injury	-	-	1 week 5 days	51

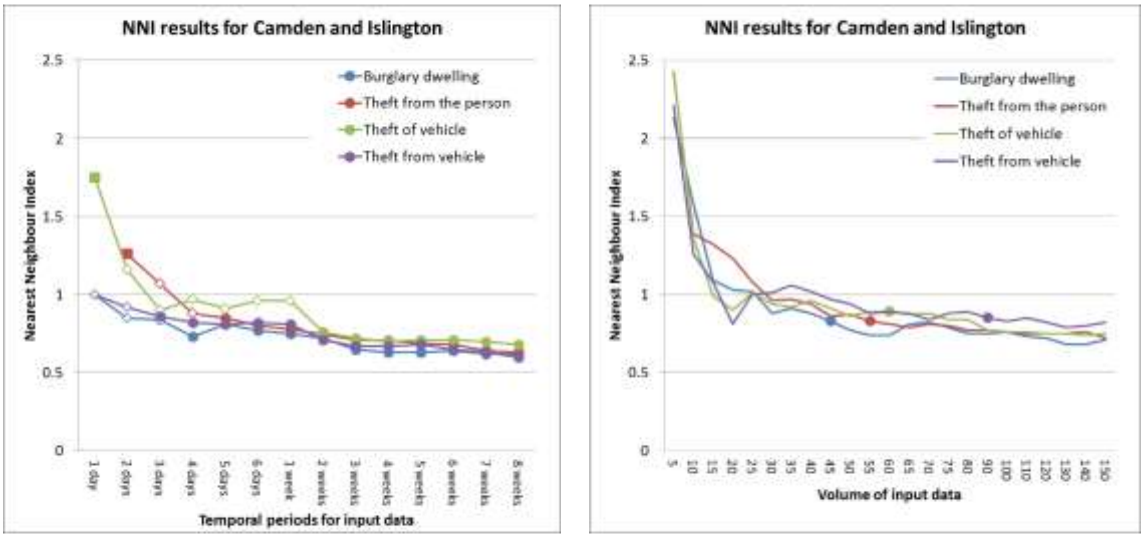

(a) (b)

Figure 5.1. NNI results for Camden/Islington burglary dwelling, theft from the person, theft of vehicles and theft from vehicles for (a) a set of temporal periods of input data, and (b) the volume of input data. Markers in (a) that are shaded represent a result that was significant to 95%; circle markers represent clustering, diamond markers refer to random spatial distribution and square markers represent a uniform spatial distribution. Circular markers in (b) represent the point when clustering was determined to be

significant to 95%. Before this point, clustering was not significant. After this point, all sets of data were significantly clustered.

In Newcastle, 8 days to 16 weeks and 5 days of data were required before clustering was evident across all crime types (see Table 5.3). Similar to Camden/Islington, theft from vehicles required the fewest number of days of recorded crime data before clustering was statistically significant (8 days), followed by burglary dwelling (2 weeks 4 days) and theft from the person (2 weeks 5 days). One week and 5 days of assaults with injury data were required for clustering to be evident. Over 16 weeks of theft of vehicles data were required for clustering to be evident. Before these temporal thresholds, the crime data displayed patterns of random variation or were significantly uniformly distributed (p = 0.05) in most cases.

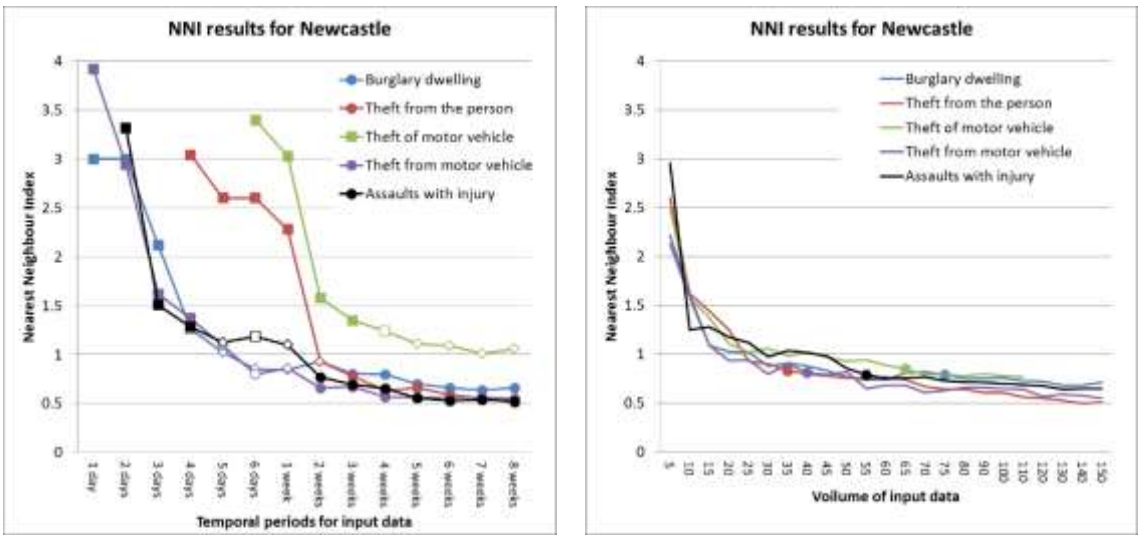

(a) (b)

Figure 5.2. NNI results for Newcastle burglary dwelling, theft from the person, theft of vehicles, theft from vehicles and assaults with injury for (a) a set of temporal periods of input data, and (b) the volume of input data. Markers in (a) that are shaded represent a result that was significant to 95%; circle markers represent clustering, diamond markers refer to random spatial distribution and square markers represent a uniform spatial distribution. Circular markers in (b) represent the point when clustering was determined to be significant to 95%. Before this point, clustering was not significant. After this point, all sets of data were significantly clustered. The theft from vehicles line stops at 110 because there were only 110 records in the full six months of input data for Newcastle.

The results from using the NNI to test for clustering are illustrative of the different levels of crime in Camden/Islington and Newcastle. Figure 5.2 and Table 5.4 show the exact number of crime events that were required in Newcastle before clustering was significant to 95%. This ranged from 34 offences for theft from the person and 65 offences for burglary dwelling. All data input periods prior to the point at which clustering was found to be significant also showed evidence of clustering. For example, all longer data input periods for burglary dwelling, prior to 2 weeks and 4 days before the measurement date (representing 65 offences) showed significant evidence of clustering.

5.5. Interpretation and conclusions from research study 1

In relation to the research hypothesis (hypothesis 1) - *hotspots can be identified using retrospective data for a short period of time rather than requiring retrospective data for longer periods of time* - the results show differences between study areas and between crime types in the volume of data that are required before hotspots are statistically evident and hotspot mapping can be performed. For example, while 1 week of burglary dwelling, theft from the person data and theft from vehicle data for Camden/Islington would be sufficient for performing some level of hotspot analysis, mapping 1 week of crime data for any of the crime types examined from Newcastle would just show random spatial patterns of these crime types. In all cases for the Newcastle study area, over one week of crime data were required for each crime type for clustering to be evident. In the case of theft of vehicles in Newcastle, over 16 weeks of crime data were required for hotspots to be statistically evident. Analysis of the number of crime events that were required showed that clustering became evident when 34 theft from the person offences in Newcastle were tested, while 66 theft from vehicle offences in Camden/Islington were required before clustering was statistically determined to be evident. These results show that, although in some cases short retrospective periods of crime data (of less than a week) or small volumes of crime (less than 35 crimes) are sufficient for hotspots to be statistically identified, the specific retrospective point when hotspots are evident varies. Once clustering was detected in the crime point data, all other input data for longer periods showed statistical evidence of clustering.

In practitioner terms, simply choosing a retrospective period, whether it be based on a retrospective number of days or retrospective volume of crime, and expecting hotspots to be present is not sufficient if the analyst then expects hotspots to appear on a map. The results from this research study illustrate the value in performing the NNI test as a preliminary stage to hotspot mapping to ensure hotspots are present in the data that are examined, particularly when only small volumes of data are available. Once hotspots have been identified using the NNI, the analyst can be confident that any mapping of these data will identify where these hotspots are present, rather than interpreting patterns that may look like clusters but are based on random spatial variation. A number of techniques exist for mapping hotspots. These are examined in the next research study.

6. Research study 2: A metric assessment of the prediction performance of common hotspot analysis techniques

6.1. Introduction

Research study 1 showed that the specific retrospective point when hotspots are evident in crime data varies, and that the Nearest Neighbour Index is a simple way of testing for the presence of hotspots. Once hotspots are statistically identified in crime data, a number of techniques are available for mapping crime data to determine where the hotspots exist. The mapping of standard deviation spatial ellipses, thematic mapping of geographic units and kernel density estimation (KDE) have become popular techniques for mapping hotspots of crime. Research study 2 tests the hypothesis (hypothesis 2) that common hotspot mapping techniques differ in how accurately they predict spatial patterns of crime.

The mapping of hotspots using spatial ellipses, thematic mapping of geographic units and kernel density estimation (KDE) have been subject to several reviews (see Chainey et al., 2002; Eck et al., 2005; Jefferis, 1999). However, these reviews have been little more than visual comparisons of each method or exercises that have evaluated their ease of use. Importantly, these reviews demonstrated that different hotspot mapping techniques produce different results in terms of identifying the location, size and shape of areas that are defined as hotspots. None of these reviews, however, have determined which of these techniques is best for helping to identify where spatial patterns of crime may occur in the future. The research described in this chapter assesses the predictive accuracy of the common hotspot mapping methods.

6.2. Chapter aims and structure

Examined in this chapter is the predictive accuracy of common hotspot mapping techniques - spatial ellipses, thematic mapping of geographic units, thematic mapping of grid cells and kernel density estimation. The research identifies if differences exist in the spatial prediction performance of these techniques and if one technique consistently provides the best predictions on where crime occurs in the future. This will involve conducting a number of experiments using each hotspot analysis technique, for a range of crime types, and a range of input data periods. These hotspot maps are then compared against where crime has occurred for a range of data output periods.

The approach that was adopted for the experiments is described in the method section (6.3). The method section includes a description of the data used, descriptions of the technical features about each common hotspot analysis technique, the method used for determining values that represent *hot* from each technique's hotspot analysis output and the metrics that were used for measuring prediction performance.

Section 6.4 presents the results from the hotspot analyses experiments. These results are then summarised, drawing from study 1 of the research to help interpret the results, and to inform how these results influence subsequent studies in this research.

6.3. Method

The experiments were organised in two parts. The first part involved an analysis of four commonly used mapping techniques for generating hotspot maps using data for the Camden/Islington study area. The techniques were spatial ellipses, thematic mapping of boundary areas, grid thematic mapping and kernel density estimation. These techniques were chosen as they are the most commonly used by those generating hotspot maps of crime (Weir and Bangs, 2007). The second part involved a replication of the experiments using Newcastle data to examine consistencies in findings with the Camden/Islington experiments. The experiments for Newcastle included using assault with injury data as well as the crime types that were used in the Camden/Islington study area.

6.3.1. Crime data

The analysis in this research study used the two measurement dates for Camden/Islington described in chapter 4 (1^{st} January 2010 and 11^{th} March 2010) and the Newcastle measurement date of the 1^{st} April 2010. The data input periods and data output periods used were those described in the study areas and crime data section (4.3) in chapter 4.

6.3.2. Hotspot analysis techniques

Each of the hotspot analysis techniques requires the user to input certain parameters to generate mapping output. The parameters that were set are described below for each technique, with the method that was followed being either the procedure advised in the supporting guidance provided with the software for that technique or the default options for the technique. This method was followed to replicate the approach a police analyst would most typically apply.

I. STAC spatial ellipses

CrimeStat (Levine, 2010) was used for generating STAC standard deviational spatial ellipses. The parameters that the user is required to enter are the search radius, the minimum number of points-per-hotspot, the number of standard deviational ellipses used to delineate the hotspots, the number of standard deviations to apply for the generation of spatial ellipses, and the scan type. The triangular scan type was selected due to the irregular road network within each study area (as advised by Levine, 2010). Spatial ellipses were created to one standard deviation. Setting the minimum points-per-hotspot and the number of ellipses proved to be more complex as it is difficult to universally apply a fixed number of points-per-hotspot with the different periods of input data using STAC. In the absence of any official direction, a variable minimum points-per-hotspot approach was applied, in which the minimum number of points-per-hotspot producing the highest number of hotspots under 20 was calculated using a trial and error method. This approach generates ellipses (i.e., hotspots) that reflect the underlying crime events and eliminates those ellipses that could be created to make up the numbers to reach 20.

The only guidance for applying a suitable search radius was experimentation and experience (Levine, 2010). As a result, three search radii were used to determine hotspots – 500m, 250m and a search radius equal to the default bandwidth that was to be applied to the creation of KDE maps. The rationale for choosing the default bandwidth value is explained below in the KDE section. This KDE bandwidth default derived parameter determined an alternate search radius size that varied in accordance with the spatial characteristics of the crime input data. The three approaches were called Spatial ellipses 500, Spatial ellipses 250 and Spatial ellipses HSD. Once these parameters had been set and run in CrimeStat, the outputs were imported into MapInfo for display and analysis of each output's predictive accuracy.

II. Thematic mapping of administrative geographic units

The thematic mapping of administrative geographic units technique requires crime point data to be aggregated to some commonly used geographic unit, such as police beats, wards or census geography. Census output areas are the smallest unit of Census geography in England and Wales, each covering approximately 125 households and are commonly used for aggregating crime data (Chainey and Ratcliffe, 2005; Weir and Bangs, 2007). A count of crime was generated for each output area, for each crime type. This count of crime was performed using standard point in polygon aggregation routines in MapInfo. The selection of the thematic range method and values for determining the crime hotspot threshold are explained below as they also apply to grid thematic mapping and KDE.

III. Grid thematic mapping

To perform grid thematic mapping in MapInfo (and most other GIS software) the analyst is first required to draw a grid lattice that can be positioned across the study area. The parameter the user must decide is the size of each square grid cell. There is little guidance on which cell size to select; if the cells are too large, the resulting hotspot map will only show coarse geographic patterns of crime, and if too small it can be difficult to discern any spatial patterns from the hotspot map (Chainey and Ratcliffe, 2005). Experimentation and experience are again the best advice, but for novices to hotspot mapping using the grid thematic mapping technique, Chainey and Ratcliffe (2005) suggest a useful starting point for grid cell size is to calculate the distance in the shortest extent of the study area, and divide this distance by 50. Following the approach advised by Chainey and Ratcliffe (2005) a grid cell size of 250 m was chosen for Camden/Islington and 320 m for Newcastle. In addition, the calculations for the grid thematic mapping technique were repeated using the bandwidth size determined from the KDE bandwidth default value (the applicability of this technique is discussed in the KDE section). This determined an alternate grid cell size that varied in accordance with the spatial characteristics of the crime input data. However, this cell size measure for Newcastle was calculated to be the same as the 320 m measure determined from the first grid cell size calculation method, and therefore a second set of experiments on grid cell size for the Newcastle study area was not required.

Once the grids had been calibrated, functions in MapInfo were used to calculate a count of the number of crime points within each grid cell. The selection of the thematic range method and values for determining the crime hotspot threshold are explained in detail below.

IV. Kernel density estimation (KDE)

The spatial application of kernel density estimation emerged as a popular technique in spatial epidemiology to assist the study of disease patterns[10]. Similar to disease, crime

[10] For an early example of the application of KDE, see Bithell, 1990

incidents are most usually geographically referenced as points. The KDE function is applied to these points to obtain a smooth surface estimate representing the density of the point distribution. In mathematical terms, KDE is expressed as:

$$f(x,y) = \frac{1}{nh^2} \sum_{i=1}^{n} k\left(\frac{d_i}{h}\right) \qquad (4)$$

where f(x,y) is the density value at location *(x,y)*, *n* is the number of incidents/points, *h* is the bandwidth, d_i is the geographical distance between incident *i* and location *(x, y)* and *k* is a density function, known as the kernel. *k* can take many forms although the results between different functions produce very similar density values (Bailey and Gatrell, 1995). A common choice for *k* is the quartic function (Bailey and Gatrell, 1995; Chainey et al., 2002; Chainey and Ratcliffe, 2005; Levine, 2010; Ratcliffe, 2004; Williamson et al., 1999).

KDE hotspot maps were generated using Hotspot Detective for MapInfo (Ratcliffe, 2004). The user is required to enter two parameters - the cell size and the bandwidth size. Following the default settings is an approach that most analysts take, and indeed are encouraged to take if they are not experts in spatial analysis.

These first experiments did not examine different cell sizes and bandwidth sizes. Experimentation of these parameters would follow if KDE was identified as one of the better common hotspot mapping techniques in this first round of experiments. Hotspot Detective determines default settings for these parameters after performing an analysis of the input data. KDE cell size is calculated by dividing the shorter side of the minimum bounding rectangle around the study area by 150 (Ratcliffe, 2004). Bandwidth selection is more complicated (Chainey and Ratcliffe, 2005). In Hotspot Detective, the calculation of the default bandwidth value is not divulged to users, but it is known to be a function of the shorter side of the minimum bounding rectangle surrounding the study area, divided by a number that provides a suitable enough cell resolution without requiring a significant number of iterations to generate a representative KDE surface (Ratcliffe – personal communication). Experience with using these defaults suggests they are appropriate in most cases for determining cell size and bandwidth settings applied to crime data. As the approach used by Hotspot Detective considers the spatial characteristics of the study area it was considered useful to use the KDE bandwidth default measure determined using

Hotspot Detective as one of the search radii parameters for STAC and as one of the cell sizes for the grid thematic mapping technique. Table 6.1 lists the parameter details that were used for each of the common hotspot analysis techniques.

Another parameter that users of the KDE function in Hotspot Detective are invited to enter is a weighting attribute. As each retrospective crime event was applied with equal weight, no weighting scheme was applied to the input data.

Table 6.1. Common hotspot mapping techniques and the parameter values that were used for producing hotspot maps

Study area	Common hotspot mapping techniques and the parameters used in each study area
Camden/Islington	• STAC: search radius 250 m • STAC: search radius 500 m • STAC: 225m (Hotspot Detective KDE bandwidth) • Thematic mapping of Census output areas (n = 1392) • Thematic mapping of grid cells: cell size 250 m • Thematic mapping of grid cells: cell size 225 m (Hotspot Detective KDE bandwidth) • Kernel density estimation: cell size 45 m, bandwidth 225 m
Newcastle	• STAC: search radius 250 m • STAC: search radius 500 m • STAC: 320 m (Hotspot Detective KDE bandwidth) • Thematic mapping of Census output areas (n = 888) • Thematic mapping of grid cells: cell size 320 m (Hotspot Detective KDE bandwidth) • Kernel density estimation: cell size 64 m, bandwidth 320 m

6.3.3. Determining a thematic threshold for hotspots

A final parameter to consider in hotspot map generation is a threshold value for determining which areas are *hot*. For spatial ellipses this is straightforward as it is simply the area drawn by each ellipse. Thematic mapping of output areas, grid thematic mapping and KDE produce areal values across a continuous range (e.g., for grid thematic mapping, each grid cell has a value representing the number of crimes located within the cell). A threshold value must be determined that specifies that any values above this threshold can be classified as an area that is a hotspot.

The same thematic range approach was applied to each of the three hotspot mapping techniques (thematic mapping of output areas, grid thematic mapping and KDE) for the purpose of simplicity and consistency in methodology. Five thematic classes were used and default values generated using the quantile method in MapInfo were applied to the thematic mapping of output areas, grid thematic mapping and KDE for determining thematic classes. This approach for thematic classification was used because the number of classes falls within the upper and lower settings specified by Dent (1999) and Harries (1999), and the quantile method was chosen because it distributes the data in an approximately equal balance between the classes, resulting in a visually balanced map pattern (Monmonier 1996). This approach for determining thematic classes is also a common approach that many practitioners apply for generating hotspot maps (Eck et al., 2005).

Hot was then determined by the top thematic class. Figure 6.1 illustrates an example of this approach - it shows a hotspot map generated using KDE where the thematic ranges were grouped into five classes and arranged following the quantile range method default in MapInfo. Cells with values in only the top thematic class were then selected, with these areas determined as the hotspots.

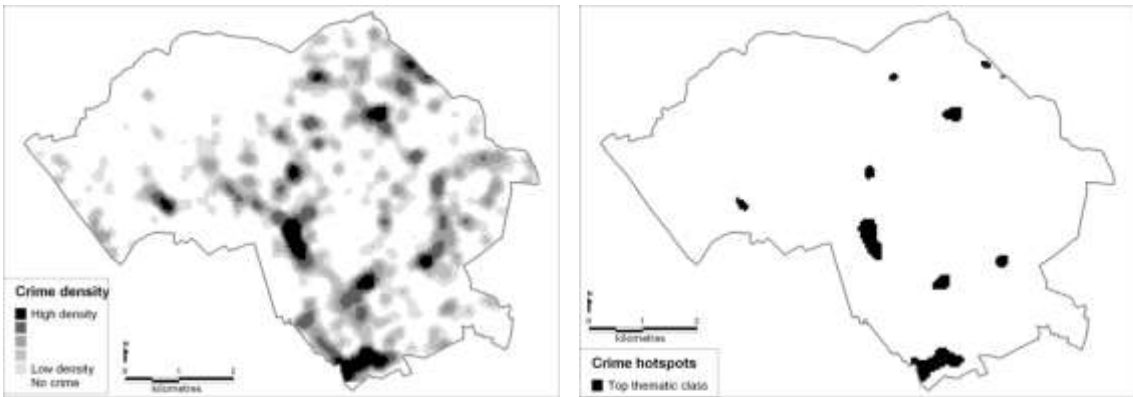

(a) Hotspot map (b) Top thematic class of hotspot map

Figure 6.1. Hotspots were determined by selecting the uppermost thematic class calculated using five classes and the default values generated from a quantile thematic range method

6.3.4. Measuring the prediction performance of common hotspot mapping output

This research study used the PAI to compare the prediction performance of the common hotspot analysis techniques, rather than using accuracy concentration curves, area under the curve and the crime prediction index. This was due to the multiple data input periods and data output periods that were used in the experiments and the large number of experiments this, therefore, required. In Johnson et al.'s (2008b) study that compared results between mapping techniques using an accuracy concentration curve, the analysis was only conducted for one input data period (two months) and one output data period (seven days). Calculating an accuracy concentration curve is practical for comparing one set of data input and output for two different techniques (i.e., two experiments). The current research used ten different input data sets, and ten different output data sets for seven different types of hotspot analysis, and for two study areas, therefore, involving 1400 experiments. The use of accuracy concentration curves, area under the curve and the CPI for each experiment was, therefore, considered neither practical nor proportionate to the aims of this particular research study (i.e., to compare the spatial crime prediction performance of the common hotspot mapping techniques). The use of the PAI has since been discussed further by Pezzuchi (2008), Levine (2008) and Chainey et al. (2008b; 2008c), with researchers concluding it to be a useful measure for comparing multiple hotspot mapping outputs.

The PAI was, therefore, used to compare the spatial crime prediction performance of the different common hotspot mapping techniques. Hotspot maps of different crime types are typically regarded as being similar in their accuracy for predicting crime patterns. As the experiments were also conducted across a range of crime types, the PAI was used to identify if there were differences between the common hotspot mapping techniques for the different types of crime.

6.4. Results

This research study examined each of the common hotspot mapping techniques to determine if differences exist in their ability to predict future patterns of crime. Each of the techniques was applied using data from Camden/Islington, and Newcastle[11].

[11] The method and majority of the results from this research study have already been published in the Security Journal: Chainey, S.P., Tompson, L., and Uhlig, S. (2008), "The utility of hotspot mapping for predicting spatial patterns of crime", *Security Journal* 21:1-2. This PhD research, however, uses a more up-to-date version of the Camden/Islington data but the results are consistent with the published paper. The published paper did not include the results for Newcastle.

Table 6.2. Prediction Accuracy Index values for different hotspot mapping techniques. Results in bold relate to the highest PAI values and results in italics relate to the lowest PAI values. Results are presented for each of the dates when hotspot maps were generated for Camden/Islington, and the single date for Newcastle.

Hotspot mapping technique	**Camden and Islington**		**Newcastle**
	Average PAI (01/01/2010)	**Average PAI (11/03/2010)**	**Average PAI (01/04/2010)**
Spatial ellipses 250m	1.74	2.25	10.6
Spatial ellipses 500m	*1.24*	*1.52*	*8.3*
Spatial ellipses HSD	1.69	2.03	9.4
Thematic mapping of Output Areas	1.91	2.38	16.4
Thematic mapping of grids	2.00	2.34	-
Thematic mapping of grids HSD	2.06	2.63	34.5
Kernel density estimation	**2.90**	**3.41**	**41.5**

Table 6.2 shows the PAI results for the different hotspot mapping techniques. These results are presented in three columns to show differences between the two measurement dates in Camden/Islington and to compare against the results for Newcastle. Average PAI values were calculated from the individual PAI values from hotspot maps for the different input periods, the different output periods and for the different crime types. These results show that there were differences between hotspot mapping techniques in their ability to predict patterns of crime. KDE consistently proved to be the best hotspot mapping technique for predicting where crimes may occur in the future and spatial ellipses were the worst. The PAI results for Newcastle were also much higher than those for Camden/Islington suggesting that the hotspot maps produced for Newcastle were more accurate than those produced for Camden/Islington. This is likely to be due to the more intense clustering of crime in Newcastle in comparison to Camden/Islington. However, the differences in the PAI values between the two areas do not necessarily mean that the prediction performance of the hotspot maps produced for Newcastle are more accurate to the magnitude indicated by the differences in PAI values. The PAI is an indicative measure of prediction performance, meaning that a map with a PAI of 10 is not necessarily five times better than a hotspot map that has a PAI of 2. Differences in the number of crimes that were predicted using hotspot maps for the two study areas are examined further in section 6.3.5.

Table 6.3. PAI values for different hotspot mapping techniques, by crime type. (a) Camden/Islington, calculated from the 1st January 2010 measurement date, (b) Camden/Islington, calculated from the 11th March 2010 measurement date, and (c) Newcastle, calculated from the 1st April 2010 measurement date. Results in bold relate to the highest PAI values and results in italics relate to the lowest PAI values.

(a) Camden/Islington: 1st January 2010 measurement date

Hotspot mapping technique	Burglary dwelling	Theft from the person	Theft from vehicle	Theft of vehicle
Spatial ellipses 250m	1.4	2.4	2.2	1.7
Spatial ellipses 500m	1.3	*1.5*	1.5	*0.8*
Spatial ellipses HSD	1.4	2.5	2.1	1.3
Thematic mapping of output areas	*1.1*	4.2	*1.2*	1.2
Thematic mapping of grids 250m	1.7	4.0	1.8	1.4
Thematic mapping of grids HSD	1.7	3.5	2.1	2.1
Kernel density estimation	**2.3**	**4.7**	**2.3**	**2.3**

(b) Camden/Islington: 11th March 2010 measurement date

Hotspot mapping technique	Burglary dwelling	Theft from the person	Theft from vehicle	Theft of vehicle
Spatial ellipses 250m	1.32	2.59	2.15	2.93
Spatial ellipses 500m	1.31	*1.40*	*1.55*	1.82
Spatial ellipses HSD	1.29	2.63	2.63	*1.59*
Thematic mapping of output areas	*1.25*	3.32	2.93	2.01
Thematic mapping of grids 250m	1.67	3.58	2.43	1.66
Thematic mapping of grids HSD	1.95	4.14	2.55	1.89
Kernel density estimation	**2.33**	**4.59**	**3.66**	**3.05**

(c) Newcastle: 1st April 2010 measurement date

Hotspot mapping technique	Burglary dwelling	Theft from the person	Theft from vehicle	Theft of vehicle	Assault with injury
Spatial ellipses 250m	5.1	23.2	4.3	0.3	19.9
Spatial ellipses 500m	*4.3*	*17.1*	*3.9*	*0.2*	*16.1*
Spatial ellipses HSD	4.7	19	4.1	0.3	19.1
Thematic mapping of output areas	8.2	36.3	6.5	0.5	30.6
Thematic mapping of grids m	10.8	78.5	7.7	0.5	74.9
Thematic mapping of grids HSD	-	-	-	-	-
Kernel density estimation	**12.1**	**103.5**	**9.4**	**1.2**	**81.1**

Table 6.3 shows the average PAI values for each hotspot mapping technique for each crime type. These were calculated from averaging the PAI values for each crime type, for each hotspot mapping technique, and for all periods of input and output data. These results were generated to see if differences in the prediction performance of hotspot mapping techniques were consistent, and to explore differences between crime types. The results from Table 6.3 show that for each crime type, KDE consistently proved to be the best technique for predicting patterns of crime. The spatial ellipses technique was not, though, the poorest performer for each crime type. Thematic mapping of output areas generated the lowest PAI values in Camden/Islington for burglary dwelling and in one case for thefts from vehicles. However, in Newcastle, the spatial ellipses technique was consistently the poorest performer.

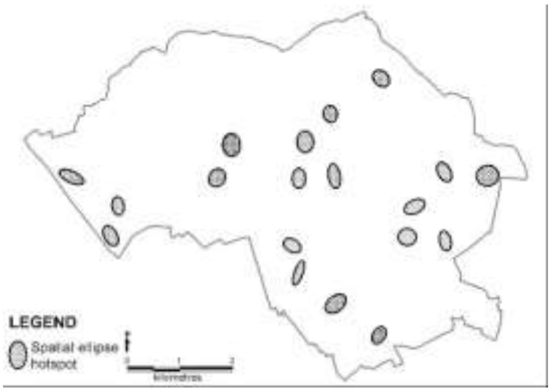

(a) STAC 250m: PAI 1.54

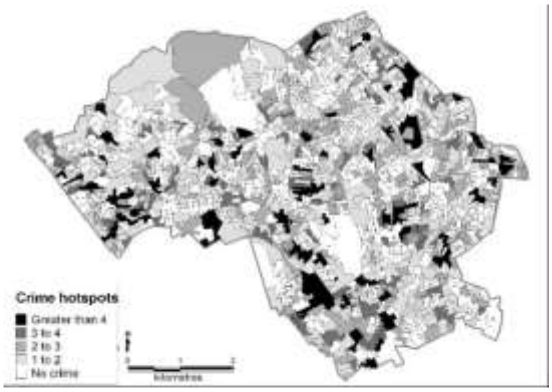

(b) Thematic mapping of OAs: PAI 1.26

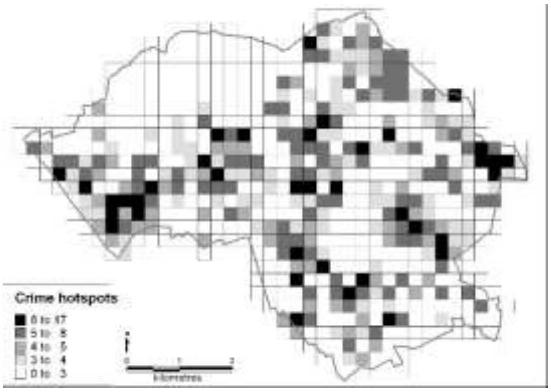

(c) Thematic mapping of grids 250m: PAI 1.64

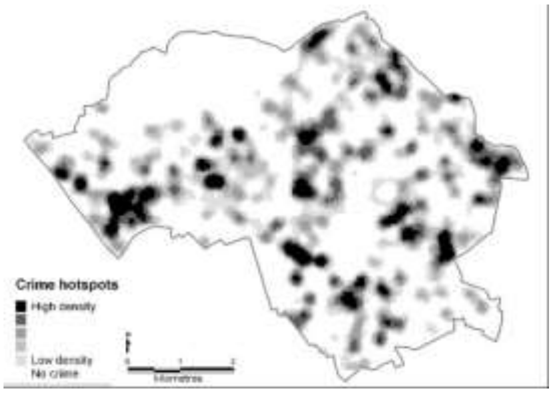

(d) KDE: PAI: 2.18

Figure 6.2. Hotspot maps generated from three months of Camden/Islington burglary dwelling input data (measurement date of the 1st January 2010) using (a) STAC, (b) thematic mapping of output areas, (c) grid thematic mapping, and (d) kernel density estimation. Each map is shown with its PAI value measured for one month of output data.

Figure 6.2 shows hotspot maps generated for each of the common hotspot analysis techniques from three months of Camden/Islington burglary dwelling input data when the measurement date was the 1st January 2010. The figures show that each technique identified similar areas, but in terms of the ability to predict spatial patterns of burglary dwelling over the next one month, the KDE map had a higher PAI value (2.18 compared to 1.64, the next highest PAI value for thematic mapping of grid cells output). That is, KDE mapping output was better than the others at predicting where burglary dwellings occurred.

Table 6.4. PAI values for burglary dwelling, theft from the person, theft from vehicles, theft of vehicles in Camden/Islington and Newcastle, and for assaults with injury in Newcastle. The higher the PAI, the greater the ability of the hotspot map to predict where future crimes (of the same respective type) will occur. Results in bold relate to the highest PAI values and results in italics relate to the lowest PAI values.

	Camden/Islington				**Newcastle**	
Crime type	**Average PAI (01/01/2010)**	*Standard deviation of PAI*	**Average PAI (11/03/2010)**	*Standard deviation of PAI*	**Average PAI (01/04/2010)**	*Standard deviation of PAI*
Burglary dwelling	1.56	0.39	*1.59*	0.42	7.5	3.4
Theft from the person	**3.24**	1.17	**3.18**	1.07	**46.3**	36.2
Theft from Vehicle	1.89	0.41	2.56	0.65	6.0	2.3
Theft of Vehicle	*1.53*	0.52	2.14	0.60	*0.5*	0.3
Assault with injury	-	-	-	-	40.3	29.7

Table 6.4 summarises the PAI values for each crime type. These are presented in the table as three sets of results – for Camden/Islington when the measurement date was the 1st January 2010 and when it was the 11th March 2010; and for Newcastle using the measurement date of 1st April 2010. These results show there were differences among the hotspot maps for different crime types in their ability to predict spatial patterns. Hotspot maps of theft from the person consistently produced the highest PAI values, and were clearly higher than PAI values for the other crime types. This is with the exception of the PAI values for assault with injury from Newcastle that also generated high PAI values. PAI values for burglary dwelling and theft of vehicles in Camden/Islington were similar to each other; however, in Newcastle the PAI value for theft of vehicles was very low. This low PAI value for theft of vehicles was perhaps a reflection of the low volume of crimes of this type in Newcastle, and (in following the results from research study 1 –

chapter 5) the lower extent of spatial clustering – over 16 weeks of theft from vehicle data were required before any clustering was significant. The consequence of this is that hotspots determined using retrospective data on theft of vehicles in Newcastle appear not to be that useful for predicting where these types of crime will occur in the future. The standard deviation values for the crime types indicated the variability in the results generated by hotspot mapping techniques.

Many practitioners make two assumptions from hotspot maps: all techniques are as good as each other, and no differences exist between crime types (Eck et al., 2005). These results counter these two assumptions. Firstly, hotspot maps generated for different crime types were found to differ in their performance for predicting patterns of crime. These results suggest that hotspot maps of theft from person are better at predicting where these types of incidents occur in comparison to hotspot maps of burglary dwelling, theft from vehicles and theft of vehicles. The results from Newcastle also showed the high level of prediction performance of retrospective data on assaults with injury. Secondly, the results show that different techniques vary in their performance to predict where crime may occur in the future, with KDE consistently outperforming the other common hotspot analysis techniques. Additionally, analyses using data for two measurement dates for Camden/Islington, and comparing to a different measurement date for Newcastle, have shown the results on the prediction performance of common hotspot analysis techniques and the results between crime types to be consistent.

6.4.1. How many crimes can a hotspot map predict?

The PAI provides a useful comparative measure, but more useful to practitioners is a measure that compares how many crimes a hotspot map is likely to predict. To measure and compare the number of crimes a hotspot map is likely to predict requires the area that the mapping technique determines as *hot* to be controlled for in size. That is, there would be little use in saying that a technique can predict 100% of all future crimes if the area it determines as *hot* is the entire study area. For the purposes of demonstrating how many crimes a hotspot mapping technique can predict, the area determined as *hot* was controlled to be 3% of the entire study area[12]. As KDE consistently produced the highest PAI values,

[12] 3% was chosen as this was considered to be representative of the size of area to which police resources could practically be allocated and targeted.

the KDE technique was selected for comparing how many crimes could be predicted for each of the different crime types.

Table 6.5 shows PAI results for KDE when three months of input data were used for generating hotspot maps for each crime type to determine the number of crimes that it could predict in the month following the measurement date. Table 6.5a shows the results for Camden/Islington using a measurement date of the 1st January 2010. Table 6.5b shows the results for Newcastle using a measurement date of the 1st April 2010.

Table 6.5. PAI and actual crimes predicted using KDE to generate a hotspot map from the previous three months of crime data and determine where crimes in the next month may occur, using (a) a measurement date of the 1st January 2010 for Camden/Islington, and (b) a measurement date of the 1st April 2010 for Newcastle. The area determined as *hot* was controlled to represent the top 3% of KDE values within the study area.

(a) Camden/Islington

Crime type	PAI	Crimes committed in January 2010	Number of crimes in hotspots	Percentage of crimes in hotspots
Burglary dwelling	2.77	470	39	8%
Theft from the person	6.59	460	91	20%
Theft from vehicle	3.98	962	115	12%
Theft of vehicle	3.26	307	30	10%

(b) Newcastle

Crime type	PAI	Crimes committed in April 2010	Number of crimes in hotspots	Percentage of crimes in hotspots
Burglary dwelling	6.7	130	21	16%
Theft from the person	389.1	60	41	68%
Theft from vehicle	9.4	190	48	25%
Theft of vehicle	0.0	11	0	0%
Assault with injury	79.2	154	68	44%

The results shown in Table 6.5 again highlight the differences between crime types in their PAI values and show the relatively high PAI results that KDE generated.

Interestingly, these PAI values were higher than the previous average values (Table 6.4) indicating that the currency of input data and the output period may influence differences in PAI results. When the size of the area that was determined as *hot* was controlled to represent the top 3% of KDE values in the study area, it was possible to compare the number of crimes for each crime type that occurred in the next month in the areas determined as hotspots. The results from this analysis show that for Camden/Islington, in all cases at least 8% of crime took place in the month that followed the measurement date, in an area representing only 3% of the study area. This result, suggesting that least 8% of crime took place in 3% of the study area, is useful for validating the role of hotspot mapping for helping to determine where to focus police and crime prevention resources. The proportion of crime in the hotspots was higher for theft from vehicles (12%) and theft from the person (20%), indicating that for the latter, if crime prevention resources are targeted to just those areas determined as hotspots, there is the chance of tackling a fifth of all crime that is committed across the whole study area.

The results for Newcastle show that up to 68% of theft from the person and 44% of assault with injury offences were predicted to occur in the *hottest* 3% of the study area. The levels of crime prediction in Newcastle were also higher than Camden/Islington's results for burglary dwelling (16%) and theft from vehicles (25%). However, for theft of vehicles, none of the crimes that took place in April 2010 were committed in the areas that were identified as hotspots. This is likely to reflect the low volume of theft of vehicle offences that occurred in Newcastle in April 2010 (n=11), and the fact that three months of theft of vehicle input data were not clustered (i.e., the areas identified as hotspots were not statistically defined as being hotspots).

6.5. Interpretation and conclusions from research study 2

With reference to this research study's hypothesis (hypothesis 2) – *common hotspot mapping techniques differ in how accurately they predict spatial patterns of crime* - the results show that kernel density estimation consistently outperformed the other common hotspot mapping techniques in predicting spatial patterns of crime. This consistency in KDE outperforming the other techniques was not only across the two study areas and for different measurement dates, but also for the range of different crime types. The results also showed how hotspot maps differed in their prediction performance for different crime types, with theft from the person hotspot maps proving to be better predictors of where these types of crime are likely to occur in the future than hotspot maps for predicting vehicle crime and burglary dwelling.

When the hotspot area was controlled to represent the top 3% of KDE values (i.e., covering 3% of the study area), the results showed the high levels to which KDE hotspot analysis can predict crime. For example, 68% of thefts from the person and 44% of assaults with injury were predicted to occur in the areas representing just 3% of the study area of Newcastle. However, KDE hotspot maps were not excellent predictors of crime in all cases. For example, none of the Newcastle thefts of vehicle offences that occurred in the prediction measurement period of April 2010 took place in the areas representing the top 3% of KDE values determined using crime data on where thefts of vehicles had previously occurred. However, it is important to recall that three months of retrospective crime data were used in the creation of each KDE map where the top 3% of KDE values were analysed for the number of crimes that were predicted. In research study 1, the results showed that 16 weeks and 5 days of data were required before the theft of vehicle input data for Newcastle showed significant evidence of clustering. For all other crime types, hotspots were evident from only a few weeks of retrospective crime data. In research study 2, the hit rate calculations for the top 3% of KDE values were determined from KDE maps generated using three months of crime data. Statistically, hotspots of theft of vehicles were not present in three months of retrospective crime data, and instead, the theft of vehicle data were randomly distributed for this period. This means that the KDE map generated from three months of theft of vehicles data was in fact only a geographical representation of a random spatial pattern, albeit with a range of KDE values representing this pattern. This leads to the conclusion that if hotspots do not exist in the crime data that are used for producing KDE maps, it is unlikely that these maps will perform well in predicting where crime may occur in the future. In contrast, when hotspots are identified from retrospective crime data and these are shown geographically using KDE, these hotspot maps are good predictors of where crime is likely to occur.

The KDE technique requires the user to enter two input parameters to generate mapping output – the cell size and the bandwidth size. Different KDE cell sizes and bandwidth sizes were used for the two study areas, and in part may help explain the differences in the PAI results for the two areas. The next research study will examine the influence that KDE cell size and bandwidth size have on the spatial crime prediction performance of KDE hotspot maps.

7. Research study 3: A metric comparison of the influence that technical parameters used in hotspot analysis can have on spatial crime prediction performance

7.1. Introduction

Research study 2 showed that kernel density estimation consistently outperformed the other common hotspot analysis techniques in predicting where crime is most likely to occur. KDE has also increasingly become the hotspot analysis technique of choice by police and public safety analysts and researchers (Chainey, 2013; Eck et al., 2005). Because of these research findings and KDE's popularity with researchers and practitioners, the focus of research study 3 is on KDE, testing whether the parameters a user is required to enter when producing KDE hotspot mapping have an influence on the technique's spatial crime prediction performance (hypothesis 3).

Like many spatial analysis techniques, KDE requires the researcher to input values of certain parameters in order to produce mapping output. The two main parameters for KDE are the value for the cell size (sometimes referred to as the resolution) and the value for the bandwidth (often referred to as the search radius). An alternative method to specifying a fixed bandwidth is the adaptive KDE approach where the bandwidth varies based on a user-determined number of neighbours to include in the kernel density calculation. The adaptive kernel approach is rarely used by crime mapping practitioners. This is because the study of crime hotspots is often towards crime concentrations in urban areas, rather than rural areas where, for the latter, an adaptive bandwidth may be more appropriate (Chainey et al., 2008). The focus of this next research study was towards the more commonly used fixed kernel bandwidth approach.

There is currently very little guidance on cell size and bandwidth size selection for the practical application of KDE hotspot mapping for policing and public safety, with the researcher either giving little thought to these values and their influence, settling for the default values determined by their KDE software application, or drawing from their own particular whims, fancies or experience (Chainey and Ratcliffe, 2005; Eck et al., 2005). Also, in a number of studies that have used KDE (such as Jefferis' 1999 assessment of KDE and other hotspot analysis techniques, and Johnson et al.'s 2009 and 2012 analyses of KDE in comparison to prospective mapping) either little thought has been given to parameter selection, or little attention has been placed on the influence that KDE input parameters may have on the research findings.

7.2. Chapter aims and structure

Examined in this research study is the influence that cell size and bandwidth size have on the prediction performance of KDE hotspot mapping outputs. Section 7.3 describes the method for these experiments and includes a detailed examination of the KDE equation to identify the role that cell size and bandwidth size parameter values have on KDE output. The current guidance on determining values for these parameters and the values used for this research study's experiments are then reviewed. The experiments conducted in this research study use a different set of input data periods and data output periods than those used in study 2. These are described in the method section along with the prediction measures that were used to determine the influence that different cell sizes and bandwidth sizes have on the prediction performance of KDE hotspot analysis outputs.

Section 7.4 presents the results from these KDE cell size and bandwidth size experiments. These results are then considered collectively to interpret how this research study's findings may influence practice and how the results influence subsequent research parts.

7.3. Method

This research study follows the general methodological process that was used in study 2 (that compared the spatial prediction measures of commonly used hotspot mapping techniques) by comparing the influence that different cell size and bandwidth size values have on the spatial prediction performance of KDE hotspot mapping outputs[13]. The research also uses the more detailed and complete measures of spatial prediction that were described in the method chapter (e.g., accuracy concentration curves) and compares these to PAI values.

7.3.1. An examination of the KDE function to explore the influence that cell size and bandwidth size has on the density values that are calculated

Recall from chapter 2 that the kernel density estimation function is applied to a dataset of points to obtain a smooth surface estimate representing the density of the point

[13] A paper describing this method and the results has been published: Chainey, S.P. (2013), "Examining the influence of cell size and bandwidth size on kernel density estimation crime hotspot maps for predicting spatial patterns of crime", *Bulletin of the Geographical Society of Liege* 60: 7-19.

distribution. In mathematical terms, KDE is expressed as:

$$f(x,y) = \frac{1}{nh^2} \sum_{i=1}^{n} k\left(\frac{d_i}{h}\right) \qquad (5)$$

where $f(x,y)$ is the density value at location (x,y), n is the number of incidents/points, h is the bandwidth, d_i is the geographical distance between incident i and location (x, y) and k is a density function, known as the kernel. Although k can take many forms (with there being little difference in the results between different functions, e.g., Bailey and Gatrell, 1995), a common choice for k is the quartic function (Bailey and Gatrell, 1995; Chainey et al., 2002; Chainey and Ratcliffe, 2005; Levine, 2010; Ratcliffe, 2004; Williamson et al., 1999).

Examination of the components of the KDE equation show that the density value for each location is affected by the number of points, the spatial distribution of these points, and the bandwidth size. For the purpose of generating a KDE hotspot map of crime for a single study area using data for one particular retrospective period of previous incidents, the number of crime incidents across the area would remain the same, the spatial distribution of the crime incidents would also remain the same, therefore neither would have an effect on generating different density estimates. The size of the bandwidth is determined by the user, therefore, for a KDE hotspot map of crime for a single study area using data for one particular retrospective period of previous incidents, different values of f may be calculated at each x,y location when different bandwidth sizes are used, consistent with volume preservation across the study area as a whole. Each x,y location is represented as a grid cell (the coordinates referring to the centroid of that cell), with the calculated density value f attributed to each cell. The cell size chosen by the researcher can vary, resulting in many calculations of f if the cell size is small or much fewer calculations if the cell size is large. While cell size is not an input to the KDE equation, the representation of these density values for areas of different size will be subject to the Modifiable Areal Unit Problem (Openshaw, 1984) – different size cells may produce different results of the KDE distribution of crime.

There is currently very little guidance on the choice of cell size a researcher should select and no research that investigates the impact it can have on the prediction performance of KDE crime hotspot maps. The little guidance that is offered is by Chainey and Ratcliffe (2005) who recommend that a suitable KDE cell size to choose for crime hotspot mapping is to divide the shorter side of the study area's minimum bounding rectangle (MBR) by

150. Whilst simple to calculate and used to determine the default cell size in the Hotspot Detective MapInfo add-on (Ratcliffe, 2004), this approach has not been rigorously evaluated.

The choice of bandwidth size for crime researchers to select is similarly not prescribed. For applications where these is a need to determine the number of neighbouring geographic polygon units to include in a calculation, such as with local spatial regression where local regression equations are applied to geographic units that make up defined sets of neighbourhoods, several bandwidth size optimisation routines such as the Mean Integrated Square Error (Bowman and Azzelini, 1997; Fotheringham et al., 2000), the Akaike Correlation Coefficient and the Cross Validation method (Brunsdon, 1995; Fotheringham et al., 2002; Silverman, 1986) can be used. However, these bandwidth optimisation routines are not appropriate for point pattern analysis. As an alternative, Bailey and Gatrell suggest a value derived from calculating $h = 0.68n^{-0.2}$ as a 'rough choice' (Bailey and Gatrell, 1995: 86) for the bandwidth (where n is the number of observed events across the study area). However, as Bailey and Gatrell's measure does not consider the spatial distribution of the observed events, this can mean that two sets of data for the same study area, with the same spatial coverage of events, but much different volumes of events, can produce very different bandwidth values. Instead, it is often advised that bandwidth size should be specified from experience and in ways that are sensitive to the nature of the application and the context in which it is applied (Bailey and Gatrell, 1995; Chainey and Ratcliffe, 2005). Even while experience may guide an appropriate choice of bandwidth size, to date, a detailed examination of the influence that bandwidth size may have on KDE hotspot mapping output has not been conducted, with many crime researchers, instead, not considering whether their choice of bandwidth size influenced their findings. For example, Johnson et al. (2008b) chose a bandwidth size of 200 m in their study that compared KDE to a prospective mapping approach, with the choice of this bandwidth not being subject to any further investigation as to whether its size had an impact on their findings.

Chainey (2011) recommends a good starting bandwidth size to apply for KDE is to measure the shorter side of the study area's MBR, divide by 150 (i.e., the calculation that is used to determine an appropriate cell size as described above), and multiply this value by 5. Whilst simple to calculate, the choice of this bandwidth size has not been evaluated, but is commonly applied - this bandwidth size calculation is used by Hotspot Detective for MapInfo (Ratcliffe, 2004). Others suggest experimenting with different sizes of bandwidth (Bailey and Gatrell, 1995; Chainey and Ratcliffe, 2005; Eck et al., 2005). Whilst experimenting in this way encourages the researcher to explore their data under different bandwidth conditions, it often leaves the researcher choosing the mapping output that 'looks the best' (Chainey and Ratcliffe, 2005: 159) in the context that KDE is applied (e.g., rural versus urban areas), rather than being more scientifically informed on the influence that bandwidth size selection may have on the prediction performance of KDE hotspot mapping output.

7.3.2. Data inputs, data outputs and measurement date used for generating KDE maps of different cell sizes and bandwidth sizes

Kernel density estimation hotspot maps were created using MapInfo version 10.5 and the Hotspot Detective v2.1 add-on (Ratcliffe, 2004). One study area was chosen for this research study - Newcastle - and two crime types were chosen for analysis – burglary dwelling and assault with injury. It was believed the study did not need replicating for the Camden/Islington study area, nor for a larger number of crime types because the results of the experiments would provide sufficient information to make conclusions on the impact that cell size and bandwidth has on KDE mapping output. Newcastle was selected as the study area rather than Camden/Islington because of the focus on burglary dwelling and assault with injury that was planned in subsequent studies in the current research.

In following the method described in the method chapter (Chapter 4) a suitable date had to be chosen within the Newcastle data time period as the day on which retrospective data were selected to generate hotspot maps against which *future* events could be compared. For simplicity, the 1st April 2010 was used to maximise the use of six months of retrospective data for generating KDE hotspot maps, and to use the complete set of six months of data after this date for measuring the hotspot maps' performance for predicting future events. Different measurement dates were not applied because research study 2 that compared different measurement dates found no difference in the results. Confidence was, therefore, placed in the selection of the 1st April 2010 as a measurement date that would generate representative results.

The retrospective time data were organised into six time periods and used as input data to generate KDE hotspot maps. This meant that rather than using just one retrospective time period (e.g., the three months prior to the measurement date) which may generate an anomalous result, the use of a number of retrospective time periods would form a more reliable basis on which to draw conclusions. Retrospective input data were organised into the time periods shown in Table 7.1a, for each crime type. This approach in using different periods of data as the input data was also followed through to the analysis involving the output data. Six time periods of output data were used. This meant that rather than using just one output data period for the research (e.g., the three months after the measurement date), multiple output data time periods were used in order to generate results from which more reliable conclusions could be made. Output data were organised into the time periods shown in Table 7.1b. This meant that KDE hotspot maps that were generated for each period of input data would be measured for their performance to predict spatial patterns of crime, when the prediction period was the next month, the next two months, and to the next six months.

Table 7.1. (a) The temporal periods of input data for generating hotspot maps, for a measurement date of the 1st April 2010 and (b) the temporal periods of output data for calculating the performance of KDE hotspot maps for predicting spatial patterns of crime

(a)

Time periods of data used to create KDE hotspot maps					
1 month	2 months	3 months	4 months	5 months	6 months
1 Mar 2010 - 31 Mar 2010	1 Feb 2010 - 31 Mar 2010	1 Jan 2010 - 31 Mar 2010	1 Dec 2009 - 31 Mar 2010	1 Nov 2009 - 31 Mar 2010	1 Oct 2009 - 31 Mar 2010

(b)

Time periods of data used to measure the prediction performance of KDE hotspot maps					
1 month	2 months	3 months	4 months	5 months	6 months
1 Apr 2010 - 30 Apr 2010	1 Apr 2010 - 30 May 2010	1 Apr 2010 – 31 Jun 2010	1 Apr 2010 - 31 Jul 2010	1 Apr 2010 - 31 Aug 2010	1 Apr 2010 - 30 Sep 2010

7.3.3. Cell and bandwidth sizes used for determining the influence these parameters have on KDE hotspot analysis output

Eight cell size values were chosen for comparison: 30 metres, 60 m, 90 m, 120 m, 150 m, 180 m, 210 m and 240 m. A value that is often used for the cell size (as referred to in section 7.3.1) is the result from measuring the shortest side of the minimum bounding rectangle of the study area, and dividing this distance by 150. Although the choice of 150 is rather arbitrary, in practice it provides a useful starting measure. This calculation gave the value of 89.6 (rounded up to 90 m). It was, therefore, thought useful to generate results for this measure in comparison to other cell size values, using multiples of 30 m in the cell size experiments. For each cell size experiment, the bandwidth was controlled to a single size: a bandwidth of 450 m was used (five times 90 m), as per the guidance described in section 7.3.1.

Five bandwidth size values were chosen for comparison: 100 m, 200 m, 300 m, 400 m, and 500 m. If the recommendations of Chainey (2011) were followed (i.e., five times the cell size) this would have suggested a bandwidth value of 450 m. Rather than use multiples of 150 m, multiples of 100 m were used instead in order to explore the influence of a small bandwidth (100 m) and to help more simply present the results. This approach would also still enable a comparison between the outputs generated between 400 m and 500 m as an indication of the effectiveness of the rather crude selection of 450 m as a bandwidth size. For each bandwidth-size experiment, the cell was controlled to a single size: a cell size of 90 m was used, as per the approach described in the paragraph above.

A final parameter to consider for KDE hotspot map generation was the threshold value for determining which areas were *hot*. For purposes of research comparison, the method used in study 2 of the research was followed. This involved using five thematic classes and default values generated from using the quantile thematic classification method in MapInfo. *Hot* was then determined by the top thematic class (see Figure 6.1 in the method section of research study 2, chapter 6, for an illustration of this process).

7.3.4. Measuring the prediction accuracy of KDE hotspots produced using different cell sizes and bandwidth sizes

A combination of PAI, accuracy concentration curves, area under the curve and CPI measures were used for measuring the prediction performance of KDE hotspot maps produced using different cell sizes and bandwidth sizes. PAI measures were aggregated and averaged for the periods of input data and for the periods of output data. This meant that the PAI measures could be compared, with any differences being explained in relation to the cell size and bandwidth size rather than different periods of input and output data. Aggregating and averaging PAI results was applied separately to the two crime datasets: burglary dwelling and assault with injury. The standard deviation and coefficient of variation of the PAI for each crime type across the eight different cell size values and eight bandwidth values were also calculated.

Accuracy concentration curves were calculated for both burglary dwelling and assault with injury KDE hotspot maps, for the full range of cell sizes and bandwidths, for a single input data period of six months and a single output data period of six months. Area under the curve and CPI values were calculated for the sub-sections of the accuracy concentration curves as described in the method chapter (chapter 4) and listed in Table 4.3. These were compared to their equivalent PAI values to examine whether the simple to calculate PAI missed any important details in assessing the prediction performance of mapping outputs.

7.4. Results

7.4.1. The influence of cell size on KDE hotspot maps for predicting where crime may occur

Table 7.2 shows the PAI results for Newcastle burglary dwelling and assault with injury KDE hotspot maps for different cell sizes. The PAI results for burglary dwelling varied between 6.6 for a cell size of 240 m to 7.1 for 30 m and 60 m cell sizes. The PAI results for assault with injury were much higher than those for burglary dwelling, but again showed only a small amount of relative variation from 59.9 for a cell size of 210 m to 68.5 for a cell size of 60 m. These results suggest that although PAI values decreased with increases in cell size, this difference was marginal. These results are also shown in Figure 7.1. There was little statistical variation in the results for each cell size, as indicated by the low coefficient of variation (CV) values, and little difference in the CV values between cell sizes.

Table 7.2. KDE hotspot map PAI, standard deviation (SD) and coefficient of variation (CV) results for burglary dwelling and assault with injury for different cell sizes

	Burglary dwelling			**Assault with injury**		
Cell size	**PAI**	**SD**	**CV**	**PAI**	**SD**	**CV**
30	7.1	0.60	0.08	68.4	3.06	0.04
60	7.1	0.66	0.09	68.5	2.96	0.04
90	6.7	0.53	0.08	65.1	2.66	0.04
120	6.9	0.57	0.08	64.5	2.50	0.04
150	6.7	0.53	0.08	63.0	3.17	0.05
180	6.8	0.64	0.10	64.3	2.95	0.05

210	6.7	0.46	0.07	59.9	2.39	0.04
240	6.6	0.48	0.07	60.2	2.85	0.05

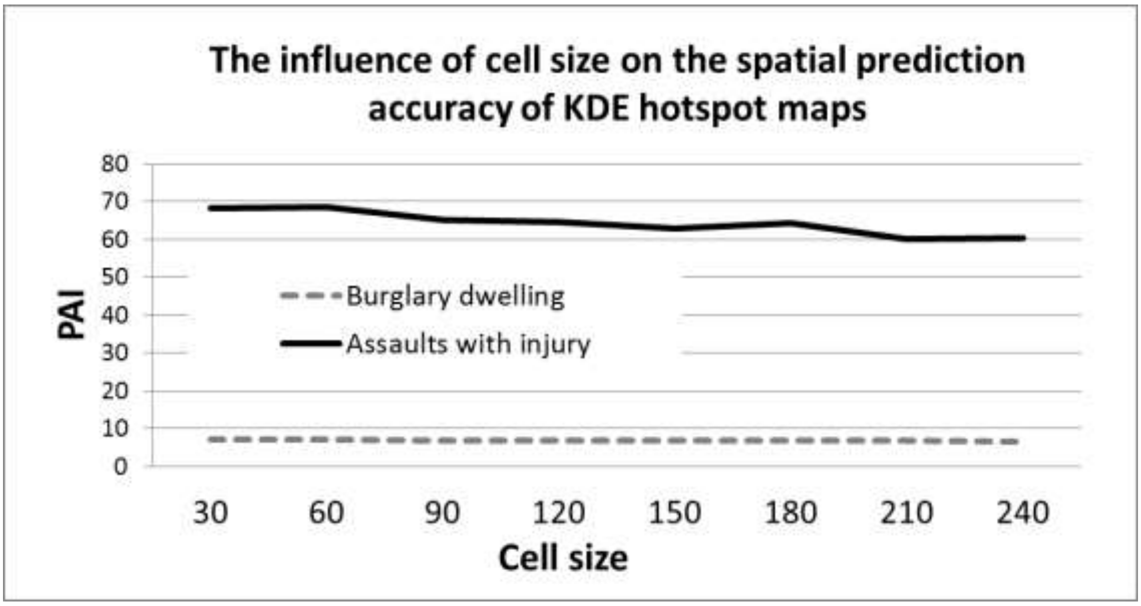

Figure 7.1. The influence of cell size on KDE hotspot map PAI values for burglary dwelling and assault with injury

Table 7.3. Crimes predicted using KDE outputs of different cell sizes for burglary dwelling and assault with injury, using three months of input crime data (January – March 2010) and three months of output data (April – June 2010). The area determined as *hot* was controlled to cover 1% of the study area's total area.

Crime type and cell size (m)	Crimes committed April – June 2010	Number of crimes in hotspots (1% of area)	Percentage of crimes in hotspots
Burglary dwelling: 30 m	329	29	8.8%
Burglary dwelling: 240 m	329	28	8.5%
Assaults with injury: 30 m	459	158	34.4%
Assaults with injury: 240 m	459	153	33.3%

The similarity in results for different cell sizes was further illustrated by the difference in the number of crimes predicted in KDE generated hotspot areas that were produced using different cell sizes (Table 7.3). When the KDE hotspot areas were controlled to identify 1% of the total study area (i.e., the 1% of areas with the highest KDE values), generated from 3 months of input data using cell sizes of 30 m and 240 m to predict where crimes would occur in the next 3 months, very similar results were produced. For burglary dwelling, KDE outputs generated using a 30 m cell size predicted 29 crimes, in comparison to 28 crimes using a cell size of 240 m. For assaults with injury, KDE outputs generated using a 30 m cell size predicted 158 crimes, in comparison to 153 crimes using a cell size of 240 m. That is, as the spatial resolution of the KDE hotspot map began to degrade, the ability of the map to predict where crime occurred in the future reduced, but only marginally.

7.4.2. The influence of bandwidth size on KDE hotspot maps for predicting where crime may occur

Table 7.4 shows the PAI results for burglary dwelling and assault with injury for different bandwidth sizes. The PAI results for burglary dwelling varied between 5.6 for bandwidth sizes of 700 m to 13.1 for 100 m bandwidth sizes. The PAI results for assault with injury were much higher than those for burglary dwelling, but also showed large variation from 42.9 for bandwidth sizes of 800 m to 142.8 for bandwidth sizes of 100 m. These results suggest that as bandwidth size increases, the performance of the KDE hotspot map to predict spatial patterns of crime degrades. These results are also shown in Figure 7.2. With the exception of burglary dwelling KDE hotspot maps generated using a bandwidth of 100 m, there was little statistical variation in the PAI results for each bandwidth size and little difference in the CV values between bandwidth sizes.

Table 7.4. KDE hotspot map PAI, standard deviation (SD) and coefficient of variation (CV) values for burglary dwelling and assault with injury for different bandwidth sizes

	Burglary dwelling			**Assault with injury**		
Bandwidth size (m)	**PAI**	**SD**	**CV**	**PAI**	**SD**	**CV**
100	13.1	2.8	0.22	142.8	11.53	0.08
200	11.1	1.3	0.12	91.7	4.65	0.05
300	8.7	1.0	0.12	79.4	3.03	0.04
400	7.1	0.7	0.10	68.3	3.54	0.05
500	6.5	0.6	0.09	60.2	2.60	0.04
600	5.9	0.6	0.11	54.3	2.52	0.05
700	5.6	0.6	0.11	48.6	2.23	0.05

800	5.7	0.5	0.09	42.9	1.98	0.05

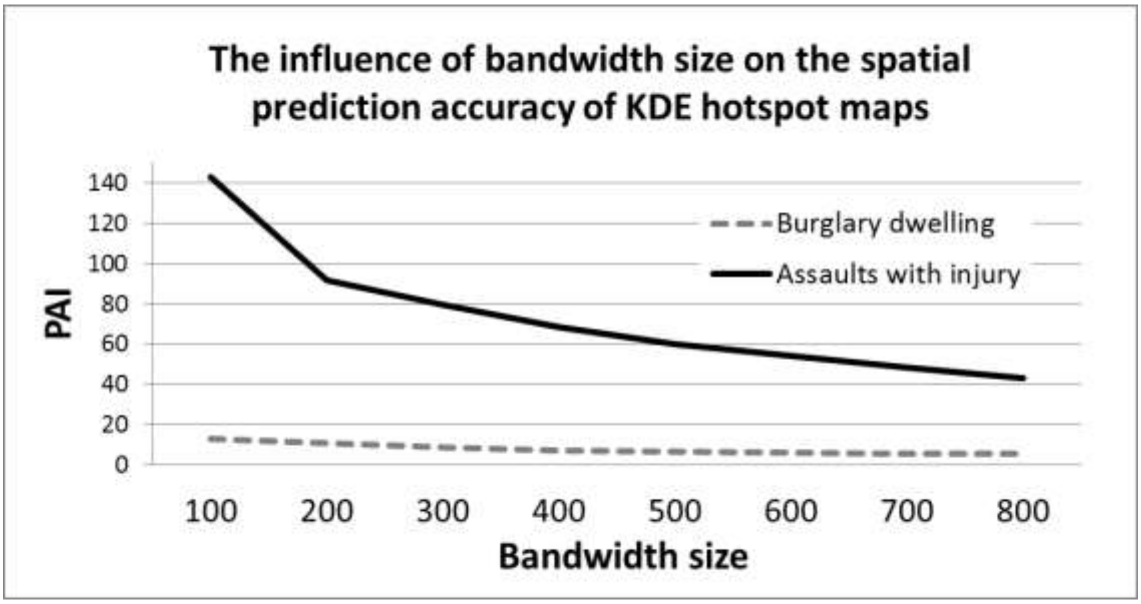

Figure 7.2. The influence of bandwidth size (m) on KDE hotspot map PAI values for burglary dwelling and assault with injury

Table 7.5. Crimes predicted using KDE outputs of different bandwidth sizes for burglary dwelling and assault with injury, based on using three months of input crime data (January – March 2010) and three months of measurement data (April – June 2010). The area determined as *hot* was controlled to cover 1% of the study area's total area

Crime type and bandwidth size (m)	Crimes committed April – June 2010	Number of crimes in hotspots (1% of area)	Percentage of crimes in hotspots
Burglary dwelling: 100 m	329	35	10.6%
Burglary dwelling: 800 m	329	22	6.7%
Assault with injury: 100 m	459	166	36.2%
Assault with injury: 800 m	459	137	29.8%

The difference in results for different bandwidth sizes is further illustrated by the difference in the number of crimes predicted from KDE generated hotspot maps of different bandwidth sizes (Table 7.5). To illustrate this, the KDE hotspot areas were controlled to identify only the top 1% of density values (i.e., the 1% of areas with the highest KDE values), generated from 3 months of input data using bandwidth sizes of

100 m and 800 m to predict where crimes would occur in the next 3 months. For burglary dwelling, KDE outputs generated using a 100 m bandwidth size predicted 35 crimes (i.e., 11% of all burglaries in just 1% of the study area), in comparison to 22 crimes using a bandwidth size of 800 m. For assaults with injury, KDE outputs generated using a 100 m bandwidth size predicted 166 crimes (i.e., 36% of all violent assaults in 1% of the study area), in comparison to 137 crimes using a bandwidth size of 800 m. That is, as the smoothing of the KDE hotspot map increased (caused by increases in bandwidth size), the performance of the map to predict where crime occurred degraded. These results also illustrate the relatively high proportion of crime that KDE hotspot maps can predict.

7.4.3. Using accuracy concentration curves, the area under the curve and the CPI to measure the prediction performance of KDE hotspot analysis output

I. Cell size results

The PAI provides a simple measure for comparing prediction performance of hotspot analysis output. A more detailed measure is the accuracy concentration curve. This allows the prediction of mapping output to be compared across the full areal extent of the study area (i.e., by comparing the number of crimes predicted at very small areal coverage levels (e.g., 1%) to larger areal coverage levels (e.g., 25%, 50%, 75% and to 100%)).

Figure 7.3 shows accuracy concentration curves for KDE hotspot maps of different cell sizes, of burglary dwelling for Newcastle, generated using six months of input data and six months of output data. These graphs are shown at different axis levels to examine in detail the variation between cell sizes. The graphs are interpreted by identifying the area that needs to be *searched* (by the proportion of the study area) for a certain proportion of offences to be identified. The area searched is ordered from high to low kernel density values. For example, in Figure 7.3a the top 20% of KDE values generated from six months of burglary dwelling data identified where approximately 60% of all offences took place in the following six months.

Figure 7.3a shows there was very little difference between the range of cell sizes used for the areas searched and the areas where offences occurred for the section of the graph between 0% to 35% of the study area coverage. That is, the prediction performance of the KDE hotspot maps generated for different cell sizes showed very little difference for the top 35% of KDE values. From approximately 35% of the study area, the prediction performance of KDE output generated using a cell size of 120 m began to degrade in comparison to the other KDE outputs for other cell sizes. This is illustrated in Figure 7.3a by the line representing the KDE output generated using a 120 m cell size being less vertical and separating from the other lines. However, from approximately 40% of the study area, each of the lines representing the KDE outputs for other cell sizes also flattened and began to show some differences in their prediction performance (as shown by the separation in the lines in Figure 7.3a).

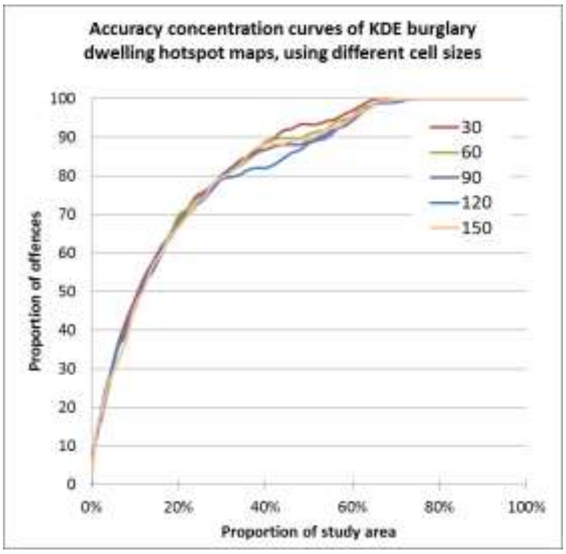

(a)

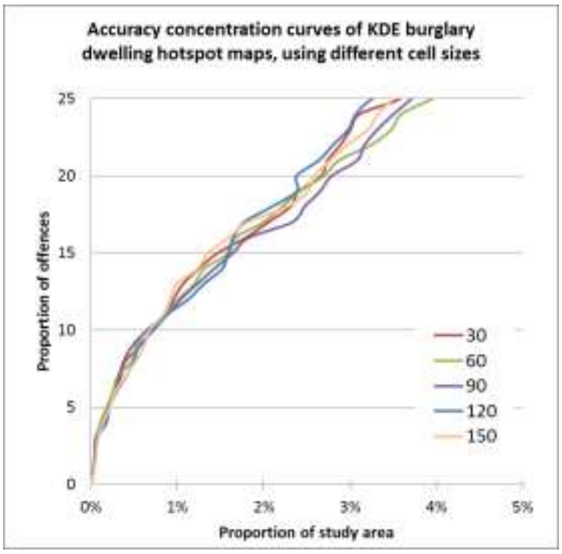

(b)

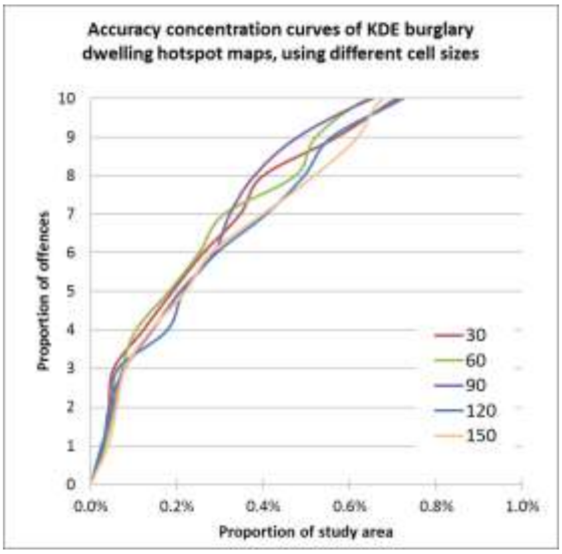

(c)

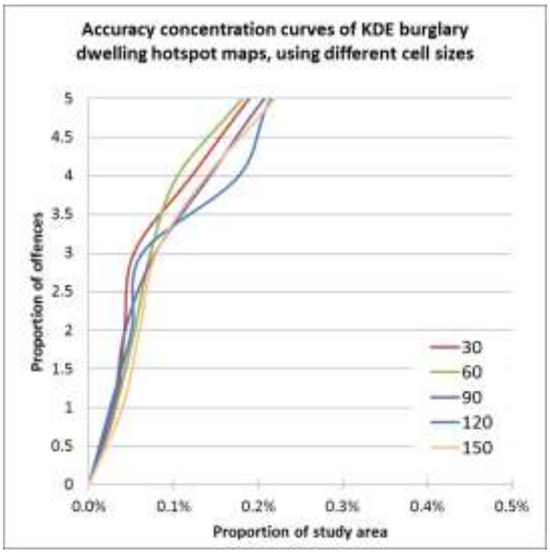

(d)

Figure 7.3. Accuracy concentration curves of Newcastle KDE burglary dwelling hotspot maps, generated using different cell sizes (and using a fixed bandwidth of 450 m)

Closer examination of the KDE hotspot outputs for each cell size explains why these differences occurred at these study area proportion levels in Figure 7.3a. The last column in Table 7.6 lists the proportion of the study area that was covered by KDE values. The KDE hotspot maps for each cell size are also shown in Figure 7.4. KDE values generated using a cell size of 120 m covered 32% of the study area (see Figure 7.4d). The remaining 68% of the cells covering the study area had density values of zero. This spatial coverage of KDE values is reflected in the line drawn for the 120 m cell size in Figure 7.3a that separates from the lines for the KDE outputs for other cell sizes. Up to a coverage of 32% of the study area, the KDE values generated using a cell size of 120 m provided some indication of where burglary dwelling may occur in the future, but after this point, the line reflects random variation. That is, the cells at this point were not arranged in order of where crime was predicted to occur (all cell values were zero), and, therefore, the likelihood that an offence fell inside a cell is down to random chance. Similarly, KDE values generated using cell sizes of 30 m, 60 m, 90 m and 150 m covered 44%, 43%, 40% and 40% of the study area respectively (see Figures 7.4a, b, c and e). It is again at these points on the chart in Figure 7.3a that the lines for each cell size flatten and follow a trend that reflects random variation.

Table 7.6. The proportion of the Newcastle study area searched across hotspots maps generated using KDE (six months of input data), relative to 5%, 10%, 25%, 50%, and 80% of burglary dwelling offences, for different cell sizes. Values in bold represent the smallest area that was searched for identifying the relevant proportion of crime.

Cell size (m)	**% of offences and the area searched (burglary dwelling)**					**% of study area with KDE values**
	5% of offences	**10% of offences**	**25% of offences**	**50% of offences**	**80% of offences**	
30	0.19%	0.71%	3.57%	**11%**	**30%**	44%
60	**0.18%**	**0.64%**	3.96%	**11%**	32%	43%
90	0.21%	0.65%	3.71%	**11%**	32%	40%
120	0.22%	0.72%	**3.26%**	11.5%	34%	32%
150	0.22%	0.68%	3.51%	11.5%	32%	40%

Closer examination of KDE hotspot analysis outputs for different cell sizes at smaller levels of the proportion of the study area again show there to be little difference in their prediction performance. Figure 7.3b, c and d show the lines representing the different cell sizes to vary very little, with there being no consistency in cells of a particular size showing any differences to others. This is supported by visual inspection of the hotspot maps for the different cell sizes in Figure 7.4, showing the same areas identified as hotspots, and with the main difference being the greater pixilation of the hotspot maps as cell size increases.

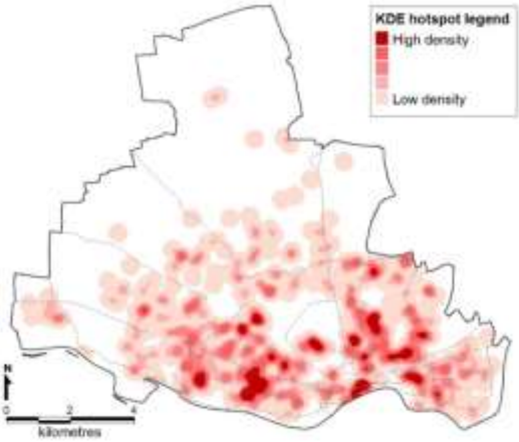

(a) Cell size: 30 m

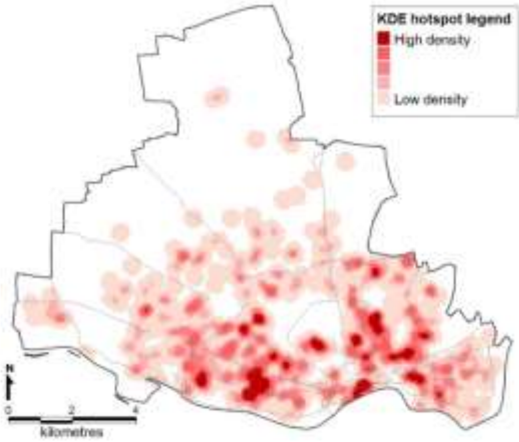

(b) Cell size: 60 m

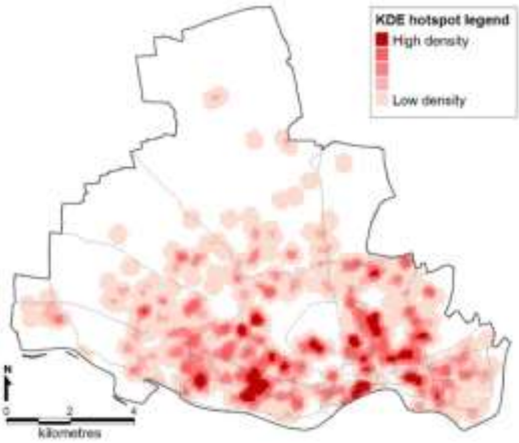

(c) Cell size: 90 m

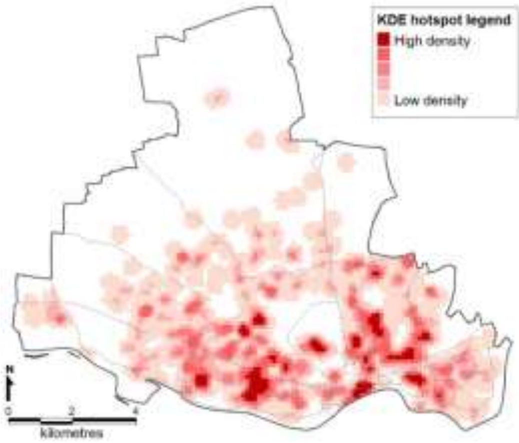

(d) Cell size: 120 m

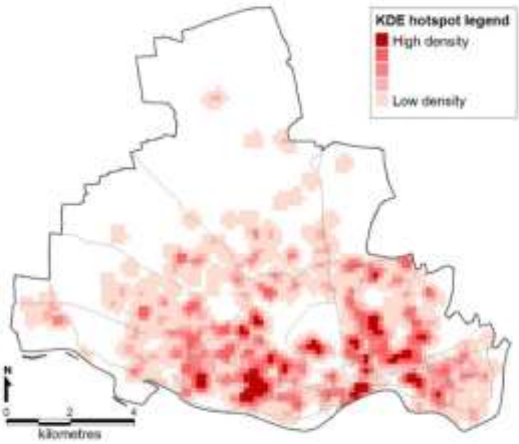

(e) Cell size: 150 m

Figure 7.4. KDE burglary dwelling hotspot maps of Newcastle, generated using six months of data and a fixed bandwidth of 450 m for cell sizes of (a) 30 m, (b) 60 m, (c) 90 m, (d) 120 m, and (e) 150 m

Table 7.6 lists the proportion of the area that would need to be searched for 5%, 10%, 25%, 50% and 80% of all offences to be identified (i.e., the size of the area where 5%, 10% … 80% of all burglaries were predicted to occur). These results identify the highly predictive nature of KDE hotspot maps. For example, KDE hotspot maps identified that 25% of all burglary dwelling offences that occurred during the output data period took place in less than 4% of the study area (identified from the top 4% of KDE values).

Accuracy concentration curves provide a useful means of examining in detail the prediction performance of mapping output across a full study area and across the full proportional range of offence volume. The area under the curve and Crime Prediction Index add to this detailed visual representation of the prediction performance of mapping output by providing statistical measures that allow for further interpretative analysis. Also, by splitting the area under the curve into sections helps to determine if there are differences in the prediction performance of KDE hotspot maps at different levels of study area coverage.

Table 7.7 lists the area under the curve and CPI values for KDE burglary dwelling hotspot mapping output for different cell sizes. These results help to further show that cell size has little impact on influencing the predictive performance of KDE hotspot mapping output. For example, for the section of the accuracy concentration curve representing 0%-0.5% of the study area and up to 5% of all offences, the area under the curves for the different cell sizes ranged between 0.000207 (for a cell size of 150 m) to 0.000216 (for a cell size of 30 m). For the section of the accuracy concentration curve representing 0%-5% of the study area and up to 25% of all offences, the area under the curves for the different cell sizes ranged between 0.008917 (for a cell size of 90 m) to 0.009219 (for a cell size of 120 m).

The CPI provides a better means of directly comparing the sub-sections of the area under the curve results (a CPI value of 1 represents a perfect prediction). CPI values for the full coverage and full offence extent ranged from 0.815 to 0.833, indicating the KDE burglary dwelling hotspot maps generated using different cell sizes were good predictors of where crime occurred in the future (see Table 7.7b). Closer examination of CPI values for different study area coverages and offence proportions showed that the KDE burglary dwelling hotspot maps performed better at predicting crime for the highest KDE values (i.e., the top 0.5% of KDE values). CPI values for the top 0.5% of KDE values ranged from 0.829 to 0.863. CPI values then fell to 0.55 for the top 20% of KDE values. These values reflect the gradual reduction in gradient of the accuracy concentration curve as it extends from 0% to 20% of the coverage of the study area (as shown in Figure 7.3a).

Table 7.7. (a) Area under the accuracy concentration curves, and (b) CPI values for Newcastle burglary dwelling KDE hotspot maps of different cell sizes (using a fixed bandwidth of 450 m). Values in bold relate to the largest area and largest CPI values.

(a)

Cell size (m)	**0.5% x 5%**	**1% x 10%**	**5% x 25%**	**10% x 50%**	**20% x 80%**	**100% x 100%**
30	**0.000216**	0.000763	0.009191	**0.029892**	**0.088646**	**0.833007**
60	0.000215	0.000768	0.009021	0.028862	0.087532	0.826522
90	0.000210	**0.000771**	0.008917	0.028702	0.087136	0.821530
120	0.000208	0.000736	**0.009219**	0.029465	0.08689	0.814785
150	0.000207	0.000727	0.009184	0.027834	0.085292	0.821723
Max values	*0.00025*	*0.001*	*0.0125*	*0.05*	*0.16*	*1*

(b)

Cell size (m)	**0.5% x 5% CPI**	**1% x 10% CPI**	**5% x 25% CPI**	**10% x 50% CPI**	**20% x 80% CPI**	**100% x 100% CPI**
30	**0.863**	0.763	0.735	**0.598**	**0.554**	**0.833**
60	0.858	0.768	0.722	0.577	0.547	0.827
90	0.841	**0.771**	0.713	0.574	0.545	0.822
120	0.830	0.736	**0.738**	0.589	0.543	0.815
150	0.829	0.727	0.735	0.557	0.533	0.822

The results using the metrics of accuracy concentration curves, the area under the curve and the CPI are comparable to the results generated using the PAI – that cell size has little influence on the prediction performance of KDE hotspot analysis output. This, therefore, shows the PAI provides a simple, yet effective measure of the prediction performance of hotspot analysis outputs. However, not only do the additional metrics of the accuracy concentration curves, the area under the curve and the CPI provide further detail in helping to examine hotspot mapping output prediction performance, they also appear to provide a better means of comparing the prediction performance results between different mapping outputs. In addition, these more detailed measures appear to offer a better means of examining the prediction performance of mapping output across a range of coverage areas and offence proportions. A weakness of the PAI is that there is no easy way to determine how much better a high PAI value is over a lesser value. While higher PAI values indicate a better prediction performance in the mapping output, it is a relative value rather than one that provides an absolute measure of good or not so good. The CPI provides a better means of determining direct comparability between hotspot mapping output – the closer the CPI value is to 1, the better the prediction. For example, the CPI value of 0.863 for the sub-section of the accuracy concentration curve from 0% to 0.5% of the study area for a cell size of 30 m suggests that KDE performs very well at predicting spatial patterns of crime for this very small area. In comparison, the CPI value of 0.735 for the sub-section of the accuracy concentration curve between 0% to 25% of the study area for the same cell size suggests the KDE hotspot mapping output does not perform as well in its prediction performance.

Generating CPI values across the range of coverages for the study area and offence proportions helps determine at which place across the coverage the mapping output predictions perform best. This is of value because in most cases police and crime prevention interest is towards the focusing of resourcing to small geographic areas. CPI values generated across the range of areal coverage levels and offence proportions provide a better means of determining if the output is good for predicting crime for very small coverage levels (the top proportion of mapping output values) and whether this then degrades as the predictions from the mapping output then decrease in value. That is, most practitioners are more likely to find value from a mapping technique that is very good at predicting crime for very small areas that can be specifically targeted rather than how the mapping technique performs in predicting crime for 25% or 50% of the coverage area.

Figure 7.5 shows the accuracy concentration curves for Newcastle KDE hotspot maps of different cell sizes for assault with injury. These again show that cell size had little influence on the prediction performance of KDE hotspot maps, and all of the main variation that occurred above study area proportion levels of 30% was a reflection of the coverage extent of KDE values and cells that had values of zero. For example, Table 7.8

shows the KDE assault with injury values extended to 30% to 42% of the coverage area for the different cell size outputs. Above these KDE study area extent levels, the accuracy concentration curves follow a trend of random variation, reflecting the fact that cells for these areas were not populated with KDE density values (i.e., these areas did not experience any crime). These KDE value coverage extents are shown in the Figure 7.6 hotspot maps with large areas containing no values. Similar to the hotspot maps for burglary dwelling shown in Figure 7.4, the KDE hotspot maps for assault with injury show very little variation in density representation of the spatial distribution of these offences, with the main differences being the greater pixilation of the hotspot maps as cell size increases.

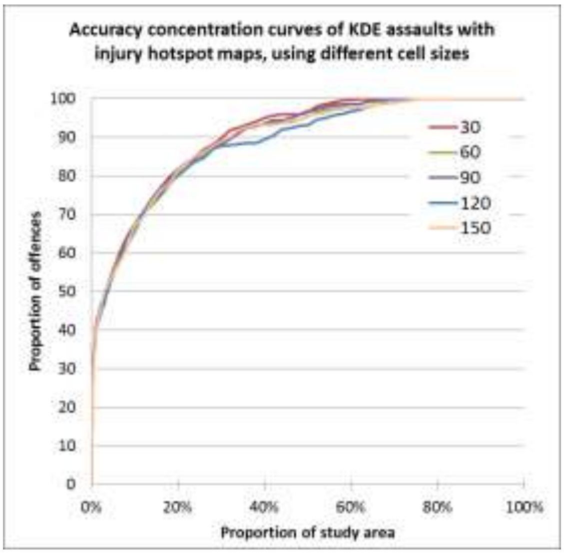

(a)

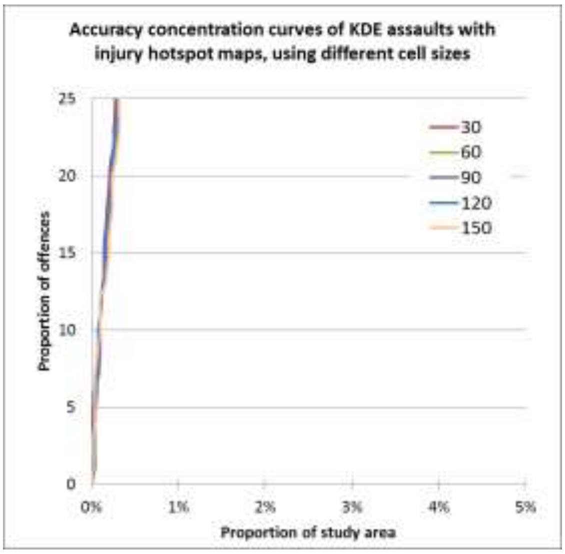

(b)

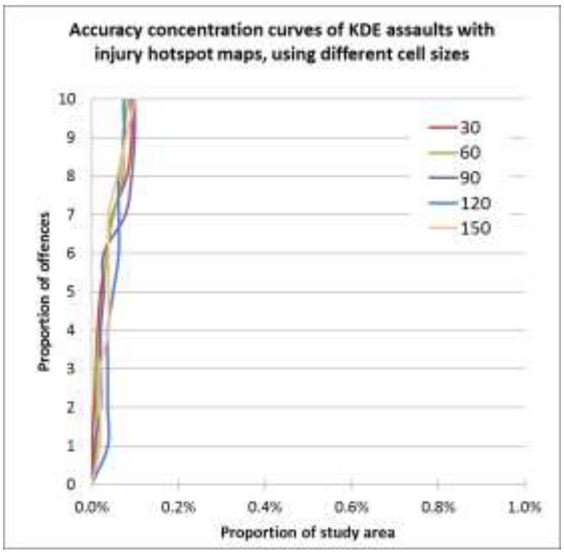

(c)

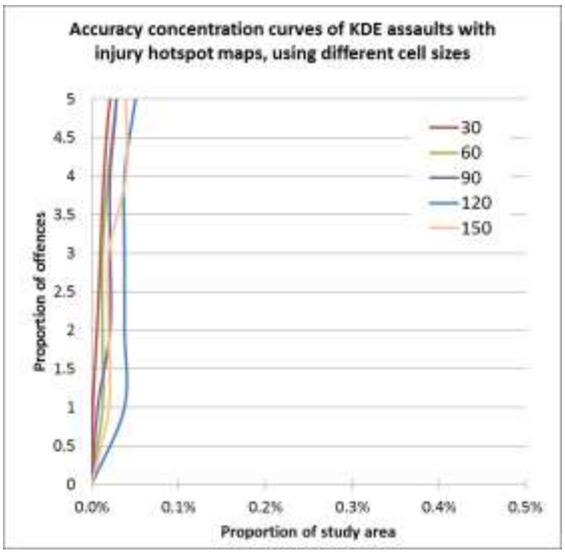

(d)

Figure 7.5. Accuracy concentration curves of Newcastle KDE assault with injury hotspot maps, generated using different cell sizes (and using a fixed bandwidth of 450 m)

The more detailed examination of the accuracy concentration curves up to 5%, 1%, and 0.5% of the study area (Figure 7.5b, c and d respectively) shows how much more vertical the curves for assault with injury were in comparison to the curves generated for burglary dwelling (Figure 7.3). This is indicative of a higher level of prediction performance in the assault with injury KDE hotspot maps. This is illustrated in Table 7.8, which shows that only 0.02% to 0.05% of the study area needed to be searched to identify 5% of all assaults, and no more than 0.32% of the study area needed to be searched to identify 25% of all assaults (compared to 0.2% and 3.5% of the study area respectively for burglary dwelling offences).

Table 7.8. The proportion of the Newcastle study area searched across hotspots maps generated using KDE (six months of input data), relative to 5%, 10%, 25%, 50%, and 80% of assault with injury offences, for different cell sizes. Values in bold relate to the smallest area that was searched for identifying the relevant proportion of crime.

Cell size (m)	**% of offences and the area searched (assaults with injury)**					**% of study area with KDE values**
	5% of offences	**10% of offences**	**25% of offences**	**50% of offences**	**80% of offences**	
30	**0.02%**	0.09%	**0.27%**	3.37%	**20%**	42%
60	0.03%	0.09%	0.29%	3.40%	**20%**	40%
90	0.03%	0.10%	0.29%	3.42%	**20%**	37%
120	0.05%	**0.08%**	0.32%	3.78%	**20%**	30%
150	0.04%	0.10%	0.32%	**3.07%**	**20%**	37%

The area under the curve values for the different cell sizes presented in Table 7.9 also showed very little variation. For example, for the section of the accuracy concentration curve representing 0%-5% of the study area and up to 25% of all offences, the area under the curves for the different cell sizes ranged between 0.01213 (for a cell size of 150 m) to 0.01218 (for a cell size of 90 m).

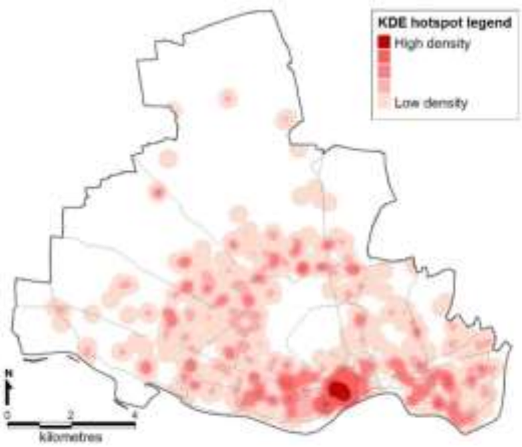

(a) Cell size: 30 m

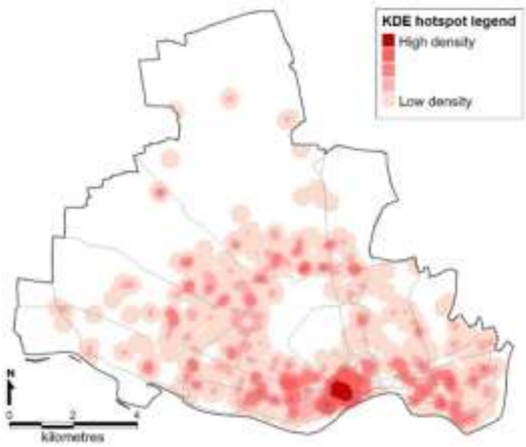

(b) Cell size: 60 m

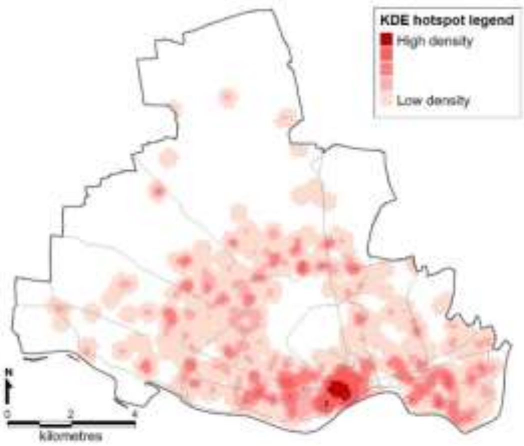

(c) Cell size: 90 m

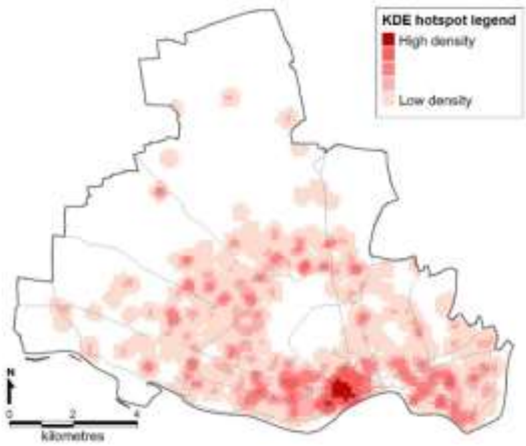

(d) Cell size: 120 m

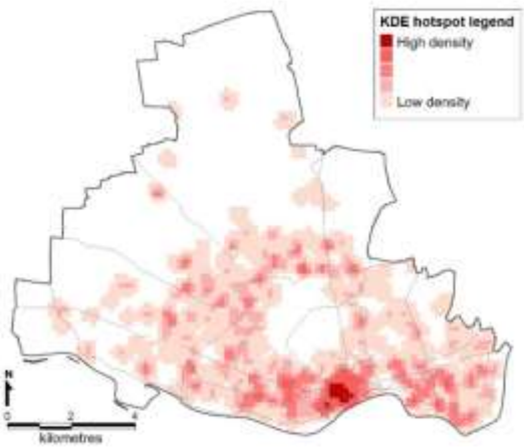

(e) Cell size: 150 m

Figure 7.6. KDE assault with injury hotspot maps of Newcastle, generated using six months of data and a fixed bandwidth of 450 m for cell sizes of (a) 30 m, (b) 60 m, (c) 90 m, (d) 120 m, and (e) 150 m

The high degree of prediction performance in the KDE assault with injury hotspot maps is also illustrated in the CPI results. Recall that a perfect prediction is 1. CPI values for the sections of the accuracy concentration curve representing up to 10% of the study area and up to 50% of all offences were no lower than 0.93 (see Table 7.9b). Indeed, CPI values were particularly high (up to 0.98) for the section of the accuracy concentration curve representing 0%-0.5% of the top KDE values covering the study area. Similar to the KDE hotspot maps produced for burglary dwelling, the CPI values fell as KDE values reduced, albeit at a lower rate of degradation (see Figure 7.7), suggesting that the highest values produced by KDE for hotspot analysis (i.e., covering the top 0.5% of the study area) provided a more accurate prediction of where crime may occur in the future than lower KDE values. Examination of the prediction performance of KDE assault with injury hotspot maps for these different cell sizes showed that smaller cell sizes tended to produce the better results. However, the differences in these results were marginal.

Table 7.9. (a) Area under the accuracy concentration curves, and (b) CPI values for Newcastle assault with injury KDE hotspot maps of different cell sizes (using a fixed bandwidth of 450 m). Values in bold relate to the largest area and largest CPI values.

(a)

Cell size (m)	**0.5% x 5%**	**1% x 10%**	**5% x 25%**	**10% x 50%**	**20% x 80%**	**100% x 100%**
30	**0.000246**	**0.000964**	0.012178	**0.046834**	**0.128185**	**0.901299**
60	0.000243	**0.000964**	0.012175	0.046573	0.125711	0.892741
90	0.000241	0.000955	**0.012180**	0.046573	0.127434	0.896380
120	0.000232	0.000949	0.012148	0.046270	0.125633	0.882987
150	0.000238	0.000959	0.012130	0.046513	0.125964	0.891772
Max values	*0.000250*	*0.001*	*0.0125*	*0.05*	*0.16*	*1*

(b)

Cell size (m)	**0.5% x 5% CPI**	**1% x 10% CPI**	**5% x 25% CPI**	**10% x 50% CPI**	**20% x 80% CPI**	**100% x 100% CPI**
30	**0.983**	**0.964**	0.974	**0.937**	**0.801**	**0.901**
60	0.971	**0.964**	0.974	0.931	0.786	0.893
90	0.966	0.955	**0.974**	0.931	0.796	0.896
120	0.929	0.949	0.972	0.925	0.785	0.883
150	0.952	0.959	0.970	0.930	0.787	0.892

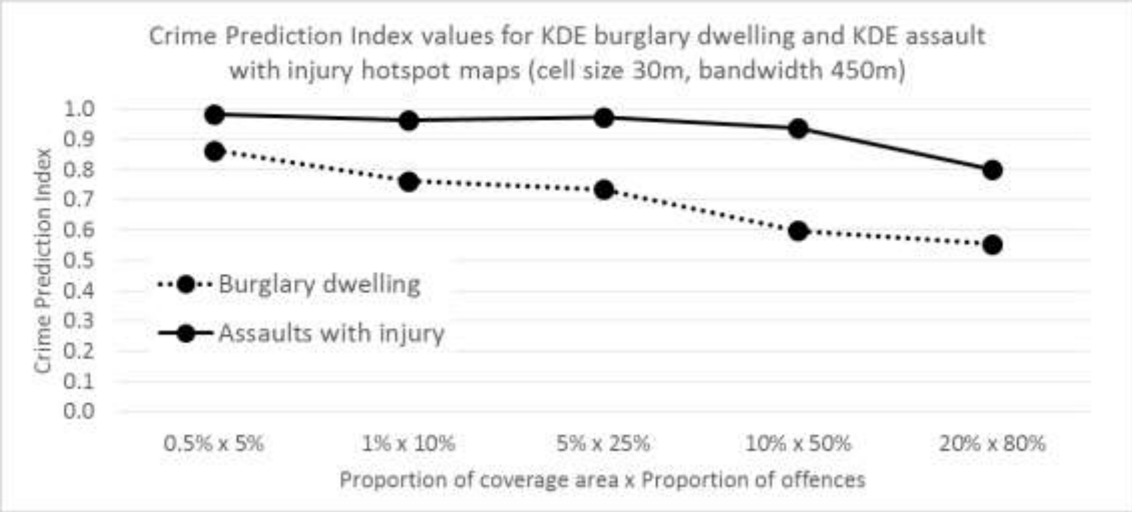

Figure 7.7. CPI values for KDE burglary dwelling and KDE assault with injury hotspot maps of Newcastle, created using a cell size of 30 m and bandwidth of 450 m

II. Bandwidth size results

Figure 7.8 shows the accuracy concentration curves for KDE burglary dwelling hotspot maps of Newcastle for different bandwidth sizes. Figure 7.8a shows a greater degree of variation in the lines when compared to the cell size charts for the same crime type. However, similar to the cell size charts (Figure 7.3), the points where each line tends to follow a flatter gradient reflects the coverage extent of KDE values and cells that have values of zero. For example, Table 7.10 shows the KDE values for burglary dwelling that were calculated using a bandwidth of 100 m extended to cover only 4% of the study area. This is reflected in the graph as this is the point where the curve flattens, and then takes a course that follows a trend of random variation. That is, from this point, every cell contains a value of zero and is not organised into any form of hierarchy, meaning that the chance of a crime being located in one cell is the same as it occurring in any other cell. Similarly, the point on the graph where the curve representing a bandwidth of 300 m flattens (Figure 7.8a) is at 39% of the proportion of the study area, reflecting the extent of KDE values for this bandwidth covering 39% of the study area (Table 7.10).

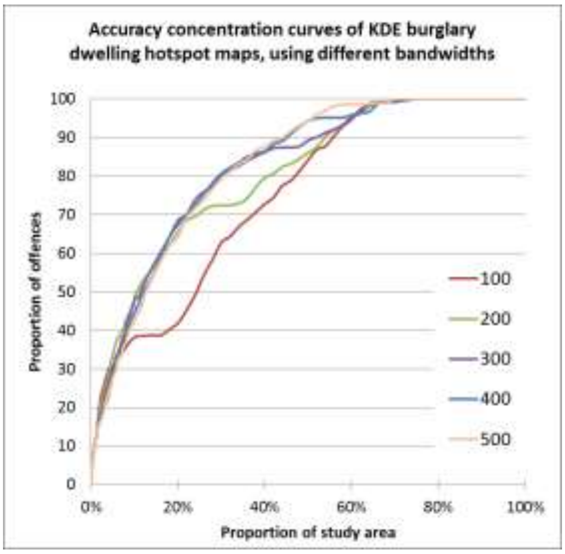

(a)

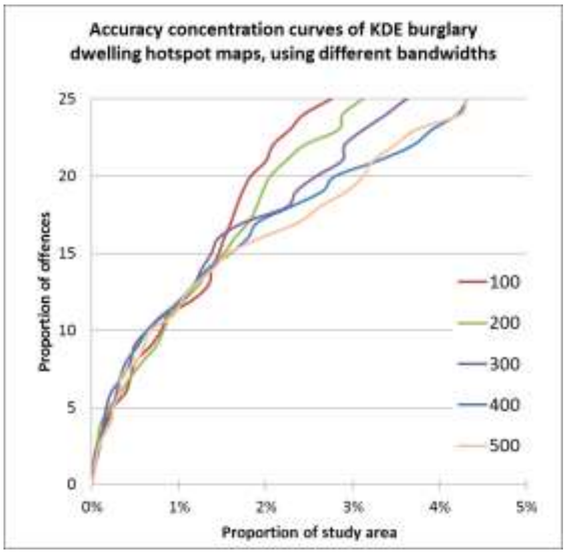

(b)

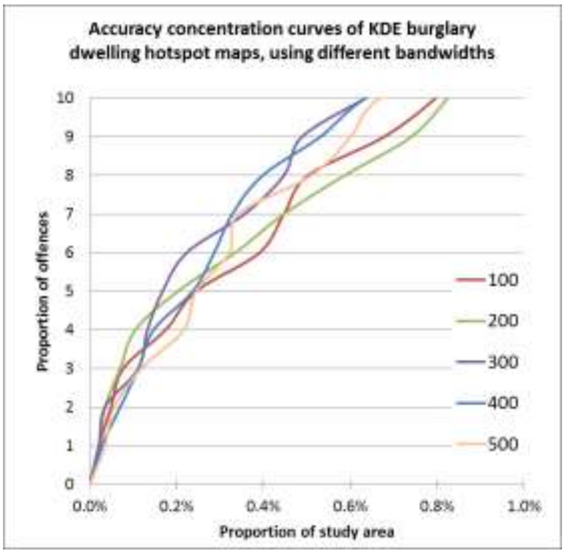

(c)

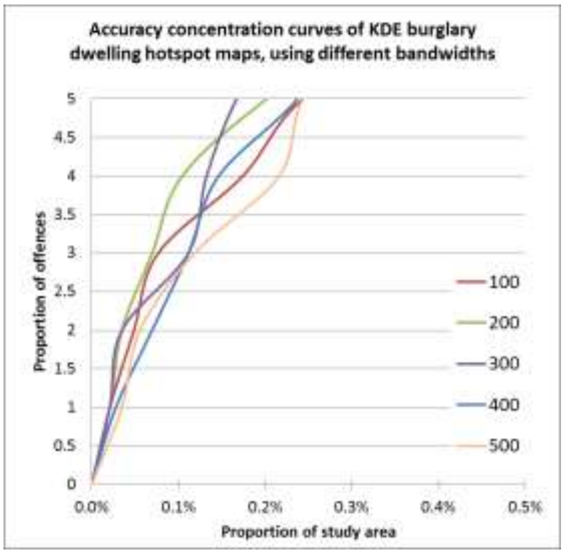

(d)

Figure 7.8. Accuracy concentration curves of Newcastle KDE burglary dwelling hotspot maps, generated using different bandwidth sizes (and using a fixed cell size of 90 m)

Examination of the sub-sections of the accuracy concentration curves up to 5%, 1% and 0.5% of the study area (Figure 7.8b, c, and d) does, though, show some variation in curve gradients that cannot be explained purely by the extent of KDE values. While the proportion of offences that are predicted changes little in a manner that is consistent between 0%-0.5% and 0%-1% of the study area, it appears that the smoothing created with larger bandwidths has more of an impact from approximately 1.5% of the study area coverage. That is, from this point, KDE burglary dwelling hotspot maps generated using smaller bandwidths produced better predictions of where crime may occur when compared to KDE hotspot maps using larger bandwidths. This is reflected in the findings presented in Table 7.10 that show the area that needed to be searched to identify different proportion levels of offences. At small coverage area levels (i.e., reflecting the highest KDE values), the differences were marginal. For a KDE map produced using a bandwidth of 300 m, 0.17% of the area needed to be searched to identify 5% of burglary dwellings, compared to a search of 0.24% of the area for KDE maps produced using bandwidths of 100 m, 400 m and 500 m. At larger coverage area levels, the differences were greater. For a KDE map produced using a bandwidth of 100 m, 2.75% of the area needed to be searched to identify 25% of burglary dwellings, compared to a search of 4.32% of the area for KDE maps produced using bandwidths of 400 m and 500 m.

Table 7.10. The proportion of the Newcastle study area searched across hotspots maps generated using KDE (six months of input data), relative to 5%, 10%, 25%, 50%, and 80% of burglary dwelling offences, for different bandwidth sizes. Values in bold represent the smallest area that was searched for identifying the relevant proportion of crime.

Bandwidth size (m)	**% of offences and the area searched (burglary dwelling)**					**% of study area with KDE values**
	5% of offences	**10% of offences**	**25% of offences**	**50% of offences**	**80% of offences**	
100	0.24%	0.80%	**2.75%**	26%	48%	4%
200	0.20%	0.82%	3.12%	**10.5%**	42%	23%
300	**0.17%**	**0.64%**	3.62%	11.5%	**30%**	39%
400	0.24%	**0.64%**	4.32%	12.%	**30%**	50%
500	0.24%	0.67%	4.32%	12%	32%	57%

Figures 7.9a-e show the smoothing effect in KDE hotspot mapping output created from using larger bandwidth sizes. These maps show that for small bandwidths, a large number of small areas are shown to be hotspots. As bandwidth size increased these small areas began to merge together to form fewer, albeit larger hotspots. This is as a direct result of a larger number of crime points across a larger local area being used in the kernel density calculation of crime points as the size of the bandwidth increases.

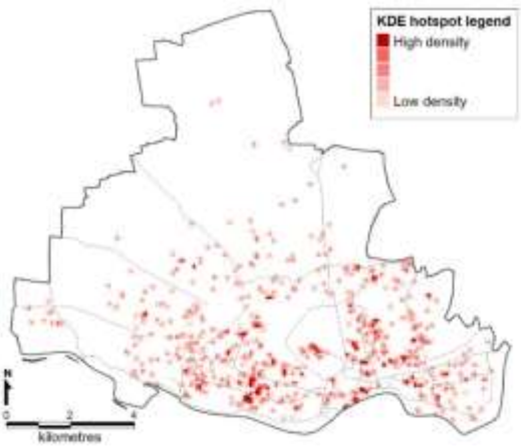

(a) 100 m bandwidth

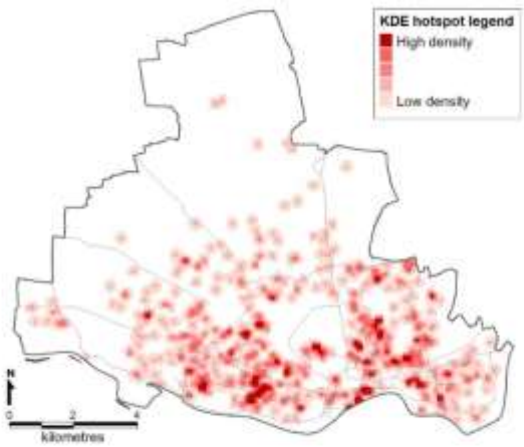

(b) 200 m bandwidth

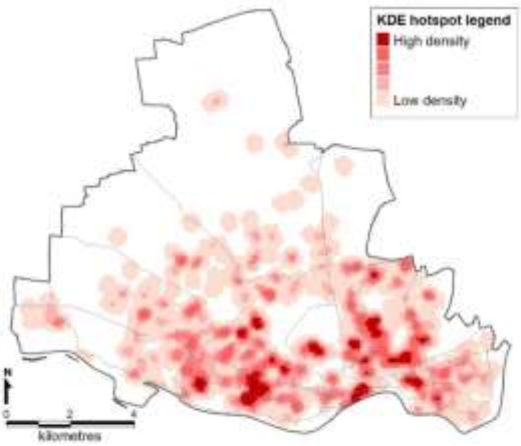

(c) 300 m bandwidth

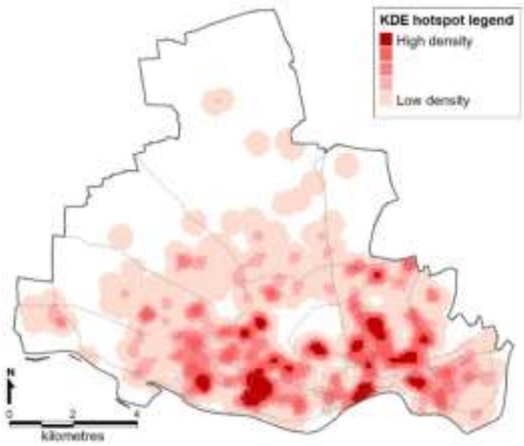

(d) 400 m bandwidth

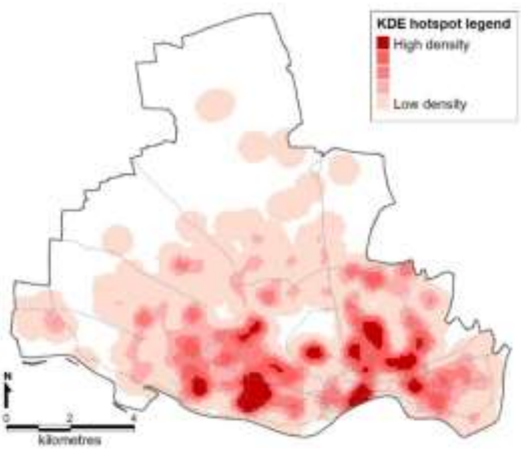

(e) 500 m bandwidth

Figure 7.9. KDE burglary dwelling hotspot maps of Newcastle, generated using six months of data, a fixed cell size of 90 m for bandwidth sizes of (a) 100 m, (b) 200 m, (c) 300 m, (d) 400 m, (e) 500 m

The variation in the prediction performance of KDE maps produced using different bandwidths is also shown in the area under the curve and CPI results (Table 7.11). These results show that KDE hotspot maps generated using smaller bandwidths produced better predictions - all the highest CPI values for the sub-sections under the accuracy concentration curves were for smaller bandwidth sizes. The CPI values also indicated the good level of prediction performance of KDE hotspot maps for burglary dwelling, but that the strength of the predictions reduced as the KDE values also reduced (and the size of area these KDE values represented increased). For example, KDE hotspot maps produced using a bandwidth of 200 m generated a CPI value of 0.867 for the section of the accuracy concentration curve presenting 0%-0.5% of the study area and up to 5% of burglary dwellings. The CPI value fell to 0.551 for the section of the accuracy concentration curve presenting 0%-20% of the study area and up to 80% of burglary dwellings.

Table 7.11. (a) Area under the accuracy concentration curves, and (b) CPI values for Newcastle burglary dwelling KDE hotspot maps of different bandwidth sizes (using a fixed cell size of 90 m). Values in bold relate to the largest area and largest CPI values.

(a)

Bandwidth size (m)	**0.5% x 5%**	**1% x 10%**	**5% x 25%**	**10% x 50%**	**20% x 80%**	**100% x 100%**
100	0.000207	0.000702	**0.009687**	0.027860	0.067119	0.752322
200	**0.000217**	0.000705	0.009426	**0.030776**	**0.088112**	0.804257
300	0.000212	**0.000770**	0.009222	0.028702	0.087085	0.821714
400	0.000203	0.000755	0.008757	0.027383	0.083872	0.824681
500	0.000195	0.000718	0.008590	0.026488	0.081396	**0.824780**
Max values	*0.00025*	*0.001*	*0.0125*	*0.05*	*0.16*	*1*

(b)

Bandwidth size (m)	**0.5% x 5% CPI**	**1% x 10% CPI**	**5% x 25% CPI**	**10% x 50% CPI**	**20% x 80% CPI**	**100% x 100% CPI**
100	0.828	0.702	**0.775**	0.557	0.419	0.752
200	**0.867**	0.705	0.754	**0.616**	**0.551**	0.804
300	0.846	**0.770**	0.738	0.574	0.544	0.822
400	0.810	0.755	0.701	0.548	0.524	0.825
500	0.781	0.718	0.687	0.530	0.509	**0.825**

The CPI values for the results shown in Table 7.11 reveal the danger of relying on just a value representing 100% of the coverage area and 100% of offences. For KDE maps of burglary dwelling in Newcastle, the CPI full coverage value was highest for a bandwidth of 500 m (0.825) and was lowest for KDE hotspot maps generated using a bandwidth of 100 m (0.752). This, however, is a reflection of the coverage extent of KDE values under the different bandwidth sizes. For smaller bandwidths, the extent of the study area that was covered in KDE values was significantly smaller. For example, for a 100 m bandwidth, only 4% of the study area was covered by KDE values (see Figure 7.9a), leaving the prediction of crime in the area not covered by KDE values (e.g., the remaining 96% of the coverage area of a 100 m bandwidth KDE hotspot map) to follow a trend of random variation. This is shown in the flattening of the accuracy concentration curve for a bandwidth of 100 m at the 4% of the study area point in Figure 7.8a. At larger bandwidths, KDE hotspot maps produced values for cells that covered a larger geographic extent. For example, for a KDE burglary dwelling hotspot map produced using a bandwidth of 500 m, KDE values were generated for 57% of the coverage area, leaving only 43% of the study area to follow a trend of random variation. Hence, by the 4% study area coverage mark, the KDE hotspot map produced using a bandwidth of 100 m had *done its job* in generating a good level of crime prediction, whereas the KDE hotspot map generated using a bandwidth of 500 m was continuing to *do its job* for up to 57% of the study area coverage. The area under the curve calculations and CPI values reflect these differences across the full coverage of the study area.

Figure 7.10 shows the accuracy concentration curves for KDE assault with injury hotspot maps for different bandwidths. These results again show that the main separation and flattening of each curve occurred when the KDE values reached the extent of their coverage (as shown in Table 7.12). For example, KDE assault with injury values calculated using a bandwidth of 100 m populated only 4% of the cells covering Newcastle. Figure 7.10a shows that at the 4% coverage level, the accuracy concentration curve for a 100 m bandwidth begins to flatten and from this point follows a trend of random variation. Similarly, the flattening point for the 200 m bandwidth accuracy concentration curve was at 21%, reflecting the geographic extent of KDE values for this bandwidth.

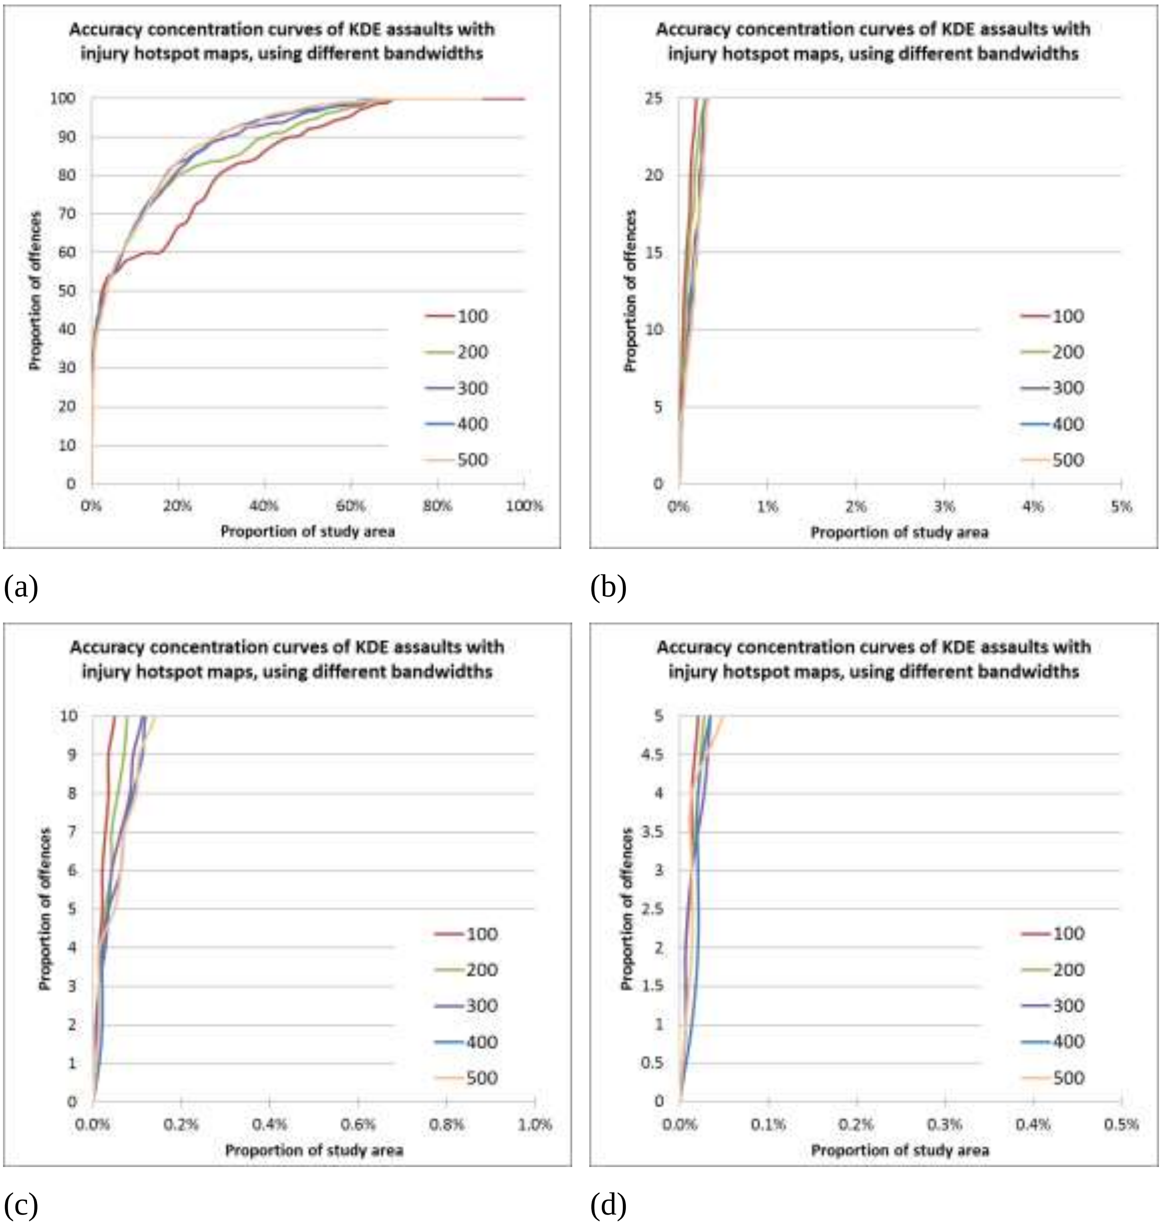

Figure 7.10. Accuracy concentration curves of Newcastle KDE assault with injury hotspot maps, generated using different bandwidth sizes (and using a fixed cell size of 90 m)

Examination of the sub-sections of the accuracy concentration curves up to 5%, 1%, and 0.5% of the study area (Figure 7.10b, c and d respectively) show again how much more vertical the curves for KDE maps of assault with injury were in comparison to the curves generated for burglary dwelling (Figure 7.8). The differences in the gradient of each curve for different bandwidths were marginal. However, a pattern that was evident was that the smaller bandwidths had steeper curve gradients than larger bandwidths reflecting the greater prediction performance of small bandwidths. This is also illustrated in Table 7.12, which shows that for each sub-section of the accuracy concentration curves, the area that needed to be searched for the upper limit volume of assaults to be identified was smallest for the lowest bandwidth size of 100 m. This is with the exception of the area that needed to be searched to identify 80% of assaults due to the KDE hotspot map for the 100 m bandwidth only covering 4% of the coverage area.

Table 7.12. The proportion of the Newcastle study area searched across hotspots maps generated using KDE (six months of input data), relative to 5%, 10%, 25%, 50%, and 80% of assault with injury offences, for different bandwidth sizes. Values in bold relate to the smallest area that was searched for identifying the relevant proportion of crime.

Bandwidth size (m)	**% of offences and the area searched (assaults with injury)**					**% of study area with KDE values**
	5% of offences	**10% of offences**	**25% of offences**	**50% of offences**	**80% of offences**	
100	**0.02**%	**0.05**%	**0.20%**	**2.50%**	30.0%	4%
200	0.03%	0.08%	0.29%	3.53%	20.0%	21%
300	0.03%	0.11%	0.31%	3.27%	20.0%	37%
400	0.03%	0.12%	0.33%	3.21%	**18.0%**	47%
500	0.05%	0.14%	0.32%	3.41%	**18.0%**	54%

Similar to the KDE hotspot maps for different bandwidth sizes of burglary dwelling, the KDE hotspots maps of assaults with injury in Figure 7.11a to e show the smoothing effect created from using larger bandwidth sizes. In contrast to the burglary dwelling KDE hotspot maps, there was one main location (covering Newcastle city centre) that was an assault with injury hotspot, with the spatial extent of this area changing less significantly than the merging of the large number of small hotspots that were observed in the burglary dwelling KDE maps. Some merging, caused by the smoothing effect of larger bandwidths did take place, but this marginal change in the spatial extent of this city centre hotspot was reflected in the more marginal change in the areas searched to identify 5%, 10%, 25% and 50% of assault with injury offences.

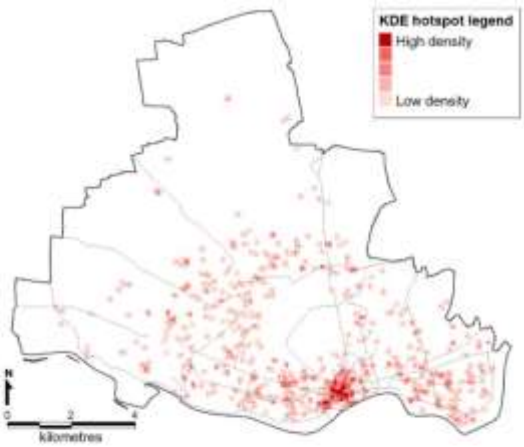

(a) 100 m bandwidth

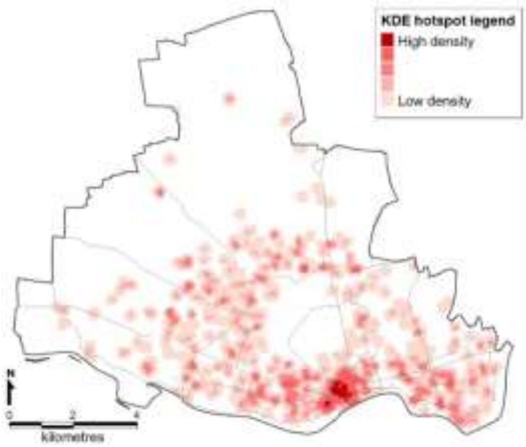

(b) 200 m bandwidth

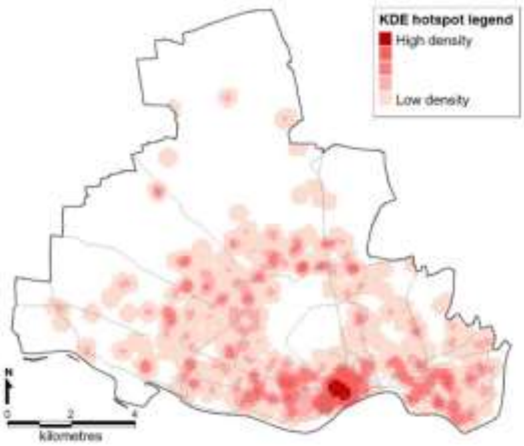

(c) 300 m bandwidth

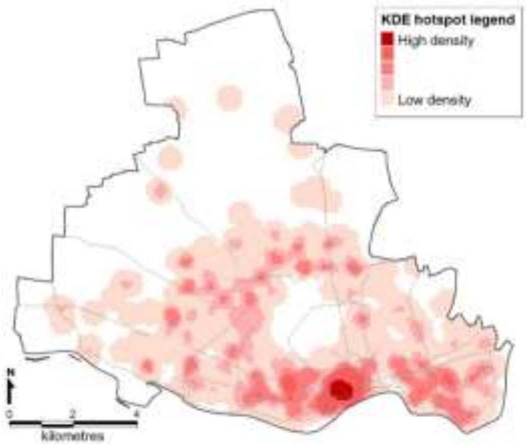

(d) 400 m bandwidth

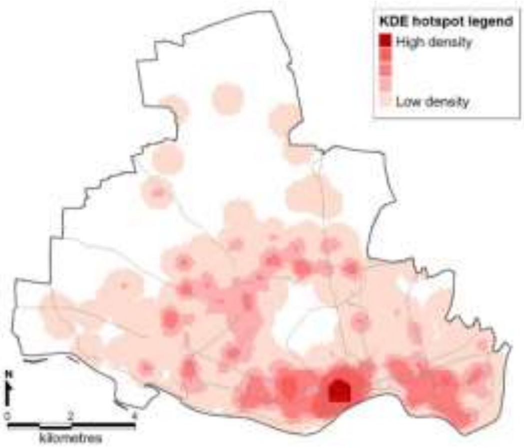

(e) 500 m bandwidth

Figure 7.11. KDE assault with injury hotspot maps of Newcastle, generated using six months of data, a fixed cell size of 90 m for bandwidth sizes of (a) 100 m, (b) 200 m, (c) 300 m, (d) 400 m, (e) 500 m.

The high degree of prediction performance in the KDE assault with injury hotspot maps is again shown in the CPI results (see Table 7.13). CPI values for the sub-sections of the accuracy concentration curves representing up to 10% of the study area and up to 50% of all offences were no lower than 0.95. The highest value of 0.985 was for the section of the accuracy concentration curve representing 0%-5% of the top KDE values covering the study area, for a bandwidth of 100 m. CPI values also showed little degradation across the small coverage area ranges (i.e., up to 5% of the study area), with the main degradations only occurring when the KDE values had reached the extent of their geographic coverage.

Table 7.13. (a) Area under the accuracy concentration curves, and (b) CPI values for Newcastle assault with injury KDE hotspot maps of different bandwidth sizes (using a fixed cell size of 90 m. Values in bold relate to the largest area and largest CPI values.

(a)

Bandwidth size (m)	0.5% x 5%	1% x 10%	5% x 25%	10% x 50%	20% x 80%	100% x 100%
100	**0.000245**	**0.000979**	**0.012310**	**0.047361**	0.112659	0.850360
200	0.000244	0.000968	0.012235	0.046888	0.126895	0.883171
300	0.000243	0.000957	0.012149	0.046627	0.126728	0.895842
400	0.000241	0.000949	0.012114	0.046368	**0.127309**	0.901356
500	0.000243	0.000950	0.012105	0.045896	0.126611	**0.902336**
Max values	*0.000250*	*0.001*	*0.0125*	*0.05*	*0.16*	*1*

(b)

Bandwidth size (m)	0.5% x 5% CPI	1% x 10% CPI	5% x 25% CPI	10% x 50% CPI	20% x 80% CPI	100% x 100% CPI
100	**0.979**	**0.979**	**0.985**	**0.947**	0.704	0.850
200	0.975	0.968	0.979	0.938	0.793	0.883
300	0.971	0.957	0.972	0.933	0.792	0.896
400	0.962	0.949	0.969	0.927	**0.796**	0.901
500	0.971	0.950	0.968	0.918	0.791	**0.902**

7.5. Interpretation and conclusions from research study 3

The aim of research study 3 was to test whether the technical parameters used in hotspot analysis techniques influence the techniques' spatial crime prediction performance (hypothesis 3). Following the results of research study 2, KDE was consistently found to generate hotspot maps with the highest level of prediction performance in comparison to the other commonly used hotspot analysis techniques, and therefore the parameter settings for KDE were the ones subject to analytical scrutiny in research study 3. The two main technical parameters that the analyst is required to enter for producing KDE mapping output are the cell size and the bandwidth size. To date, the advice given on cell size and bandwidth size to use has either not been appropriate for spatial crime analysis, or has been lacking in scientific merit. Research study 3 set out to establish if the selection of the cell size and bandwidth size selected by the crime analyst mattered, particularly in terms of the influence these parameters have on the prediction performance of KDE hotspot maps.

The results for study 3 show that cell size has little impact on the prediction performance of KDE hotspot maps. However, as small cell sizes offer the advantage in improving the resolution and visual appeal of the hotspot map (rather than the map looking pixelated) and as they also offer small marginal improvements in the prediction performance of KDE output, it is recommended that analysts use small cell sizes. A disadvantage in using small cell sizes is that extra processing is required for calculating KDE values and extra processing is also required in displaying the results. However, with increasing computer processing performance this disadvantage is becoming of less concern.

The choice of bandwidth size does, though, have an impact on the prediction performance of KDE hotspot maps, with smaller bandwidths producing the best results. The trade-off with this, however, is that KDE maps generated using small bandwidths can result in many small (very localised) hotspots being identified. The results for the Newcastle data illustrated this was more of a problem for burglary dwelling where the spatial distribution of this crime type, while concentrated at certain places, was more dispersed across the study area. The problem of many small hotspots was less of a problem for assault with injury KDE hotspot maps because only one main area of crime concentration was identified in and around the city centre of Newcastle. The practical implications of this trade-off between using small KDE bandwidths that generate maps of higher prediction accuracy but many different areas to target resources to, and the use of larger bandwidths that generate maps of lower prediction accuracy, but fewer areas to focus resource targeting is discussed further in chapter 12 (discussion and implications).

The results from the study 3 experiments and those from study 2 show that KDE hotspot mapping can produce output that offers a high level of prediction performance, particularly when care is taken in selecting cell size and bandwidth values. However, KDE is not without its weaknesses. In particular, a weakness with KDE mapping is that there is no objective method for determining the areas that are identified as *hot*. An alternative approach to identifying hotspots of crime is by using spatial significance mapping. Spatial significance mapping techniques offer the ability to identify areas that can be statistically determined as *hot*. Research study 4 will examine the application of spatial significance mapping techniques for identifying hotspots of crime and whether these perform better than KDE in predicting where crime is likely to occur.

8. Research study 4: Improving hotspot analysis using spatial significance mapping

8.1. Introduction

This research study aims to test whether spatial significance mapping methods provide an improved means of predicting where crime is likely to occur in comparison to common hotspot mapping techniques, and removes the ambiguity of defining areas that are *hot* (hypothesis 4).

Study 2 of the research has shown that kernel density estimation is the best of the common hotspot analysis techniques for predicting where crime is likely to occur, with part 3 then showing the influence that KDE parameter settings, namely the cell size and bandwidth size, can have on the prediction performance of KDE hotspot maps. The main findings were that cell size had little impact on the prediction performance of KDE hotspot mapping output, but that small bandwidth sizes tend to produce KDE mapping output that perform better in predicting where crime is likely to occur. With the use of the Crime Prediction Index, the research also showed the very good predictive performance of KDE maps for certain crime types, particularly theft from the person and assaults.

However, KDE is not without its faults. The main weakness with KDE is that the selection of values that are used to determine thematic thresholds is left to the choice and convention of the analyst, rather than the selection being determined by some statistical process. That is, the analyst subjectively determines the numerical thematic threshold for what is *hot*, rather than being guided by a robust systematic process. Choosing the numerical distinctions between thematic classes for KDE hotspot maps is very much left to trial and error, experimentation, experience, or whatever suits the analyst's circumstance. The selection of a desired thematic classification method is also a straightforward functional choice in a GIS, with the crime analyst being able to choose between equal ranges, equal count, natural break, standard deviation or quantile classification options at a single click of a button. Figure 8.1a and 8.1b illustrate examples of this problem with KDE mapping output: both figures are based on the same crime data, yet Figure 8.1b gives the impression that the crime problem is much worse. The difference between the maps is purely due to the difference in the values that are used to determine the thematic classes between high density levels of crime (i.e., the areas that can be referred to as hotspots) and low density levels of crime. Figure 8.1b was generated using the equal count thematic classification method[14], whereas Figure 8.1a was generated using the equal ranges[15] thematic classification method. The result of using two different thematic classification methods meant that the threshold value representing *high density* in Figure 8.1b was lower than the value used to represent *high density* in Figure 8.1a.

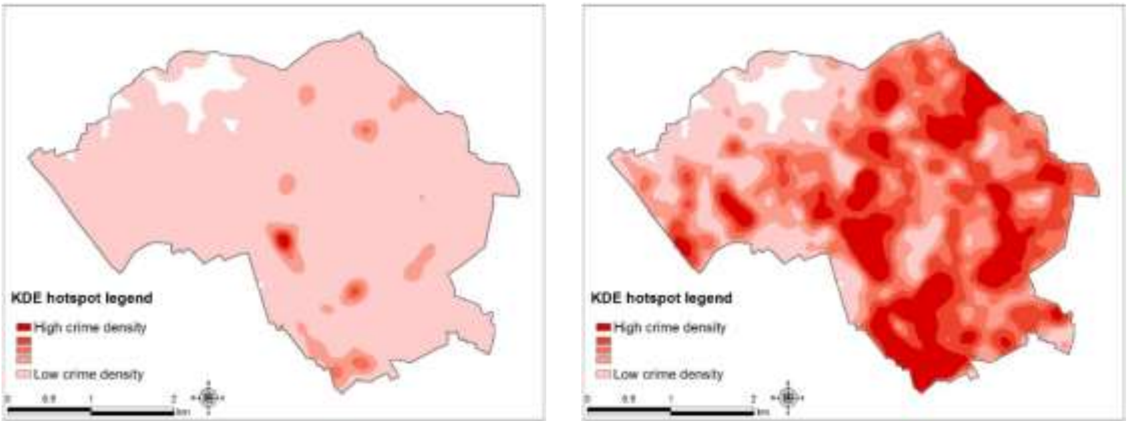

(a) (b)

Figure 8.1. Camden/Islington KDE hotspot maps produced using the same data, but different thematic classification methods: (a) equal ranges, and (b) equal count

Spatial significance mapping offers potential in removing this ambiguity in defining hotspots by using the principles of statistical significance testing. Although the nature of the KDE technique and the Gi* statistic are different – KDE seeks to redistribute occurrences under a continuous global volume preserving structure, whereas Gi* is a local test of spatial association – the popularity of the KDE technique for crime hotspot analysis and the increasing interest in the GI* statistic as an application for crime hotspot analysis are why an assessment between the two are sought. If, in a statistical sense, the areas that are hotspots can be determined using Gi*, it is also possible that these identified areas offer a more accurate means of determining where crime is likely to occur in the future.

8.2. Chapter aims and structure

This research study examines whether hotspot analysis can be improved by using spatial significance mapping. The improvements that are examined are whether the subjective selection of the area determined as *hot* on a KDE hotspot map can be removed using the

[14] The equal count method organises data into a user defined number of categories where the number of events in each thematic class is equal.
[15] The equal ranges method organises data into a user defined number of categories where the range in the values across each thematic class is equal.

principles of statistical significance testing, and whether this then improves the prediction performance of the mapping output.

The chapter begins by building on the introduction of Local Moran's I, Local Geary's C and the Gi* statistic from chapters 2 and 3 by describing further the technical features of these statistics. This focus on the technical features of these statistics includes determining if one is more suitable than the others for hotspot analysis of crime data, and examining the input parameter requirements for each statistic that the analyst is required to determine. A set of experiments are then performed that test spatial significance mapping output for its prediction performance. These experiments broadly followed the method from research parts 2 and 3. The process that was used for conducting these experiments is described in section 8.4, including a description of the prediction measures that were used for assessing the prediction performance of the spatial significance mapping outputs.

The results from these spatial significance mapping experiments are then presented and interpreted to inform how they influence subsequent research parts.

8.3. Spatial significance mapping techniques for identifying hotspots

Mapping techniques that examine the distribution and concentration of events, and employ a test for spatial statistical significance include Local Moran's I, Local Geary's C and the Getis-Ord Gi* statistic. These techniques are described as local indicators of spatial association (LISA statistics) (Anselin, 1995) in that they identify the association between a single point and its neighbours up to a specified distance from the point. LISA statistics can therefore provide an indication of the extent to which the value of an observation is similar or different to its neighbouring observations, but they do so in different ways. Chapters 2 and 3 provided an introduction to these techniques. In the following section the technical detail of each is described in more detail.

8.3.1. Local Moran's I and Local Geary's C

Local Moran's I and Local Geary's C can be used to determine for each observation the extent to which there is spatial association; that is, clustering of similar values around that observation. An exact test of significance is not possible for Local Moran's I nor Local Geary's C because the distribution of each statistic is not known, and the expected value of I or C and the variance of I or C are very complicated mathematically. Instead, high positive or high negative standardised scores of I (termed Z(Ii)) and C (termed Z(Ci)) are taken as indicators of similarity or dissimilarity (Levine, 2010). A high positive value for the Z(li)s or Z(Ci)s indicates the spatial clustering of similar values (either high or low), while a high negative score indicates clustering of dissimilar values (a high area surrounded by low areas or a low area surrounded by areas with high values). The higher the value of Z(Ii) or Z(Ci), the more the observation is similar (positive) or dissimilar (negative) to its neighbours. This means that the Local Moran's I statistic and the Local Geary's C statistics can potentially be good indicators of the presence of crime hotspots that, in a spatial sense, are statistically significant. It does, though, require the analyst to review the values for each of the geographic units that are of interest to see if they represent a statistically significant hotspot or a statistically significant coldspot (i.e., observations with low levels of crime surrounded by other observations with low levels of crime).

8.3.2. Gi* statistic

The Gi* statistic is a measure that compares local averages to global averages to identify those areas that are significantly different in comparison to what is generally observed across the whole study area (Ord and Getis, 1995). The Gi* statistical outputs are standardised Z scores. Z scores indicate the place of a particular value in a dataset relative to the mean, standardised with respect to the standard deviation. When Z equals zero, the observation is equivalent to the sample (data) mean. When Z is less than zero the observation is a value less than the mean. When Z is greater than zero, the observation is a value greater than the mean. Z scores are used extensively for determining whether an observation is statistically significant. When a normal distribution of the observations is assumed, Z score values of 1.96, 2.576 and 3.291 are used to determine 95%, 99% and 99.9% statistical significance levels respectively. That is, if a Gi* result for an observation is greater than or equal to 1.96, the observation is determined as being statistically significant to 95%. This means the Gi* statistic can potentially be used to distinguish those areas that, in a statistical sense, are *hot* from those that are *not hot* (i.e., where there is significant spatial concentration in the distribution of crime data across a study area in comparison to the spatial distribution of crime that is not significant).

When using Local Moran's I and Local Geary's C, the analyst is required to review the results for each of the geographic units that are significant to determine if they represent a statistically significant hotspot or a statistically significant coldspot. In the case of Gi* results, this is not required because areas where low levels of crime cluster are represented by negative Z scores and areas where high levels of crime cluster are represented by positive scores. Because of this, it was decided that the Gi* statistic would be the preferred method of further study. The Gi* statistic was also chosen for further study because it is a technique that has been receiving increasing interest for improving crime hotspot analysis (Chainey and Ratcliffe, 2005; Eck et al., 2005; Hart and Zandbergen, 2012), but to date no systematic study has been completed that compares whether the output the GI* statistic produces performs any better in predicting spatial patterns of crime when compared to KDE hotspot analysis output.

1	1	1	5	0	0	0	1	0	0	0	0	0	0	3	2
0	3	0	0	6	1	0	1	1	0	0	0	0	0	1	3
5	0	0	0	0	1	9	5	0	0	3	0	0	1	0	1
1	4	0	2	0	5	0	0	0	1	1	0	0	0	0	2
1	0	2	3	0	3	6	0	1	2	0	0	0	1	5	0
3	5	0	4	0	0	0	2	1	2	1	1	0	0	1	0
0	0	1	1	8	1	6	6	2	2	0	1	0	1	2	0
0	2	2	2	4	6	12	9	2	2	3	6	2	0	0	2
0	0	3	8	5	1	2	1	1	1	5	0	0	0	2	2
1	2	4	2	1	0	1	0	1	3	0	0	2	3	0	2
4	4	1	0	0	1	1	1	0	2	1	4	2	1	6	4
1	1	0	0	0	0	0	0	1	4	5	2	2	6	1	0
0	0	0	2	0	0	1	0	2	6	1	3	0	4	0	0
1	1	0	0	0	0	0	0	0	2	0	0	13	0	0	0
0	0	0	1	1	0	0	0	1	4	6	0	2	0	0	0
0	8	2	6	0	0	0	4	3	1	4	7	0	0	0	0

Figure 8.2. 16x16 matrix of cells, each holding a value representing the number of crimes within their respective areas

The statistical function of the Gi* statistic is best described with an example. Figure 8.2 shows a 16x16 grid lattice. Each cell can be identified by its centroid point and each cell has a count of the number of crimes in that area. Consider the point positioned in the eighth row of the eighth column in Figure 8.2. This point (*i*) has the value 9. The null hypothesis states that site *i* is not the centre of a group of unusually high values centred on *i* and its surrounding cells (its neighbours *j)*, such that the sum of values at all the *j* sites within a radius *d* of *i* is not more (or less) than one would expect by chance given all the values in the entire study area (both within and beyond the distance *d*). If local spatial association exists, it will be exhibited by a spatial clustering of high or low values. When there is a clustering of high values, the Gi* values will be positive, while a clustering of low values will yield a negative Gi* value (Anselin, 1995; Chainey and Ratcliffe, 2005; Getis and Ord, 1996; Ord and Getis, 1995).

The Gi* statistic typically requires the analyst to determine the value for one parameter – the lag distance. A lag distance is the radius of a moving three-dimensional sphere that visits each grid cell. In many ways it is similar to the bandwidth measure that is used for kernel density estimation, except that a suitable value is slightly easier to determine. Typically, our interest is in the local association between a geographic unit and its nearest neighbours, and most usually the immediate surrounding neighbours. The lag distance should, therefore, be at least the distance of the radius from each cell's centroid that has a coverage that will consider all of its immediate neighbours (Chainey and Ratcliffe, 2005; Eck et al., 2005). The distance between the grid cells in the example 16x16 cell grid in Figure 8.2 is 125 m. A suitable lag distance to apply would be 178 m as this is the furthest distance to include the eight immediate neighbours (the cells on the diagonal i.e., $\sqrt{(125^2 + 125^2)}$). As the lag distance is based on the calculation of the cell size, the cell size, therefore, will have an influence on the Gi* output that is generated. The current research study examines the influence that different cell sizes and, hence, different lag distance have on the Gi* outputs.

In some circumstances (depending on the software that is used) the analyst may also be required to determine the number of lags they wish to apply. A lag of 1 refers to the lag distance as above. Increasing the number of lags, as in this example, allows different multiples of 178 m to be explored. For example, a lag of 3 for the 16x16 matrix will calculate Gi* values within a distance (*d*) of 178 m, 356 m, and 534 m.

Table 8.1 shows the Gi* statistics for the cell positioned in the eighth row of the eighth column with the crime count value of 9. Higher positive values of Gi* indicate greater clustering of high values. At a lag of 1 the Gi* statistic is positive. This high and positive Gi* statistic indicates that there is positive local spatial association between this cell and its neighbouring cells. That is, this particular cell, and the eight cells immediately surrounding it and forming a 3 x 3 matrix, sum to a relatively high total count that is greater than the global average. The Gi* statistic at a lag of 2 is also positive but not as high as the Gi* statistic at a lag of 1. The Gi* value remains similar for a lag of 3. For those cells that have low values (i.e., low crime counts) and which are also surrounded by cells of low values, the Gi* statistic would be negative (Chainey and Ratcliffe, 2005; Getis and Ord, 1996; Ord and Getis, 1995).

Table 8.1. Gi* statistics for the cell positioned in the eighth row of the eighth column of the 16x16 grid (Figure 8.2), for lags 1-3

Gi* at 0 to <= 178m (Lag 1)	**Gi* at 0 to <= 356m (Lag 2)**	**Gi* at 0 to <= 534m (Lag 3)**
4.2	2.4	2.3

In the first part of this section it was explained that as the Gi* calculations are Z scores, statistical significance levels can be determined using the standard measures of 1.96, 2.576 and 3.291 for 95%, 99% and 99.9% confidence levels respectively. However, an issue with the application of this approach for spatial data is that it assumes that the determination of each Z score value is independent. With spatial analysis, treatment of statistical tests is usually (as in this case) not only applied to the observation but also to a defined set of neighbours. This means that each observation is not used just once, but a multiple number of times because it is included as a neighbour for many other observations. This issue of multiple testing results in the increased probability of incorrectly rejecting the null hypothesis – a Type I error. Failure to account for these effects results in over identification of clusters using Gi* (Anselin, 1995; Getis and Aldstadt, 2004; Ord and Getis, 1995).

A number of approaches exist to correct for multiple testing. False discovery rate (FDR) controlling procedures seek to reduce the expected proportion of Type I errors (i.e., false discoveries) when the issue of multiple testing is present (for an example of FDR applied to geographic data see Caldas de Castro and Singer, 2006). However, FDR approaches tend not to be that conservative, resulting in a higher chance of Type I errors remaining. Familywise error rate (FWER) approaches seek to reduce the probability of even one false discovery, or Type I errors, and, therefore, tend to be more conservative than FDR approaches (Charlton and Brunsdon, 2011).

Ord and Getis (1995) suggest a Bonferroni correction (a FWER procedure) to help address the problem of multiple testing associated with the application of the Gi* statistic. The Bonferroni procedure has the immediate advantage of being straight-forward to calculate (therefore, allowing for easy replication by practitioners). However, the Bonferroni correction can be very conservative if there are a large number of tests. Also, the Bonferroni correction procedure comes at the cost of increasing the probability of producing Type II errors (i.e., the failure of detecting whether an effect is present) and can result in a lack of statistical power (i.e., the probability that the null hypothesis is rejected when clustering actually occurs) (Charlton and Brunsdon, 2011). In the current research, as the number of surrounding areas used in the computation of each Gi* statistic is small (i.e., only the eight immediate neighbours are used), although conservative, the Bonferroni correction is considered to be accurate in these circumstances (Rogerson, 2004). In addition, a conservative determination of statistical significance thresholds is considered a positive feature for spatial crime analysis because it may be useful in identifying small, specific areas for focused policing and crime prevention targeting. If the use of the Bonferroni correction produces results that appear to be far too conservative (i.e., it is difficult to discern clusters of crime from Bonferroni corrected Gi* mapping output), alternative correction procedures will then be considered.

In practice, a Bonferroni correction to determine the Z score representing a 95% significance threshold involves a two-step process:

- First, the analyst must calculate the number of cells representing 5% of the study area: 0.05/number of cells covering the study area (0.05 is used as it is equivalent to 95%)
- Second, the analyst needs to determine the Z score that represents this corrected percentile. The most practical way to do this is for the analyst to use a percentile to Z score calculator to determine the Z score value that represents a 95% statistical significance threshold. Many such tools exist online. One such online tool is available at http://www.measuringusability.com/zcalcp.php

The process above can then be repeated for determining Bonferroni corrected 99% and 99.9% statistical significance threshold values, substituting the initial calculation using 0.01/number of cells covering the study area (for determining the 99% statistical significance threshold value) and 0.001/number of cells (for determining the 99% statistical significance threshold value).

For the illustration sample of a 16x16 grid matrix, a Bonferroni corrected 95% statistical significance threshold value was calculated to be 3.725 (i.e., the result of 0.05/256 (16x16 = 256), which is then converted to a standardised Z-score). That is, under a Bonferroni correction, any cell in the 16x16 matrix that has a Gi* value that is greater than or equal to 3.725 is statistically significant to 95% level of confidence. The Gi* value for the example cell (eighth row, eighth column in Figure 8.2), at a lag of 1 is 4.2. This means that there is significant statistical evidence (to 95%) of crimes clustering at the cell positioned in the eighth row of the eighth column of the 16x16 grid, up to a lag 1 distance. The Gi* values for a lag of 2 and lag of 3 were not statistically significant.

The results from the example using the 16x16 grid matrix suggest that the Gi* statistic offers a distinct advantage over techniques such as KDE for hotspot analysis by identifying those areas that can be statistically defined as *hot*. That is, if we decided to apply a 95% significance threshold to determine areas that are *hot*, any cells with Gi* values that were determined to be significant to at least 95% would be classed as hotspots. Alternatively, different thematic classes could be determined with respect to the different levels of spatial significance. This means that the application of the Gi* statistic could help remove the ambiguity of defining areas that are *hot*. It may also suggest that this additional statistical rigour to hotspot analysis using the Gi* statistic may result in better predictions of where crime may occur in the future.

A problem sometimes associated with KDE is that in some locations it may smooth across or into areas where there is no recorded crime because it is not constrained to the high detail of the underlying geography of the crime point distribution. An advantage of the Gi* statistic is that it compares local averages against global averages, and can exclude in an analysis areas where it is impossible for certain crimes to occur (e.g., for burglary dwelling this would include parkland and rivers or any other areas where there was no residential housing). Areas that can be excluded can be identified in a GIS and the grid cells that cover these areas could be extracted from the full grid coverage so they do not influence the global average. Another problem associated with KDE is that it may smooth away the peaks in areas where the specified bandwidth aggregates high values with neighbouring low values, and as a result overlooks areas where the concentration of crime is tightly compact. The Gi* statistic may potentially not overlook these areas, and, therefore, may be a more accurate technique for identifying these highly spatially concentrated crime hotspots.

The Gi* statistic, therefore, appears to offer the potential to more accurately and robustly (in a statistical sense) identify where hotspots occur, and in turn where crime may occur in the future. That is, it may be better than common hotspot mapping methods such as KDE used in policing and public safety for determining where crime is predicted to occur. To date, very little use has been made of the Gi* statistic in spatial crime analysis, and no assessment has been made of whether it performs better than KDE in predicting spatial patterns of crime. This research study aims to fill that gap.

8.4. Method

This research study involved conducting a number of experiments to determine if the Gi* statistic could improve the spatial categorisation of areas determined as *hot*, and how Gi* hotspot analysis output compared to KDE in its spatial crime prediction performance. The research study also tested whether the cell size, and hence the lag distance, influenced the prediction performance of Gi* hotspot analysis output.

The following sections explain how Gi* hotspot analysis output was compared to KDE mapping output, the crime types, data input period and data output periods that were chosen for analysis, and how prediction performance was measured.

8.4.1. Comparing the hotspot analysis outputs produced using Gi* against those produced using KDE for three different thematic classification methods

The first experiment involved producing two examples of Gi* hotspot analysis output and comparing these to KDE hotspot maps to identify if Gi* output could be used to define the spatial extent of areas that are determined as *hot*. A Gi* hotspot map using three months of burglary dwelling data was produced for the Camden/Islington study area. This Gi* hotspot map was then compared against KDE burglary dwelling hotspot maps (using the same data) produced using the equal interval thematic classification method, natural breaks thematic classification method, and quantile thematic classification method. A Gi* hotspot map using three months of theft from the person data were produced for Newcastle and compared against KDE theft from the person maps using the same three thematic classification methods that were used in the Camden/Islington examples. In each case, KDE maps were produced using a cell size of 20 m and a bandwidth size of 300 m. These parameters were chosen for consistency between the two study areas and because they provided representative KDE output.

No guidance exists on the appropriate cell size to use for generating Gi* output. Indeed, it is possible that no strict mathematical process can be devised for selecting the cell size because the cell size is dependent upon the spatial structure of the study area and of the phenomenon under investigation. After some experimentation with different cell sizes that balanced between the generation of counts for each cell and the visual output of the results, a cell size of 150 m was chosen for the experiments for both Camden/Islington and Newcastle. The interest in the Gi* statistic is for identifying local clusters, and, therefore, it was logical to only include the immediate nearest neighbours in the Gi* calculation of the local average (i.e., the eight cells surrounding the central cell for which the Gi* statistic is calculated). This meant the lag distance could be determined as a function of the cell size (as explained in section 8.3.2). The lag distance for both areas was calculated as 213 m. A Bonferroni correction was applied to determine the Gi* statistical significance levels.

8.4.2. Crime types, data input periods and data output periods

Data for Camden/Islington and Newcastle on burglary dwelling, theft from the person, theft from vehicles and theft of vehicles, and additionally Newcastle assault with injury data were used for determining the prediction performance of Gi* hotspot analysis. The data input periods and data output periods used were for the most part those described in the study areas and crime data section of the method in chapter 4. The exception was that only the 1st January 2010 was used as a measurement date for Camden/Islington, rather than the two measurement dates (i.e., the 1st January 2010 and the 11th March 2010) used previously. This was because the previous studies in the current research showed no difference in the results for different measurement dates.

8.4.3. Measuring the prediction performance of Gi* hotspots

A combination of PAI measures, accuracy concentration curves, area under the curve and Crime Prediction Index measures were used for assessing the prediction performance of Gi* hotspot maps. PAI measures were aggregated and averaged for the periods of input data and for the periods of output data. The standard deviation of the PAI for each crime type across the different data input and data output periods was also calculated. This would allow for comparisons between the PAI results for KDE and those for Gi*. Gi* statistical significance levels were calculated using a Bonferroni correction procedure. Gi* PAI values were calculated in each experiment for each of the 95%, 99% and 99.9%

significance levels in order to determine whether PAI values were different between these *hotspot* categories.

These first experiments that used the PAI as a prediction measure used a cell size of 150 m for generating Gi* hotspot maps for both study areas (the reasons for this are described in section 8.3.2.). This meant a lag distance of 213 m was used for both study areas.

A second set of experiments were then conducted to examine the influence that cell size had on the prediction performance of Gi* hotspot analysis output. This involved using cell sizes of 50 m. 100 m, 150 m, 200 m and 300 m for both Camden/Islington and Newcastle experiments. These experiments used accuracy concentration curves, the area under the curve and the Crime Prediction Index, alongside the PAI, to examine the influence that cell size had on Gi* output. The experiments were conducted using data on burglary dwellings and theft from the person for Camden/Islington, and burglary dwellings, theft from the person and assault with injury for Newcastle. For these experiments, three months of input data were used for creating Camden/Islington and Newcastle Gi* hotspot analysis outputs, and these findings were compared with crime patterns in the three months following the measurement date for both study areas. The GI* results were then compared against KDE hotspot analysis output for the same crime types and the same data input and output periods.

8.5. Results

8.5.1. Comparing the hotspot analysis outputs produced using Gi* against those produced using KDE for three different thematic classification methods

The first experiment examined how Gi* hotspot analysis output compared to KDE output, and whether the use of the Gi* statistic helped to overcome the ambiguity in defining areas that were hotspots. Figure 8.3 shows examples of KDE burglary dwelling hotspot maps for Camden/Islington generated using three standard thematic classification methods. Figure 8.4 shows examples of KDE hotspot maps of theft from the person for Newcastle using the same three thematic classification methods.

The KDE examples in Figures 8.3 and 8.4 for both study areas show different areas were identified as hotspots (with the exception of Figure 8.4a and 8.4b which show only a marginal difference due to the intense clustering of theft from the person in Newcastle). This finding highlights the ambiguous nature to defining hotspots using KDE. Figure

8.3d and Figure 8.4d show examples of Gi* hotspot analysis output, with each thematic class determined by the standardised Z scores that represent (Bonferroni corrected) 95%, 99% and 99.9% levels of spatial statistical significance.

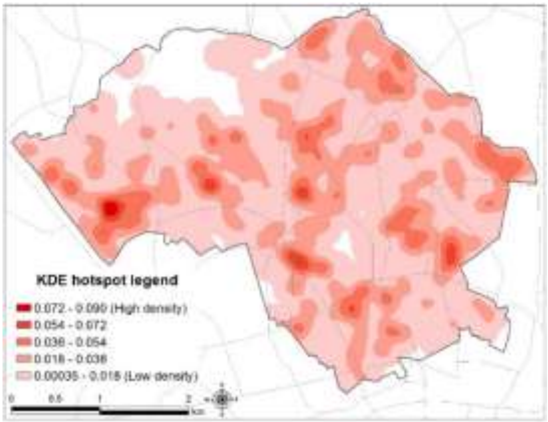

(a) KDE: equal interval thematic classification

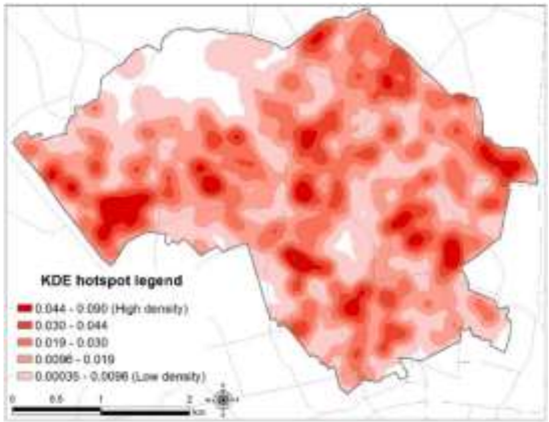

(b) KDE: natural breaks thematic classification

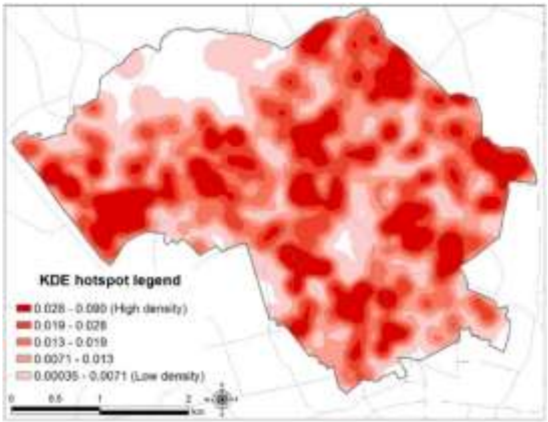

(c) KDE: quantile thematic classification

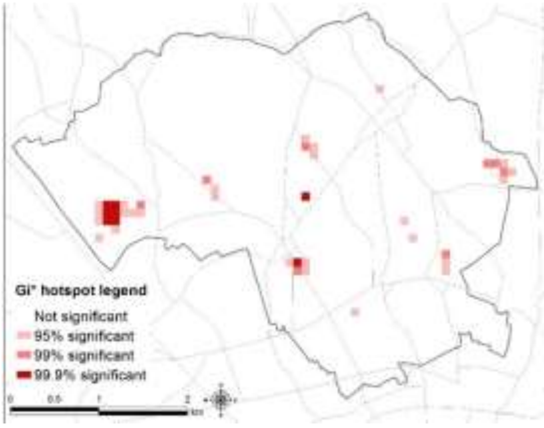

(d) Gi* hotspots

Figure 8.3. A comparison of Camden/Islington KDE burglary dwelling hotspot maps (using a cell size of 20 m and bandwidth of 300 m) using (a) equal interval thematic classification, (b) natural breaks thematic classification, (c) quantile thematic classification, and (d) a Gi* burglary dwelling map of (Bonferroni corrected) spatial significance (using a cell size of 150 m).

While the areas that are identified on the Gi* hotspot maps in Figures 8.3 and 8.4 are comparable to the KDE generated hotspots, four main differences can be observed. In the first instance, areas on the Gi* maps that are not determined to be *hot* are not included. While only the high density areas could have been shown on each of the KDE maps, this would still have identified differences in the areas identified as *hot* in each of the KDE maps.

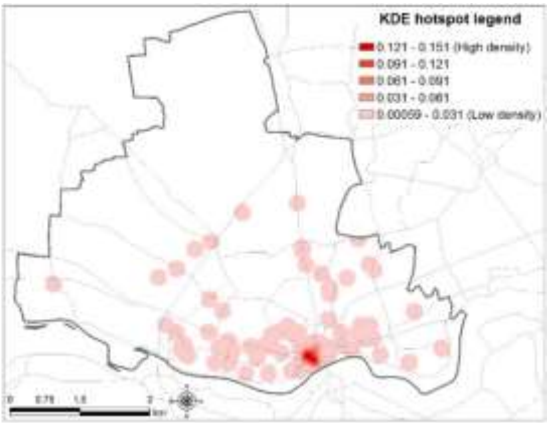

(a) KDE: equal interval thematic classification

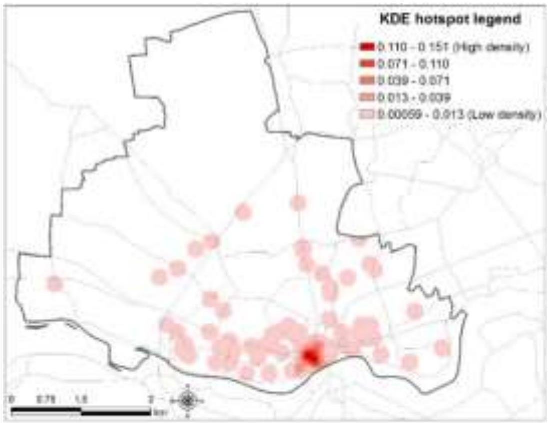

(b) KDE: natural breaks thematic classification

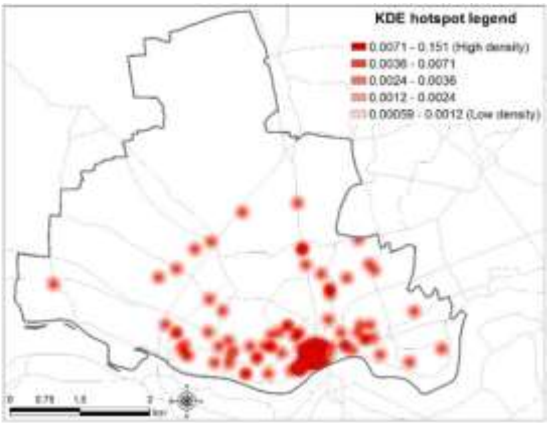

(c) KDE: quantile thematic classification

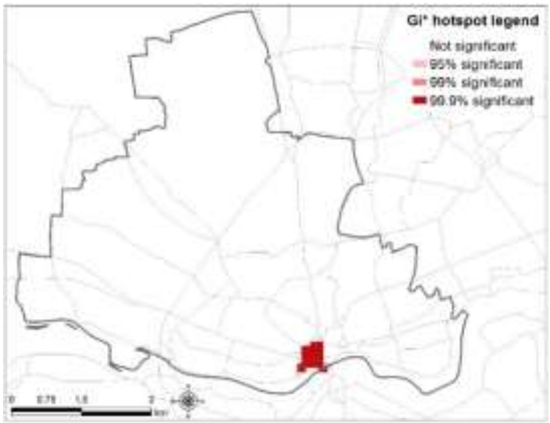

(d) Gi* hotspots

Figure 8.4. A comparison of Newcastle KDE theft from the person hotspot maps (using a cell size of 20 m and bandwidth of 300 m) using (a) equal interval thematic classification, (b) natural breaks thematic classification, (c) quantile thematic classification, and (d) a Gi* burglary dwelling map of (Bonferroni corrected) spatial significance (using a cell size of 150 m).

Secondly, the identification of *hot* areas in the Gi* maps were systematically classified using statistical principles (i.e., statistical significance), rather than through a subjective process (i.e., the analyst's choice of thematic classification method) as in the case of KDE hotspot mapping output. For example, by observing the top thematic class for each KDE burglary dwelling hotspot map in Figure 8.3, the threshold value that determines each category of high density ranges from 0.028 crimes per square km to 0.072 crimes per square km. This variation in KDE threshold values for defining the areas that are *hot* explains why the maps show differences in the areal extent of hotspots. Gi* hotspot analysis helps remove the ambiguity in defining hotspots, albeit requiring the analyst to determine if a highly focused approach is needed by defining the threshold for what is *hot* using a 99.9% statistical significance, rather than a 95% definitional threshold. That is, although the determination of the areas that are *hot* (using the Gi* statistic) follow a statistically defined classification process, the determination of *hot* areas is still somewhat subjective as it is based on the analyst's decision on which statistical significance threshold and correction procedure to apply.

Thirdly, in places, KDE appears to *smooth out* hotspots in areas where the crime concentration is very compact (Figures 8.3a and 8.3b), or over-exaggerates these compact concentrations of crime when the quantile range method is used (Figure 8.3c and Figure 8.4c). The Gi* outputs in Figure 8.3 appear to be better at identifying small compact concentrations of crime (overcoming the smoothing nature of KDE) but not to the extent of visually exaggerating the area of influence these compact concentrations have at the local level.

Finally, the KDE maps do, though, have the advantage of looking more visually appealing, a strength that has been noted by several commentators (Chainey et al., 2008; Chainey and Ratcliffe, 2005; Eck et al., 2005). In comparison to the KDE maps, the Gi* hotspot analysis output appears more blocky, influenced in part by the choice of a 150 m cell size.

8.5.2. Examining the prediction performance of Gi* hotspots

Table 8.2 lists the PAI results for each crime type for the Camden/Islington and Newcastle study areas. The PAI results were calculated by averaging the individual PAI values for each input dataset and each output dataset. A statistical significance threshold of 95% was used to define the areas that were *hot* from Gi* hotspot analysis output. These results show that the Gi* PAI results for Camden/Islington were consistently higher than the equivalent KDE PAI values across all crime types. However, PAI Gi* values in the Newcastle study area were lower than the equivalent KDE PAI values, albeit marginally, across all crime types (with the exception of theft of vehicles). These results suggest that although Gi* hotspot analysis helps remove much of the ambiguity in defining the areas that are hotspots, it does so without necessarily improving the prediction performance of the hotspot analysis output.

The finding that Gi* hotspot analysis may not necessary improve the prediction performance of hotspot analysis output is examined in further detail by reviewing the results in Table 8.3. Previous analysis in research study 2 that used KDE hotspot analysis outputs (generated from three months of input data) to give some indication of how many crimes are predicted using KDE hotspots analysis, employed a process of using the top 3% of KDE values to control for the size of the hotspots. This process was repeated for the Gi* statistic, but rather than using the top 3% of Gi* values, the proportion of the area representing a Bonferroni corrected 95% statistical significance threshold was used. While this resulted in the sizes of areas defined as *hot* differing between crime types for the two study areas, this process was considered to be useful to help indicate how many crimes were predicted using this statistical hotspot definition process in comparison to previous KDE findings.

Table 8.2. PAI values of KDE and Gi* hotspot maps for burglary dwelling, theft from the person, theft from vehicles, and theft of vehicles for Camden/Islington and Newcastle, and assault with injury for Newcastle. Values in bold relate to the highest PAI values for that study area.

	Camden/Islington			**Newcastle**		
Crime type	**KDE average PAI**	**GI* average PAI**	*Standard deviation of Gi* PAI*	**KDE average PAI**	**Gi* average PAI**	*Standard deviation of Gi* PAI*
Burglary dwelling	1.56	**2.05**	0.18	**7.5**	7.6	0.96
Theft from the person	3.24	**4.35**	0.42	**46.3**	45.8	8.57
Theft from Vehicle	1.89	**3.10**	0.34	**6.0**	5.9	0.53
Theft of Vehicle	1.53	**2.45**	0.31	0.5	**1.3**	0.08
Assault with injury	**-**		-	**40.3**	37.7	4.53

Table 8.3a shows that the size of the areas the Gi* method defined as hot in Camden/Islington were similar in three out of four crime types to the top 3% KDE controlled areas, therefore, providing some direct comparisons. In three of the four crime types, KDE PAI values were marginally higher than Gi* PAI values, with this being reflected in the smaller additional proportion of crimes that were predicted in the KDE hotspots. The one result where the Gi* prediction performance was better than that for KDE was for theft of vehicles where the area that the Gi* analysis determined to be hot was much smaller (1.4%) than the controlled 3% area used in the KDE PAI calculations. Similar results were found for Newcastle (see Table 8.3b), where the KDE PAI results for three of the five crime types were higher. The exceptions were burglary dwelling where the Gi* hotspot analysis performed better than KDE, and theft of vehicles where

the Gi* analysis of three months of crime data could not predict where the 11 vehicle thefts took place in the month following the measurement date. Of note, though, was that the areas identified as *hot* from the Gi* analysis were smaller than the KDE 3% controlled areas, and significantly so in the case of theft from the person (0.6% of the study area), theft of vehicles (0.8%) and assaults with injury (1%).

Table 8.3. PAI and actual crimes predicted using KDE and Gi* (cell size 150 m; 95% Bonferroni corrected). Each KDE and GI* output was generated using three months of crime data to determine where crimes in the next month may occur. Values in bold relate to the highest PAI values. *The Newcastle theft of vehicles hotspot maps were generated using six months of input data because clustering was not evident from three months of input data.

(a) Camden/Islington

Crime type	**Gi* PAI**	**KDE PAI**	**Crimes committed in January 2010**	**% of area determined *hot* using Gi*** (Bonferroni corrected p=0.05)	**n of crimes in Gi* hotspots**	**% of crimes in Gi* hotspots**	**% of crimes in KDE hotspots (3% of area)**
Burglary dwelling	2.3	**2.8**	470	3.0%	32	7%	8%
Theft from person	6.5	**6.6**	460	3.1%	93	20%	20%
Theft from vehicle	3.4	**4.0**	985	3.0%	100	10%	12%
Theft of vehicle	**4.4**	3.3	307	1.4%	19	6%	10%

(b) Newcastle

Crime type	**Gi* PAI**	**KDE PAI**	**Crimes committed in April 2010**	**% of area determined *hot using Gi**** (Bonferroni corrected p=0.05)	**n of crimes in Gi* hotspots**	**% of crimes in Gi* hotspots**	**% of crimes in KDE hotspots (3% of area)**
Burglary dwelling	**7.2**	6.7	130	2.4%	22	17%	16%
Theft from person	101.9	**389.1**	60	0.6%	34	58%	68%
Theft from vehicle	8.9	**9.4**	189	2.1%	35	19%	25%
Theft of vehicle*	0.0	0.0	11	0.8%	0	0%	0%
Assault with injury	37.9	**79.2**	154	1.0%	56	36%	44%

8.5.3. Significance level and prediction performance

Table 8.4 lists the Gi* PAI values for the three spatial statistical significance levels, for different crime types, for the two study areas. These results consistently show that the prediction performance of Gi* hotspot output improved as the significance level was set higher. For example, Table 8.4a shows the PAI values for burglary dwelling improved from 2.1 at the 95% significance level, to 2.3 at 99% and 2.7 at 99.9%. This trend is most likely to be due to smaller hotspot areas being identified.

Table 8.4. PAI values for Gi* hotspot analysis (cell size 150 m; Bonferroni corrected) of 95%, 99% and 99.9% statistical significance levels for (a) Camden/Islington, and (b) Newcastle dataset. Values in bold relate to the highest PAI values.

(a) Camden/Islington

	Average PAI for Gi*			
	Burglary dwelling	**Theft from the person**	**Theft from vehicle**	**Theft of vehicle**
95%	2.1	4.4	3.0	2.5
99%	2.3	5.0	3.0	3.0
99.9%	**2.7**	**5.9**	**3.3**	**3.4**

(b) Newcastle

	Average PAI for Gi*				
	Burglary dwelling	**Theft from the person**	**Theft from vehicle**	**Theft of vehicle**	**Assault with injury**
95%	7.5	110.8	9.1	0	45.3
99%	7.6	123.5	9.2	0	56.9
99.9%	**7.8**	**149.2**	**9.4**	**0**	**59.2**

8.5.4. An analysis of accuracy concentration curves, the area under the curve, CPI values and the influence of cell size on Gi* hotspot analysis

So far, the results comparing Gi* to KDE suggest that while Gi* hotspot analysis allows for a more systematic identification of *hot* areas using statistical confidence levels, the Gi* outputs did not necessarily offer an improvement in the prediction performance of KDE hotspot maps. These results have been determined using the PAI. The next set of experiments involved comparing Gi* and KDE hotspot analysis outputs using accuracy concentration curves, the area under the curve and the Crime Prediction Index. The experiments also tested the impact that cell size (and hence lag distance) has on the prediction performance of Gi* hotspot analysis outputs.

I. Camden/Islington burglary dwelling

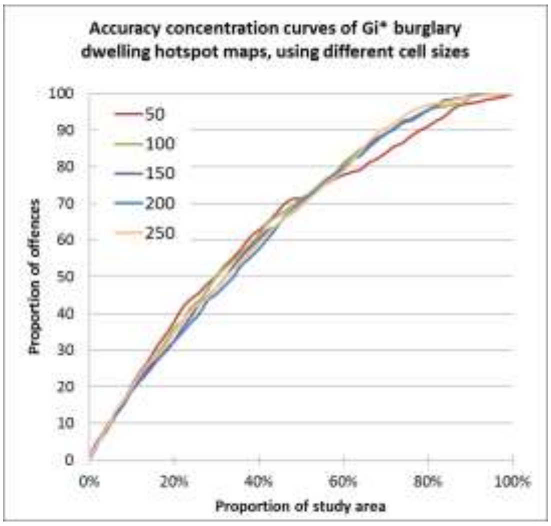

(a)

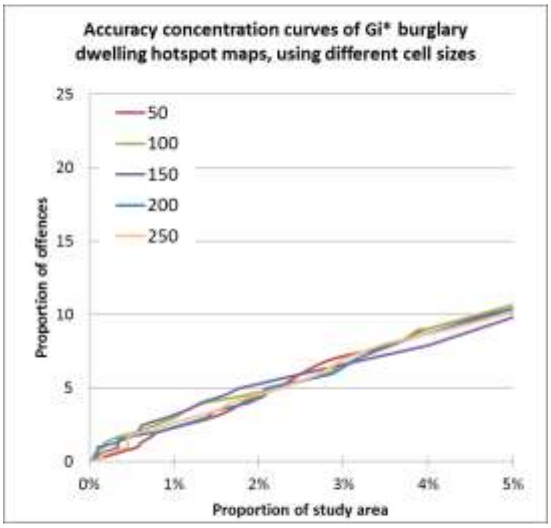

(b)

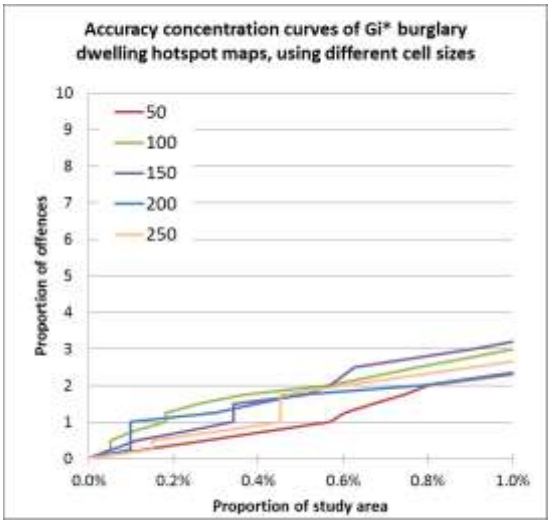

(c)

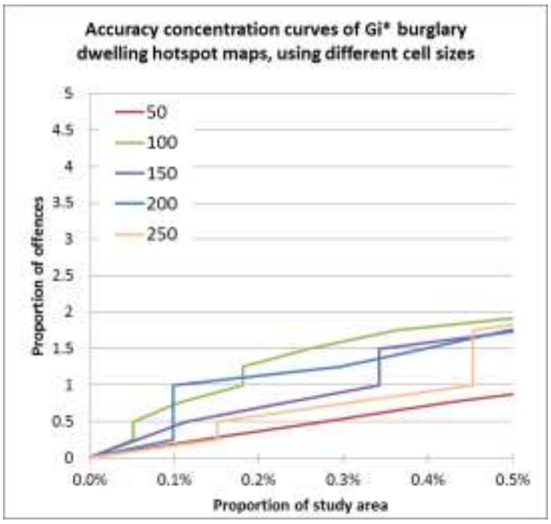

(d)

Figure 8.5. Accuracy concentration curves of Camden/Islington Gi* burglary dwelling hotspot analysis for different cell sizes. The results are shown for (a) the full study area, and sub-sections relating to (b) 5% of the study area, (c) 1% of the study area and (d) 0.5% of the study area.

Figure 8.5 shows the accuracy concentration curves for Gi* burglary dwelling hotspot maps for Camden/Islington. The first observation was that there was little difference in the prediction performance of Gi* hotspot analysis outputs for different cell sizes. The step changes in certain curves shown in Figure 8.5c and d was more noticeable for Gi* output generated using larger cell sizes. It is unlikely that these step changes indicate any optimum points of prediction performance for the crime phenomenon under investigation, and is more likely to be an artefact of the spatial distribution of the crime points and cell sizes. These sections of steep gradients on these curves occurred when a single individual cell captured a large volume of offences, with the cumulative addition of this offence volume having a knock on effect for the next immediate cells that cover the study area. This type of scenario is more likely when using larger cells than smaller cells. For example, Table 8.5 lists the Gi* hotspot analysis output for the five cells with the top Gi* values, generated using a cell size of 200 m. This shows the number of offences for 0.25%, 0.5%, 0.75% and 1% of all offences was captured in just one cell, representing 0.1% of the study area.

Table 8.5. (a) Cells with the top Gi* values (Camden/Islington burglary dwelling) and (b) their effect on measures representing the proportion of the study area searched to identify the relevant proportion of offences

(a)

Cell Gi* rank	**Gi* Z score**	**n of crimes in cell during month following measurement date**
1	7.062	15
2	6.357	2
3	6.004	3
4	5.652	4
5	5.300	3

(b)

Proportion of offences	**n of offences equivalent to proportion of offences**	**Proportion of study area searched to identify proportion of offences**
0.25	4	0.10% (i.e., one cell)
0.5	7	0.10% (i.e., one cell)
0.75	11	0.10% (i.e., one cell)
1	15	0.10% (i.e., one cell)
1.25	19	0.30% (i.e., three cells)
1.5	22	0.40% (i.e., four cells)
1.75	26	0.50% (i.e., five cells)

The marginal differences between the prediction performance of Gi* hotspot mapping outputs generated using different cell sizes are also shown in Table 8.6. The proportion of crime predicted differed little between Gi* hotspot maps of different cell sizes, although in general, smaller cell sizes performed better than larger cell sizes.

Also of note from Table 8.6 were the small sizes of areas that were determined as *hot* using spatial statistical significance measures. For example, the hot areas identified on the Gi* burglary dwelling hotpot map that used a cell size of 100 m represented only

1.8%, 0.5% and 0.4% of the study area at the 95%, 99% and 99.9% confidence levels. This finding suggests that the prediction performance of Gi* hotspot analysis should primarily be analysed for these statistically determined hotspot areas, rather than assessing the prediction performance of Gi* hotspot analysis across coverage areas that include areas where the concentration of crime is not statistically significant. Indeed, of the results shown for the areas that would need to be searched to find the proportion of offences by the categories listed in Table 8.6, only three of the twenty five results related to search areas that were statistically significant. For example, the Gi* burglary dwelling hotspot mapping output that was generated using a cell size of 150 m identified hotspots to 95% significance that covered an area representing 2.4% of the study area. The area searched to identify 5% of offences was 1.8%, and therefore all the values for this area were significant to at least 95%. This identification of relatively small areas that are hotspots highlights the conservative nature to Bonferroni correction.

Table 8.6. The proportion of the study area searched across maps generated using Gi* (using three months of input data) relative to 5%, 10%, 25%, 50%, and 80% of burglary dwelling offences, for different cell sizes. The proportion of the study area with Gi* values relative to 95%, 99% and 99.9% is also listed. Values in bold represent the smallest areas searched. Where the area searched was not greater than the proportion of the area that was statistically significant, this is highlighted as follows: *** 99.9% ** 99% * 95%.

Cell size (m)	**% of offences and area searched (burglary dwelling)**					**% of study area with significant Gi* values**		
	5% of offences	**10% of offences**	**25% of offences**	**50% of offences**	**80% of offences**	**95%**	**99%**	**99.9%**
50	2.2%	5.2%	**12.7%**	**29.6%**	65.0%	1.1%	0.5%	0.2%
100	2.2%	**4.5%**	13.5%	29.9%	52.9%	1.8%	0.5%	0.4%
150	**1.8%***	5.0%	14.4%	32.3%	53.6%	2.4%	0.9%	0.5%
200	2.1%*	4.9%	13.8%	34.1%	**52.8%**	3.6%	0.7%	0.6%
250	2.3%*	5.0%	13.3%	33.1%	54.1%	2.9%	0.4%	0.4%

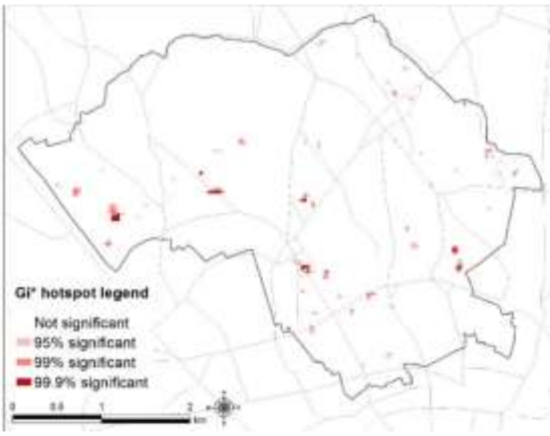

(a) Cell size: 50 m

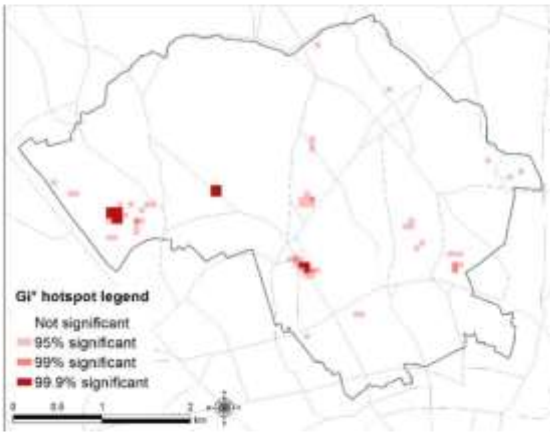

(b) Cell size: 100 m

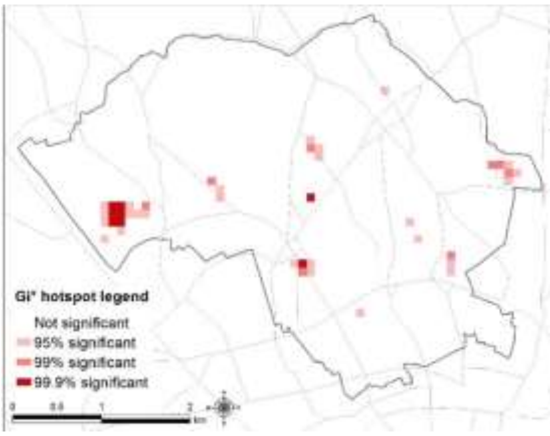

(c) Cell size: 150 m

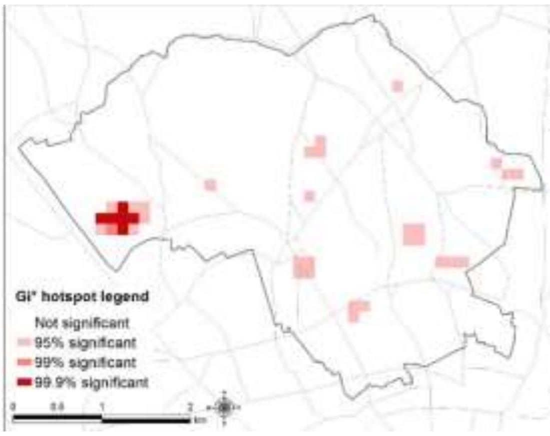

(d) Cell size: 200 m

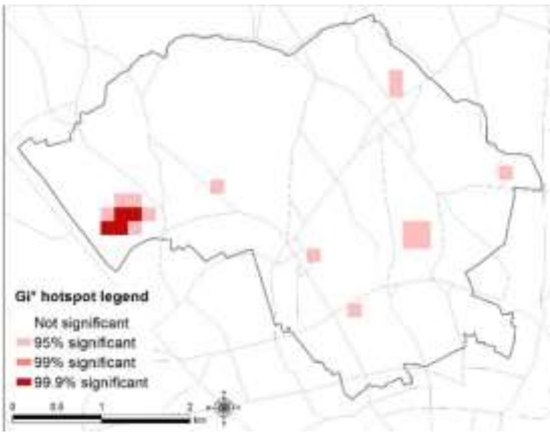

(e) Cell size: 250 m

Figure 8.6. Camden/Islington Gi* burglary dwelling hotspot maps produced using three months of input data and cell sizes of (a) 50 m, (b) 100 m, (c) 150 m, (d) 200 m, and (e) 250 m

The Bonferroni corrected Gi* burglary dwelling hotspot maps in Figure 8.6 (produced using three months of input data) also highlight the conservative nature of this technique. In each map, between 1.1% (cell size of 50 m) to 3.6% (cell size of 200 m) of the study area was identified as *hot*. Table 8.7 lists the Bonferroni corrected Z scores for 95%, 99%

and 99.9% significance levels and highlights the higher Z score significance threshold values this correction determines in comparison to uncorrected Z scores. For example, rather than using 1.96 to determine 95% significance for a cell size of 100 m, the Bonferroni corrected Z score for this significance level was 3.828. The maps in Figure 8.6 also show that as cell size increases, the number of hotspots begin to reduce. This is as a result of very small (e.g., single cell) hotspots losing their significance as cell sizes increase by taking into consideration the volume of crime across a larger local neighbourhood. For example, in Figure 8.6e seven hotspots are identified compared to the 52 hotspots (some of which are single cells) identified in Figure 8.6a.

Table 8.7. Bonferroni corrected Z score spatial statistical significance values for Camden/Islington Gi* maps of different cell sizes. Uncorrected Z scores: 95%: 1.960; 99%: 2.576; 99.9%: 3.291

Cell sizes (m)	**n of cells covering study area**	**Bonferroni corrected Z scores**		
		95%	**99%**	**99.9%**
50	15066	4.150	4.972	5.401
100	3872	3.828	4.702	5.152
150	1753	3.628	4.537	5.001
200	1010	3.483	4.419	4.489
250	662	3.369	4.327	4.810

Table 8.8 shows again that, when considering the area under the curve and CPI values for the different sub-sections of the Camden/Islington burglary dwelling accuracy concentration curves (for consistency, using the same sub-sections as the analysis in study 3), no apparent pattern consistency in prediction performance emerges between small cell sizes and larger cell sizes. Of note, however, are the low CPI values in comparison to previous analysis of KDE outputs (study 3, chapter 7). Of more interest would be how these area under the curve and CPI measures compare at the 95%, 99% and 99.9% significance levels rather than these arbitrary selected sub-sections, in comparison to KDE values. This is examined in section 8.5.5.

Table 8.8. Camden/Islington Gi* burglary dwelling hotspot analysis of (a) the area under accuracy concentration curves, and (b) CPI values, for different cell sizes. Values in bold relate to the largest area under the curve and the largest CPI values.

(a)

Cell size (m)	0.5% x 5%	1% x 10%	5% x 25%	10% x 50%	20% x 80%	100% x 100%
50	0.000023	0.000108	0.002797	0.010287	**0.039342**	0.645000
100	**0.000063**	0.000185	**0.002901**	0.010276	0.037582	**0.652613**
150	0.000046	0.000177	0.002828	0.010096	0.035721	0.644417
200	0.000054	0.000155	0.002769	0.009964	0.036068	0.640565
250	0.000033	**0.000204**	0.002798	**0.010498**	0.037784	0.648155
Max values	*0.00025*	*0.001*	*0.0125*	*0.05*	*0.16*	*1*

(b)

Cell size (m)	0.5% x 5% CPI	1% x 10% CPI	5% x 25% CPI	10% x 50% CPI	20% x 80% CPI	100% x 100% CPI
50	0.090	0.107	0.224	0.206	**0.246**	0.645
100	**0.254**	0.185	**0.232**	0.206	0.235	**0.653**
150	0.182	0.177	0.226	0.202	0.223	0.644
200	0.216	0.155	0.222	0.199	0.225	0.641
250	0.132	**0.204**	0.224	**0.210**	0.236	0.648

II. Camden/Islington theft from the person

Figure 8.7 shows the accuracy concentration curves for Camden/Islington Gi* theft from the person hotspot maps. These results show little variation in the gradients of the curves up to the point representing 40% of the study area. After this point the Gi* hotspot analysis output that was generated using a cell size of 50 m flattens in comparison to outputs generated using other cell sizes. Figures 8.7b, c and d, and Table 8.9 show that the Gi* hotspot analysis outputs that were generated using small cell sizes produced the best predictions. Table 8.9 also shows that the proportion of the areas searched to identify up to 10% of all theft from the person offences were for areas that the Gi* mapping output identified to be statistically significant to 99.9%, for all cell sizes.

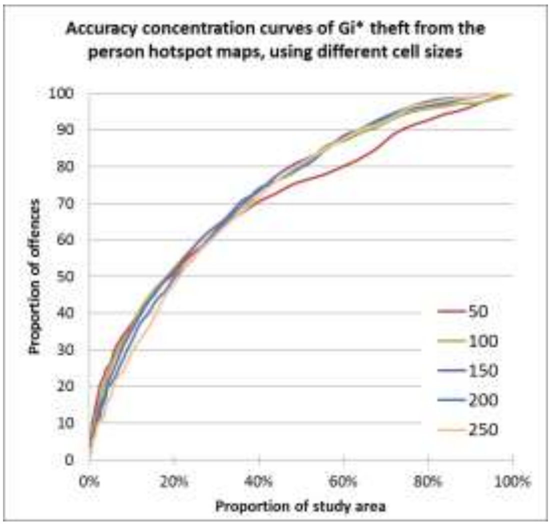

(a)

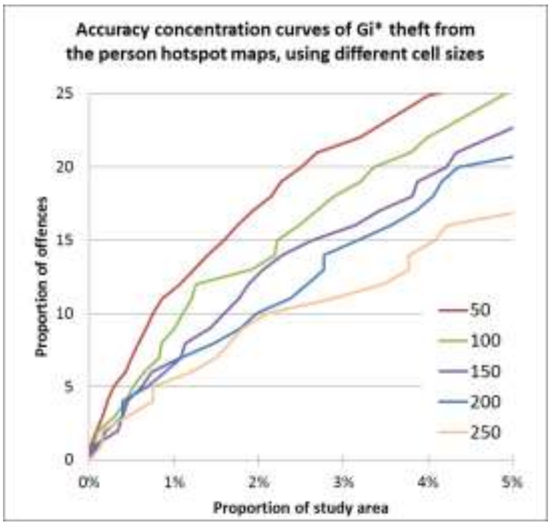

(b)

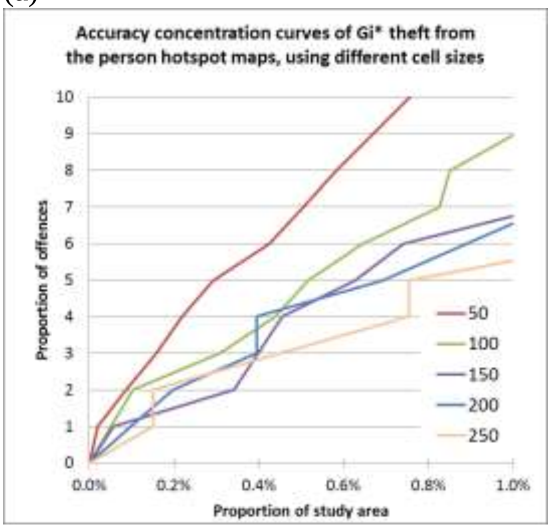

(c)

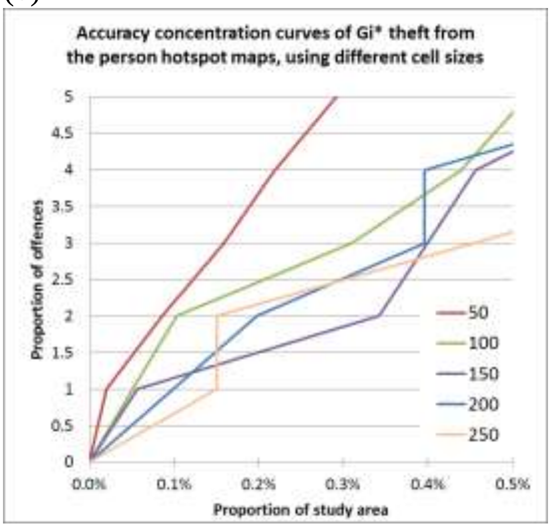

(d)

Figure 8.7. Accuracy concentration curves of Camden/Islington Gi* theft from the person hotspot analysis for different cell sizes. The results are shown for (a) the full study area, and sub-sections relating to (b) 5% of the study area, (c) 1% of the study area and (d) 0.5% of the study area.

Figure 8.8 shows the Bonferroni corrected Gi* theft from the person hotspot maps (produced using three months of input data). In each map, between 1.8% (cell size of 50 m) to 4.8% (cell size of 250 m) of the study area was identified as *hot*. The maps also show that as cell size increases, the number of hotspots begin to reduce, caused by very small (e.g., single cell) hotspots losing their significance as larger cell sizes begin to include the volume of crime across a larger local neighbourhood.

Table 8.9. The proportion of the Camden/Islington study area searched across maps generated using Gi* (three months of input data) relative to 5%, 10%, 25%, 50%, and 80% of theft from the person offences, for different cell sizes. The proportion of study area with Gi* values relative to 95%, 99% and 99.9% is also listed. Values in bold represent the smallest areas searched. Where the area searched was not greater than the proportion of the area that was statistically significant, this is highlighted as follows: *** 99.9% ** 99% * 95%.

Cell size (m)	**% of offences and area searched (theft from the person)**					**% of study area with Gi* values**		
	5% of offences	**10% of offences**	**25% of offences**	**50% of offences**	**80% of offences**	**95%**	**99%**	**99.9%**
50	**0.3%*** **	**0.8%*** **	**4.1%**	19.0%	59.7%	1.81%	1.20%	1.20%
100	0.5%***	1.1%***	4.9%	**18.6%**	50.2%	3.33%	2.09%	1.86%
150	0.5%***	1.6%***	5.9%	**18.6%**	**47.2%**	4.28%	2.80%	1.94%
200	0.7%***	2.0%***	6.8%	20.3%	50.0%	4.46%	2.77%	2.77%
250	0.8%***	2.1%***	8.5%	21.3%	49.2%	4.83%	3.17%	2.57%

Table 8.10 shows results for the area under the curve and CPI values for the different sub-sections of the Camden/Islington GI* theft from the person accuracy concentration curves. Small cell sizes generated the best results for up to 20% of the study area. Table 8.10 also shows that the best results were for the smallest areas searched. For example, the CPI value for Gi* hotspot maps, produced using 50 m cells, under the 0.5% of the study area sub-section of the accuracy concentration curve, was 0.831 compared to 0.439 for the 20% of the study area sub-section of the accuracy concentration curve. This suggests that the prediction performance of Gi* hotspot analysis was better for small areas searched, reflecting the identification of these areas as being spatially statistically significant. It also suggests that Gi* hotspot analysis generated using small cell sizes may produce better predictions of crime than those generated using larger cell sizes. Also of note were the higher CPI values for Gi* theft from the person hotspot analysis output compared to the CPI values for burglary dwelling Gi* hotspot maps. This finding suggests that retrospective data on theft from the person was better for predicting hotspots of this crime type compared to the use of retrospective data on burglary dwelling for predicting burglary dwelling hotspots.

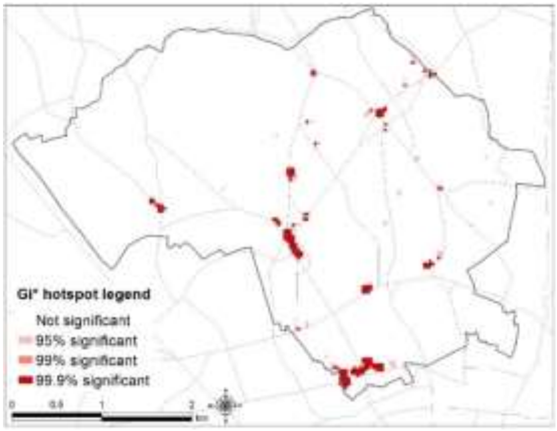

(a) Cell size: 50 m

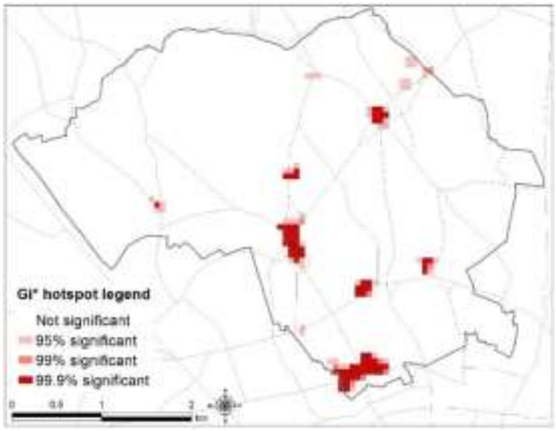

(b) Cell size: 100 m

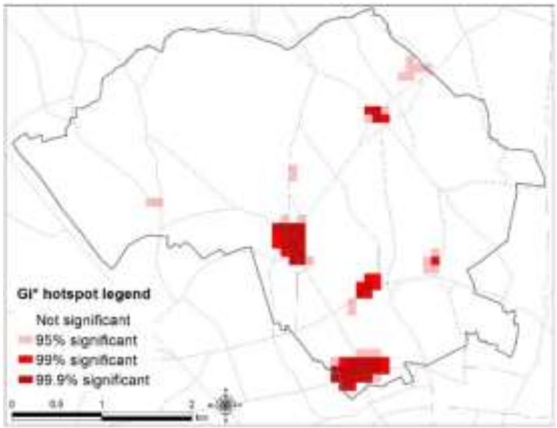

(c) Cell size: 150 m

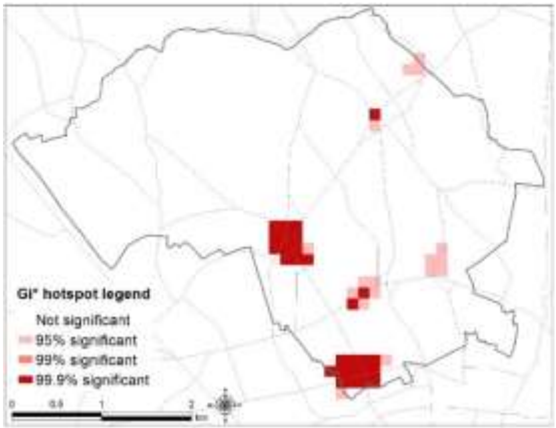

(d) Cell size: 200 m

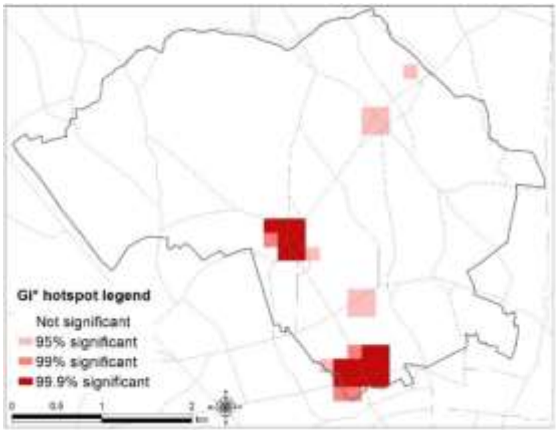

(e) Cell size: 250 m

Figure 8.8. Camden/Islington Gi* theft from the person hotspot maps produced using three months of input data and cell sizes of (a) 50 m, (b) 100 m, (c) 150 m, (d) 200 m, and (e) 250 m

Table 8.10. Camden/Islington Gi* theft from the person hotspot analysis of (a) the area under accuracy concentration curves, and (b) CPI values, for different cell sizes. Values in bold represent the largest area under the curve and the largest CPI values.

(a)

Cell size (m)	**0.5% x 5%**	**1% x 10%**	**5% x 25%**	**10% x 50%**	**20% x 80%**	**100% x 100%**
50	**0.000208**	**0.000690**	**0.008988**	**0.025317**	**0.070272**	0.71100
100	0.000134	0.000476	0.007657	0.023237	0.068515	0.734313
150	0.000098	0.000388	0.006642	0.021311	0.065648	**0.738449**
200	0.000112	0.000382	0.005973	0.01914	0.060804	0.730548
250	0.000098	0.000319	0.004967	0.014322	0.055415	0.721725
Max values	0.00025	0.001	0.0125	0.05	0.16	1

(b)

Cell size (m)	**0.5% x 5% CPI**	**1% x 10% CPI**	**5% x 25% CPI**	**10% x 50% CPI**	**20% x 80% CPI**	**100% x 100% CPI**
50	**0.831**	**0.689**	**0.719**	**0.506**	**0.439**	0.711
100	0.536	0.476	0.613	0.465	0.428	0.734
150	0.391	0.388	0.531	0.426	0.410	**0.738**
200	0.450	0.382	0.478	0.383	0.380	0.731
250	0.390	0.319	0.397	0.286	0.346	0.722

III. Newcastle burglary dwelling

Figure 8.9 shows the accuracy concentration curves for Newcastle Gi* burglary dwelling hotspot analysis outputs. These show a noticeable flattening of the curve gradients at about 8% of the study area coverage for the 50 m cell size, 25% for 100 m, 38% for 150 m, 42% for 200 m and 50% for 250 m. These inflection points relate to where Gi* values turned from positive to negative Z scores. Examination of the 0%-5% study area sub-section of the accuracy concentration curves show a noticeable difference in the gradients of each cell size curve from above 1% of the study area, with smaller cell sizes having curves with steeper gradients than larger cell sizes. However, below 1% of the study area, this pattern is more mixed, with no cell sizes consistently performing better than others (see Figures 8.9 c and d).

(a)

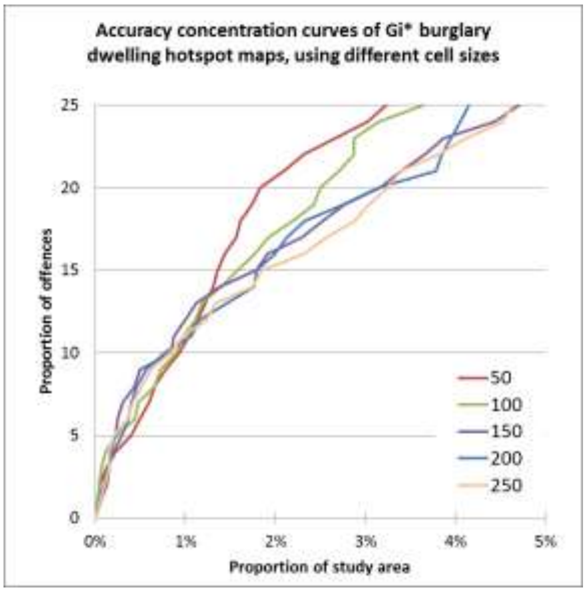

(b)

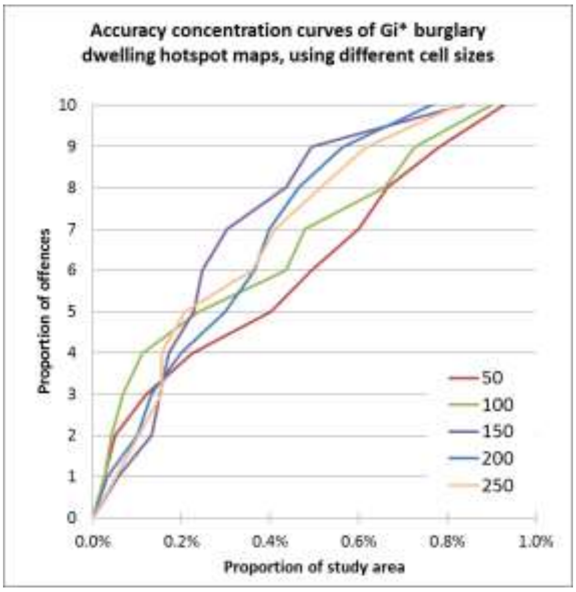

(c)

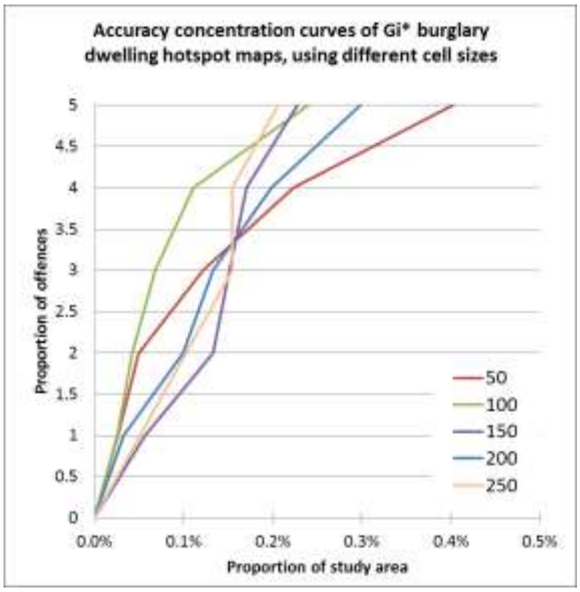

(d)

Figure 8.9. Accuracy concentration curves of Newcastle Gi* burglary dwelling hotspot analysis for different cell sizes. The results are shown for (a) the full study area, and sub-sections relating to (b) 5% of the study area, (c) 1% of the study area and (d) 0.5% of the study area.

Table 8.11 lists the proportion of the area searched to find 5%, 10%, 25%, 50% and 80% of burglary dwelling offences in the one-month output period. These results show there to be no consistency in the prediction performance between different cell sizes. Table 8.11 also shows that the areas that needed to be searched to identify 5% and 10% of all burglary dwelling offences contained Gi* values that were statistically significant to at least 99.9%.

Table 8.11. The proportion of the Newcastle study area searched across maps generated using Gi* (three months of input data) relative to 5%, 10%, 25%, 50%, and 80% of burglary dwelling offences, for different cell sizes. The proportion of the study area with Gi* values relative to 95%, 99% and 99.9% is also listed. Values in bold represent the smallest areas searched. Where the area searched was not greater than the proportion of the area that was statistically significant, this is highlighted as follows: *** 99.9% ** 99% * 95%.

Cell size (m)	**% of offences and area searched (burglary dwelling)**					**% of study area with Gi* values**		
	5% of offences	**10% of offences**	**25% of offences**	**50% of offences**	**80% of offences**	**95%**	**99%**	**99.9%**
50	0.40***	0.93***	**3.23**	11.9	41.0%	1.0%	1.0%	1.0%
100	0.24***	0.90 ***	3.65	**10.95**	36.9%	2.3%	1.5%	1.5%
150	0.23***	0.83***	4.7	12.39	**30.4%**	2.4%	2.4%	1.6%
200	0.30***	**0.76*** **	4.15	13.84	30.6%	3.6%	2.8%	2.1%
250	**0.21*** **	0.82***	4.63	13.13	31.9%	4.0%	3.2%	2.7%

Table 8.12 lists the Newcastle Gi* Bonferroni corrected Z scores for 95%, 99% and 99.9% significance levels for the different cell sizes, and highlights the higher Z score significance threshold values this correction determines in comparison to uncorrected Z scores. For example, rather than using 1.96 to determine 95% significance for a cell size of 50 m, the Bonferroni corrected Z score for this significance level was 4.875.

Table 8.12. Bonferroni corrected Z score spatial statistical significance values for Newcastle burglary dwelling Gi* maps of different cell sizes

Cell sizes (m)	**n of cells covering study area**	**Bonferroni corrected Z scores**		
		95%	**99%**	**99.9%**
50	46019	4.875	5.184	5.598
100	11685	4.597	4.922	5.355
150	5272	4.429	4.764	5.209
200	3012	4.306	4.650	5.104
250	1942	4.208	4.559	5.021

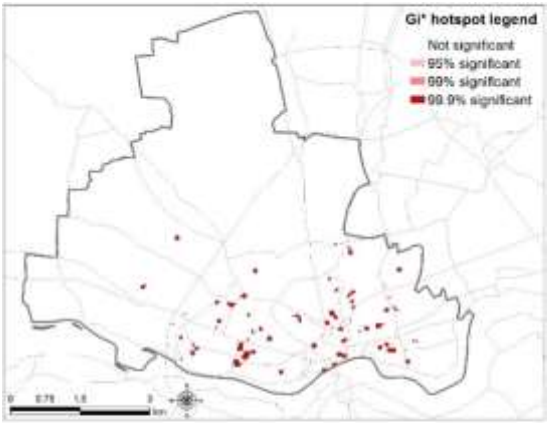

(a) Cell size: 50 m

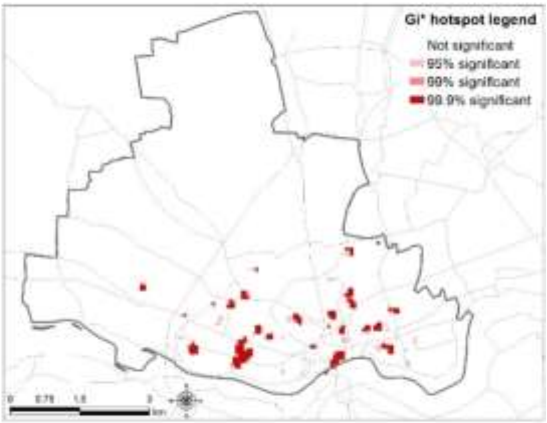

(b) Cell size: 100 m

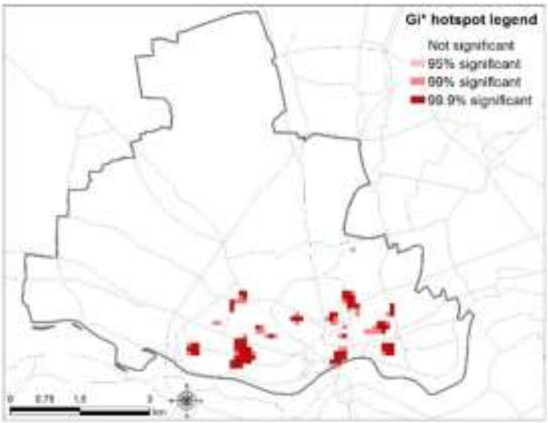

(c) Cell size: 150 m

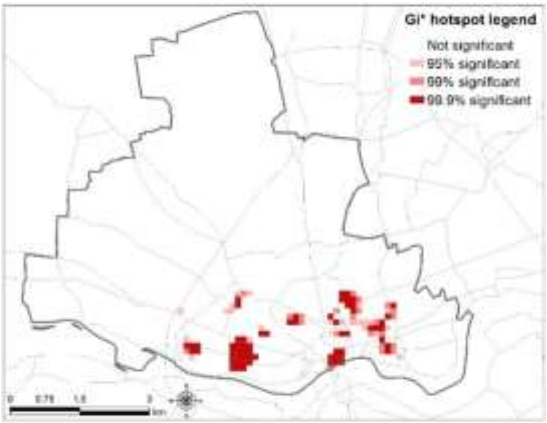

(d) Cell size: 200 m

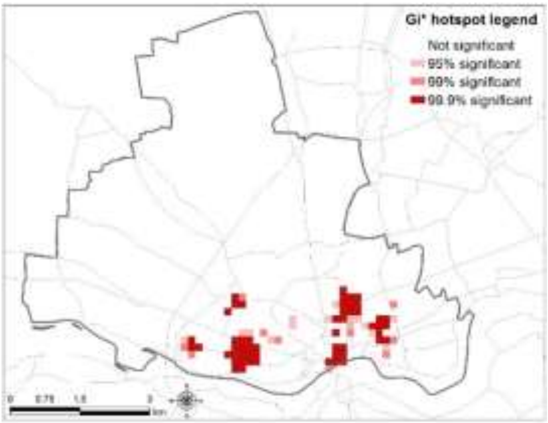

(e) Cell size: 250 m

Figure 8.10. Newcastle Gi* burglary dwelling hotspot maps produced using three months of input data and cell sizes of (a) 50 m, (b) 100 m, (c) 150 m, (d) 200 m, and (e) 250 m

Figure 8.10 shows the Newcastle Gi* burglary dwelling hotspot maps (using three months of input data) for different cell sizes. Similar to the previous Camden/Islington examples in Figures 8.6 and 8.8, these maps show how the number of hotspots identified reduced as cell size increased. The area under the accuracy concentration curves and CPI values shown in Table 8.13 also suggest that Gi* hotspot maps generated using the smaller cell sizes (50 m to 150 m) performed better in predicting spatial patterns of crime. For example, the CPI value for the 0%-0.5% sub-section of the accuracy concentration curve for a Gi* hotspot analysis using a 100 m cell size was 0.853 compared to 0.773 for a Gi* hotspot analysis using a 250 m cell size.

Table 8.13. Newcastle Gi* burglary dwelling hotspot analysis of (a) the area under accuracy concentration curves, and (b) CPI values, for different cell sizes. Values in bold represent the largest area under the curve and the largest CPI values.

(a)

Cell size (m)	0.5% x 5%	1% x 10%	5% x 25%	10% x 50%	20% x 80%	100% x 100%
50	0.000188	0.000617	**0.009474**	**0.03222**	0.083873	0.79853
100	**0.000213**	0.000676	0.009082	0.029513	**0.086828**	0.816491
150	0.000187	**0.000736**	0.008539	0.026002	0.082336	0.822238
200	0.000189	0.000706	0.008438	0.026248	0.079797	0.822238
250	0.000193	0.000701	0.008221	0.025392	0.078442	**0.822802**
Max values	0.00025	0.001	0.0125	0.05	0.16	1

(b)

Cell size (m)	0.5% x 5% CPI	1% x 10% CPI	5% x 25% CPI	10% x 50% CPI	20% x 80% CPI	100% x 100% CPI
50	0.751	0.617	**0.758**	**0.644**	0.524	0.799
100	**0.853**	0.676	0.727	0.590	**0.543**	0.816
150	0.750	**0.736**	0.683	0.520	0.515	0.822
200	0.754	0.706	0.675	0.525	0.499	0.822
250	0.773	0.701	0.658	0.508	0.490	**0.823**

IV. Newcastle theft from the person

Figure 8.11 shows the accuracy concentration curves for Newcastle Gi* theft from the person hotspot analysis outputs. These curves show a noticeable flattening of the gradients from about 5% of the study area coverage for 50 m cell size to about 22% for 250 m cell size. This flattening relates to the point that Gi* values turned from positive to negative Z scores. Examination of the 0%-5%, 0%-1% and 0%-0.5% study area sub-sections of the accuracy concentration curves show the very high curve gradients for each cell size, indicative of a high degree of prediction performance in these Gi* hotspot analysis outputs. Although there is little difference between the curves, the curve for the cell size of 50 m consistently had the steepest gradient across each of the sub-sections. A similar pattern is found in the results for the proportion of the study area searched (Table 8.14), the areas under the accuracy concentration curves and the CPI values (Table 8.15) – with the best results being for the Gi* hotspot analysis that used a 50 m cell size.

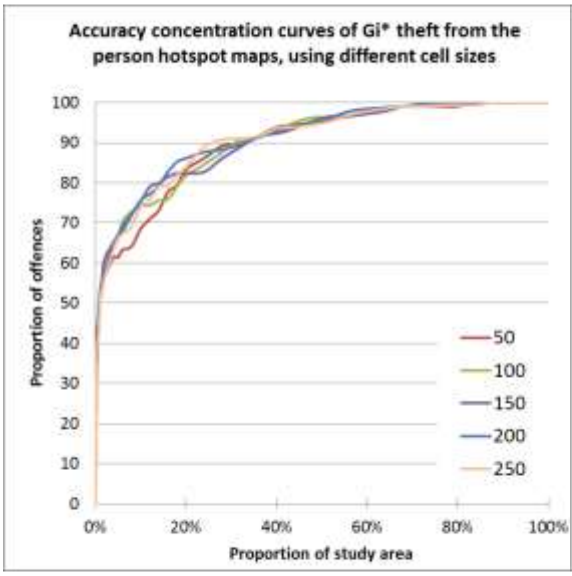

(a)

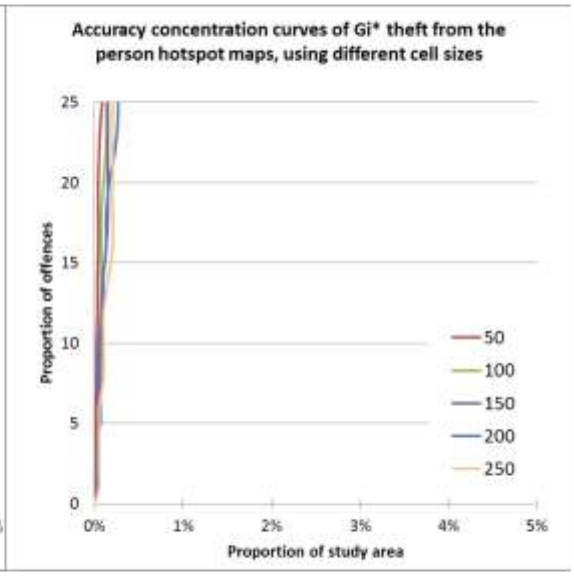

(b)

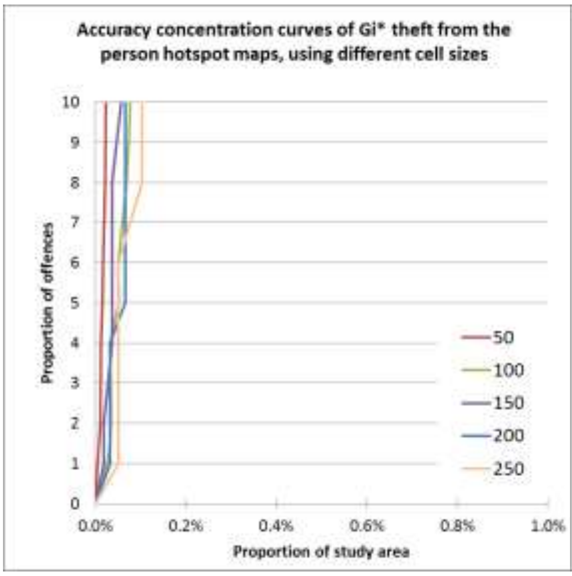

(c)

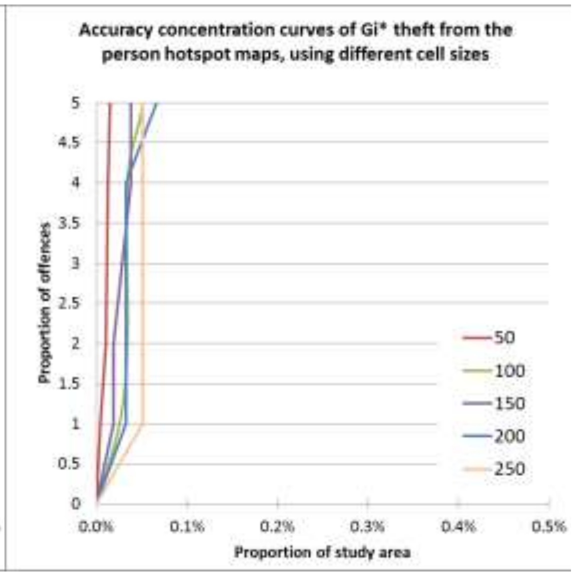

(d)

Figure 8.11. Accuracy concentration curves of Newcastle Gi* theft from the person hotspot analysis for different cell sizes. The results are shown for (a) the full study area, and sub-sections relating to (b) 5% of the study area, (c) 1% of the study area and (d) 0.5% of the study area.

Table 8.14. The proportion of the Newcastle study area searched across maps generated using Gi* (three months of input data) relative to 5%, 10%, 25%, 50%, and 80% of theft from the person offences, for different cell sizes. The proportion of study area with Gi* values relative to 95%, 99% and 99.9% is also listed. Values in bold represent the smallest areas searched. Where the area searched was not greater than the proportion of the area that was statistically significant, this is highlighted as follows: *** 99.9% ** 99% * 95%.

Cell size (m)	**% of offences and area searched (theft from the person)**					**% of study area with Gi* values**		
	5% of offences	**10% of offences**	**25% of offences**	**50% of offences**	**80% of offences**	**95%**	**99%**	**99.9%**
50	**0.02%*****	**0.02%*****	**0.06%*****	0.81%	18.2%	0.36%	0.36%	0.36%
100	0.05%***	0.08%***	0.14%***	**0.76%**	18.3%	0.73%	0.65%	0.64%
150	0.04%***	0.06%***	0.15%***	0.85%	**14.0%**	0.80%	0.72%	0.68%
200	0.07%***	0.07%***	0.24%***	0.80%***	14.2%	0.93%	0.90%	0.90%
250	0.05%***	0.10%***	0.21%***	0.98%***	16.7%	0.98%	0.98%	0.98%

Table 8.15. Newcastle Gi* theft from the person hotspot analysis of (a) the area under accuracy concentration curves, and (b) CPI values, for different cell sizes. Values in bold represent the largest area under the curve and the largest CPI values.

(a)

Cell size (m)	**0.5% x 5%**	**1% x 10%**	**5% x 25%**	**10% x 50%**	**20% x 80%**	**100% x 100%**
50	**0.000245**	**0.000985**	**0.01242**	**0.0491**	0.135	0.901
100	0.000235	0.000953	0.01231	0.0489	**0.141**	0.907
150	0.000238	0.000967	0.01228	0.0487	0.112	0.907
200	0.000233	0.000950	0.01222	0.0486	0.110	**0.913**
250	0.000227	0.000936	0.01217	0.0485	0.124	0.909
Max values	0.00025	0.001	0.0125	0.05	0.16	1

(b)

Cell size (m)	**0.5% x 5% CPI**	**1% x 10% CPI**	**5% x 25% CPI**	**10% x 50% CPI**	**20% x 80% CPI**	**100% x 100% CPI**
50	**0.981**	**0.985**	**0.993**	**0.981**	0.842	0.901
100	0.938	0.953	0.985	0.978	**0.879**	0.907
150	0.951	0.967	0.982	0.974	0.699	0.907
200	0.934	0.950	0.977	0.972	0.691	**0.913**
250	0.907	0.936	0.973	0.969	0.776	0.909

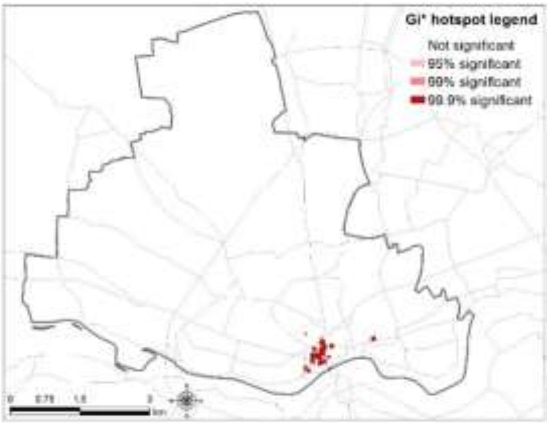

(a) Cell size: 50 m

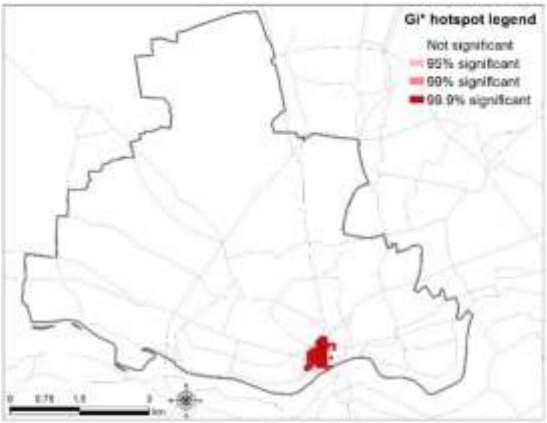

(b) Cell size: 100 m

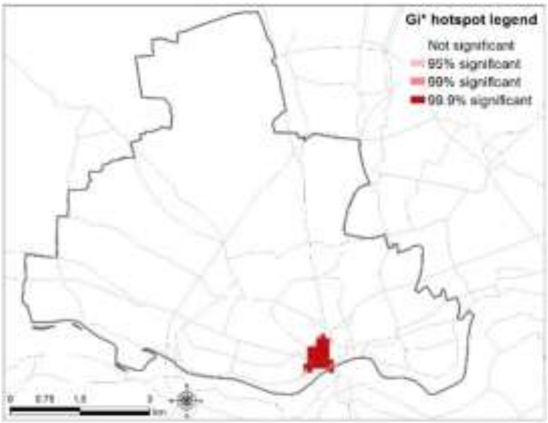

(c) Cell size: 150 m

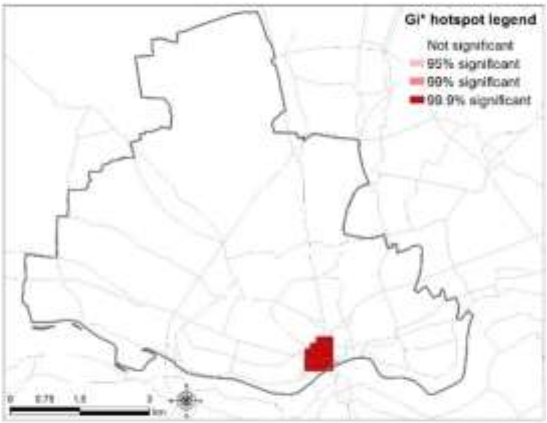

(d) Cell size: 200 m

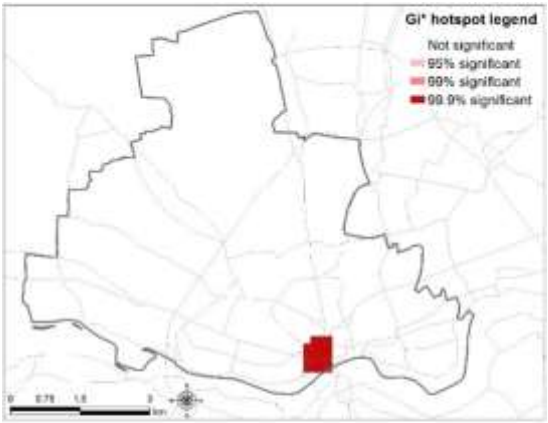

(e) Cell size: 250 m

Figure 8.12. Newcastle Gi* theft from the person hotspot maps produced using six months of input data and cell sizes of (a) 50 m, (b) 100 m, (c) 150 m, (d) 200 m, and (e) 250 m

Figure 8.12 shows the Newcastle Gi* hotspot maps generated for the different cell sizes. All five maps show that the city centre of Newcastle featured as the main hotspot, with this identified area changing little across maps of different cell sizes. A similar pattern is shown in the prediction performance of each Gi* hotspot map (i.e., the CPI values shown in Table 8.15), with little difference between the outputs generated for different cell sizes. The Gi* theft from the person hotspot map generated using a cell size of 50 m (Figure 8.12a) does, though, illustrate the more focused area that a small cell size Gi* map produces, particularly when the clustering of crime is geographically compact (as in the case of theft from the person in Newcastle).

V. Newcastle assaults with injury

Figure 8.13 shows the accuracy concentration curves for Newcastle Gi* assault with injury hotspot analysis outputs. These show a noticeable flattening of the curve gradients from about 10% of the study area coverage for 50 m cell size to about 50% for 250 m cell size. This flattening relates to the point that Gi* values turned from positive to negative Z scores. Examination of the 0%-5%, 0%-1% and 0%-0.5% study area sub-sections of the accuracy concentration curves show the very steep curve gradients for each cell size, indicative of a high degree of prediction performance in these Gi* hotspot analysis outputs. Although there is little difference between the curves, the curve for the cell size of 50 m consistently had the steepest gradient across each of the sub-sections. This is also reflected in the results for the proportion of the study area searched (Table 8.16), the areas under the accuracy concentration curves and the CPI values (Table 8.17) – with the best results being for the Gi* hotspot analysis that used a 50 m cell size.

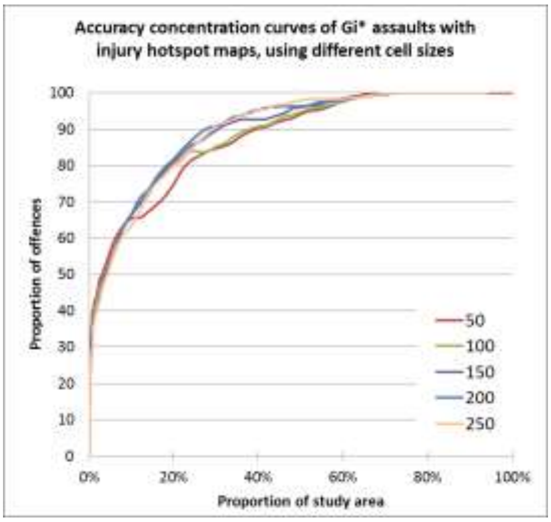

(a)

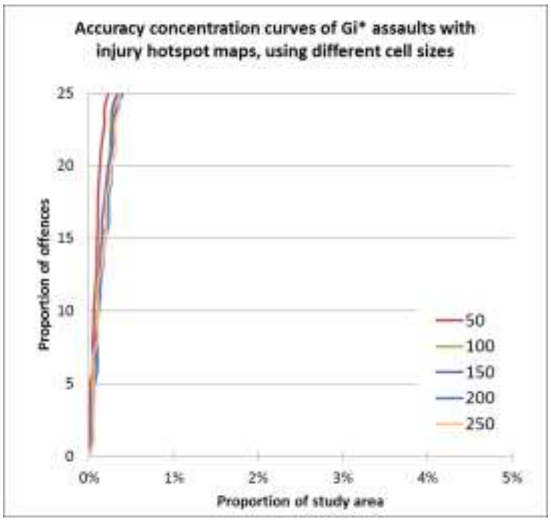

(b)

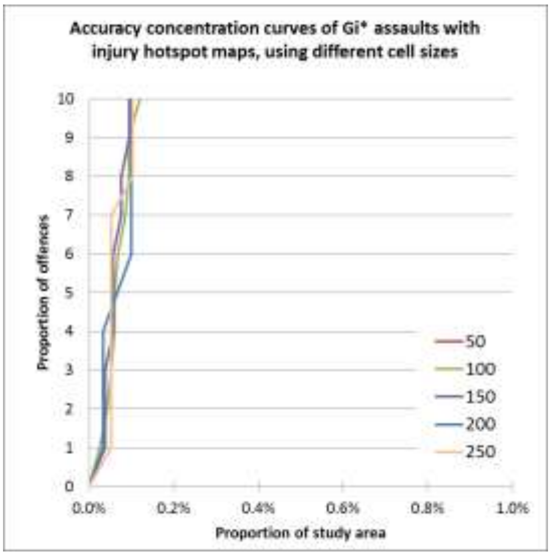

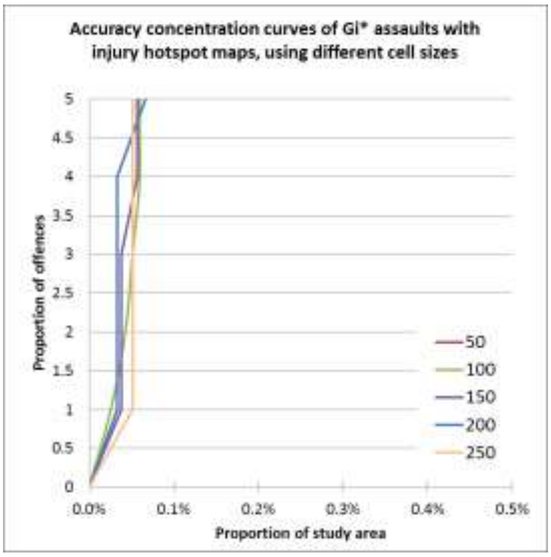

(c) (d)

Figure 8.13. Accuracy concentration curves of Newcastle Gi* assault with injury hotspot analysis for different cell sizes. The results are shown for (a) the full study area, and sub-sections relating to (b) 5% of the study area, (c) 1% of the study area and (d) 0.5% of the study area.

Table 8.16. The proportion of the Newcastle study area searched across maps generated using Gi* (three months of input data) relative to 5%, 10%, 25%, 50%, and 80% of assaults with injury offences, for different cell sizes. The proportion of study area with Gi* values relative to 95%, 99% and 99.9% is also listed. Values in bold represent the smallest areas searched. Where the area searched was not greater than the proportion of the area that was statistically significant, this is highlighted as follows: *** 99.9% ** 99% * 95%.

Cell size (m)	**% of offences and area searched (assaults with injury)**					**% of study area with Gi* values**		
	5% of offences	**10% of offences**	**25% of offences**	**50% of offences**	**80% of offences**	**95%**	**99%**	**99.9%**
50	**0.02%*** **	**0.06%*** **	**0.23%*** **	**31.0%**	22.8%	0.60%	0.60%	0.44%
100	0.06%***	0.12%***	0.33%***	36.1%	18.7%	0.80%	0.77%	0.77%
150	0.06%***	0.09%***	0.34%***	37.9%	18.5%	1.02%	0.95%	0.91%
200	0.07%***	0.10%***	0.40%***	39.8%	**18.4%**	1.20%	1.20%	1.13%
250	0.05%***	0.10%***	0.36%***	43.3%	20.2%	1.44%	1.44%	1.34%

Table 8.17. Newcastle Gi* assault with injury hotspot analysis of (a) the area under accuracy concentration curves, and (b) CPI values, for different cell sizes. Values in bold represent the largest area under the curve and the largest CPI values.

(a)

Cell size (m)	0.5% x 5%	1% x 10%	5% x 25%	10% x 50%	20% x 80%	100% x 100%
50	**0.000246**	**0.000974**	**0.01229**	**0.0472**	0.1226	0.878
100	0.000229	0.000936	0.01213	0.0463	**0.1260**	0.885
150	0.000230	0.000942	0.01215	0.0459	0.1257	0.894
200	0.000233	0.000935	0.01209	0.0457	0.1260	**0.899**
250	0.00023	0.000938	0.01210	0.0454	0.1229	0.895
Max values	0.000250	0.001	0.0125	0.05	0.16	1

(b)

Cell size (m)	0.5% x 5% CPI	1% x 10% CPI	5% x 25% CPI	10% x 50% CPI	20% x 80% CPI	100% x 100% CPI
50	**0.983**	**0.973**	**0.983**	**0.944**	0.766	0.878
100	0.916	0.936	0.970	0.927	**0.788**	0.885
150	0.920	0.942	0.972	0.918	0.785	0.894
200	0.934	0.935	0.967	0.913	0.787	**0.899**
250	0.934	0.938	0.968	0.908	0.768	0.895

Figure 8.14 shows the Newcastle Gi* assaults with injury hotspot maps generated for the different cell sizes. These show from Figure 8.14a through to 8.14e the city centre of Newcastle to feature as the main hotspot, with this identified area changing little across these five maps of different cell sizes. However, the Gi* hotspot map generated using a cell size of 50 m not only captures the statistically significant concentration of crime in the Newcastle city centre area in a more focused manner, but also shows other hotspots in other parts of the district. These features of this small cell size Gi* hotspot map are perhaps reflected in its consistently high CPI values for sub-sections of the concentration accuracy curves representing up to 50% of offences (see Table 8.17).

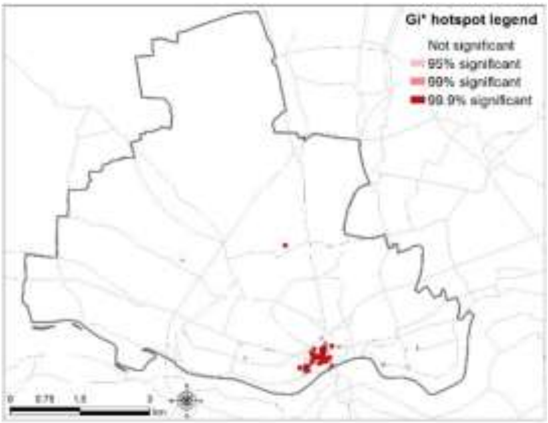

(a) Cell size: 50 m

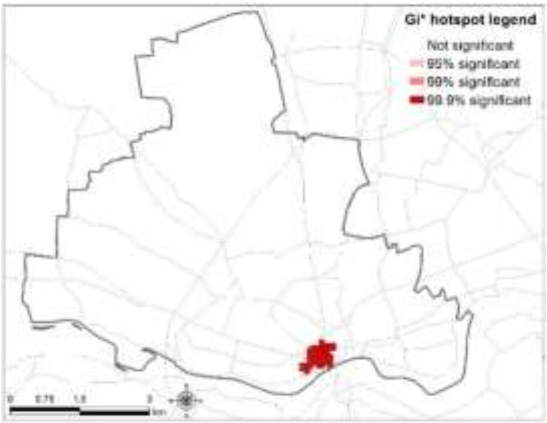

(b) Cell size: 100 m

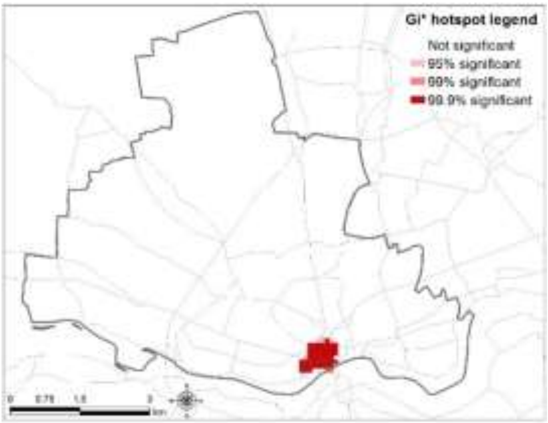

(c) Cell size: 150 m

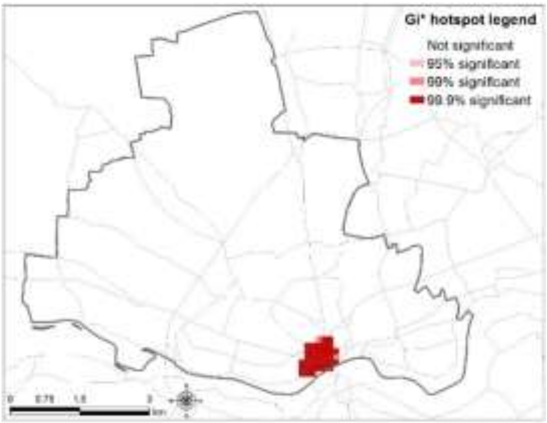

(d) Cell size: 200 m

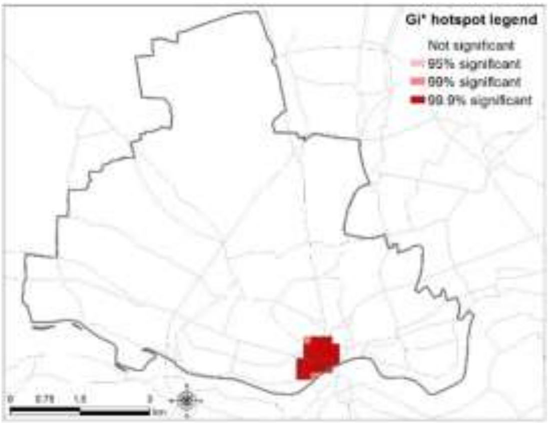

(e) Cell size: 250 m

Figure 8.14. Newcastle Gi* assaults with injury hotspot maps produced using six months of input data and cell sizes of (a) 50 m, (b) 100 m, (c) 150 m, (d) 200 m, and (e) 250 m

8.5.5. A metric assessment that compares the prediction performance of Gi* hotspot analysis to KDE hotspot analysis

The analysis so far in study 4 of the research has shown that while Gi* hotspot analysis helps remove much of the ambiguity in defining areas that are hotspots (compared to KDE and the other common thematic hotspot analysis techniques), it may not necessarily be better at predicting crime than KDE hotspot analysis output. So far, these comparisons have only been in relation to the PAI. The more complete analysis of the prediction performance of Gi* hotspot analysis using accuracy concentration curves and calculating CPI values has shown that Gi* outputs can generate very good spatial predictions of crime. This section compares these more complete and detailed prediction measures for Gi* with their equivalent KDE results.

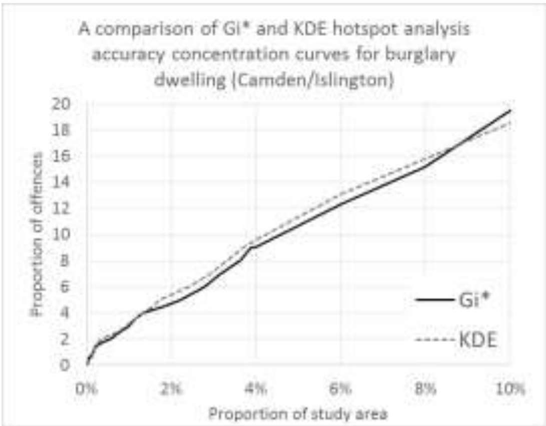

(a)

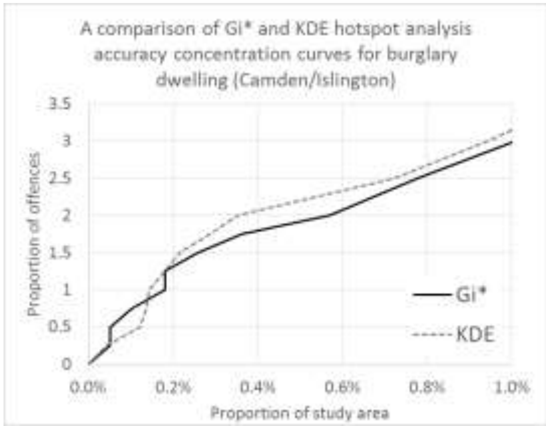

(b)

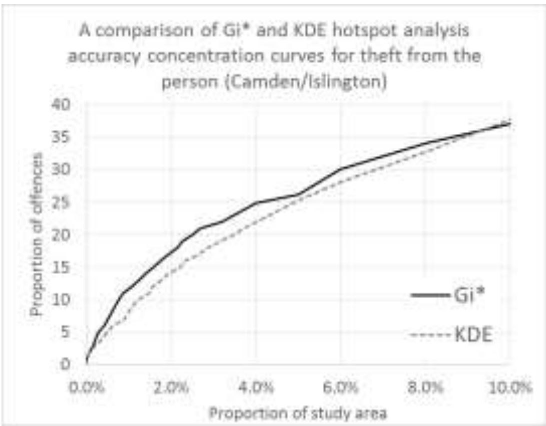

(c)

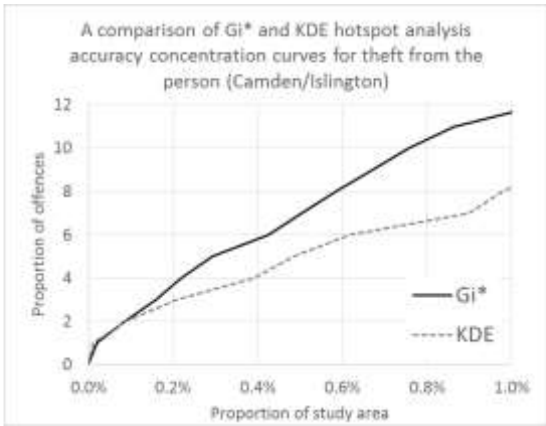

(d)

Figure 8.15. A comparison of Gi* and KDE hotspot analysis accuracy concentration curves (for sub-sections of 10% of the study area and 1% of study area) for Camden/Islington (a and b) burglary dwelling and (c and d) theft from the person

Figures 8.15 and 8.16 show the accuracy concentration curves for Gi* and KDE hotspot maps of burglary dwelling and theft from the person in Camden/Islington, and of burglary dwelling, theft from the person and assaults with injury in Newcastle. The charts depict the proportion of crime predicted by each technique within 10% and 1% of the study area (i.e., the top 10% and top 1% of Gi* and KDE values). The analysis focused on sub-sections below 10% of the study area because this is the size of the area where the highest prediction levels are of most value to support the targeted nature of policing and public safety. The KDE hotspot analysis for each crime type and study area used the bandwidths that produced the highest CPI values. Similarly, the Gi* hotspot analysis for each crime and study area used the cell size that produced the highest CPI values.

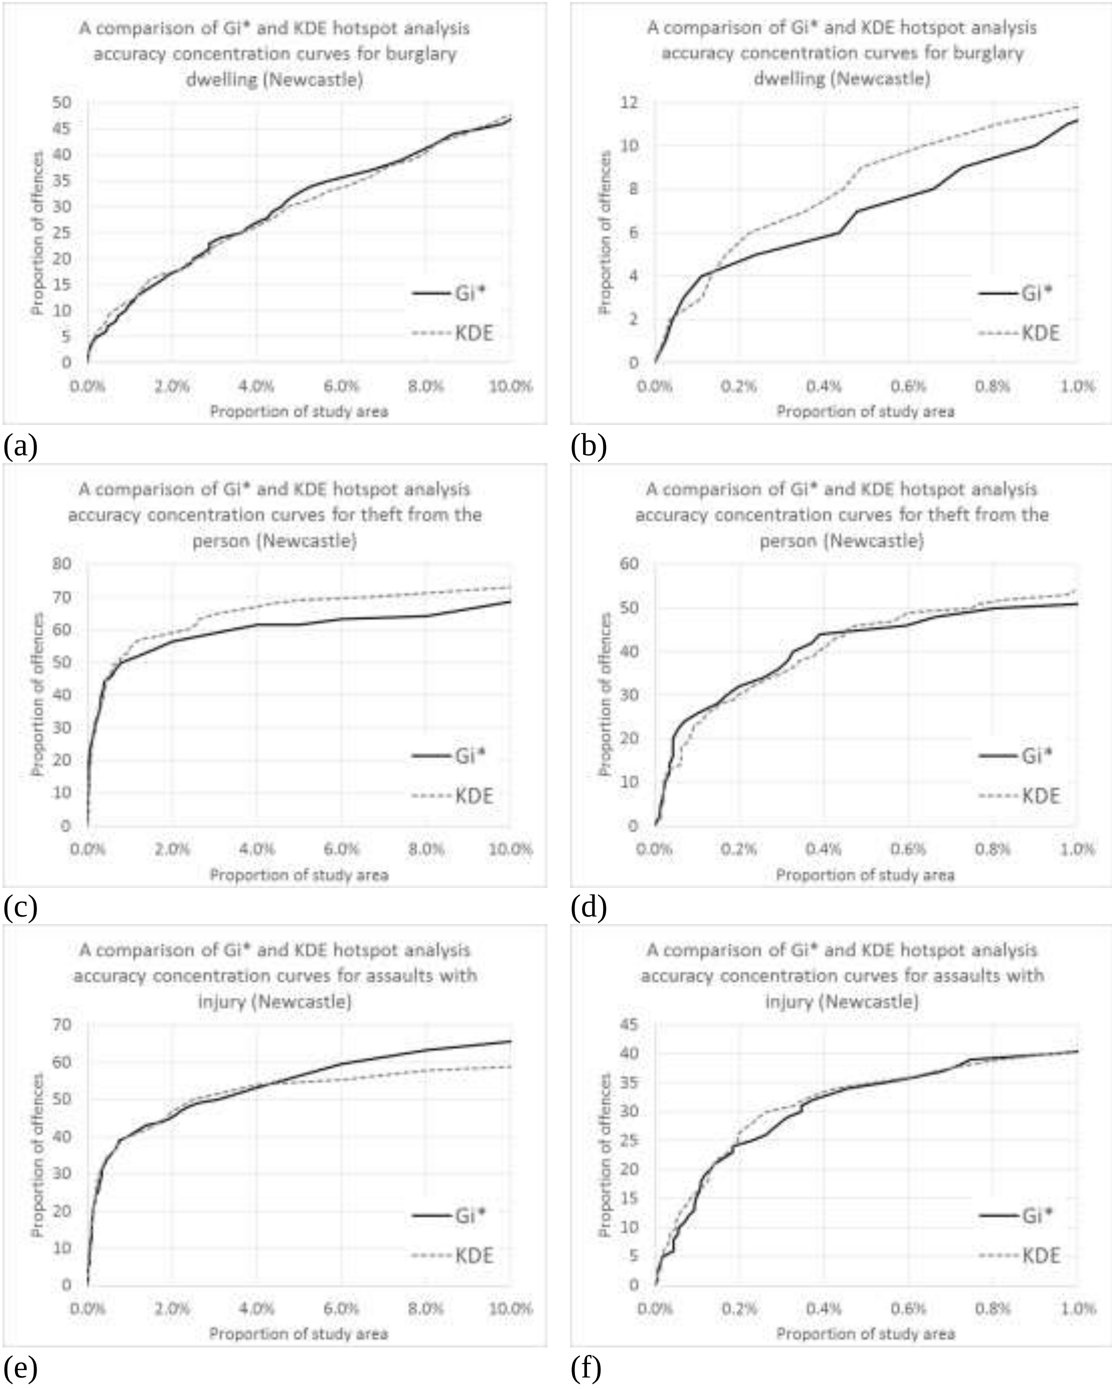

Figure 8.16. A comparison of Gi* and KDE hotspot analysis accuracy concentration curves (for sub-sections of 10% study area and 1% of study area) for Newcastle (a and b) burglary dwelling, (c and d) theft from the person, and (e and f) assaults with injury

Using the charts for 0%-10% of each study area (Figure 8.15a and c, and 8.16a, c and e), the results show the prediction performance of Gi* and KDE maps to be very similar, although in one case (Camden/Islington theft from the person – Figure 8.15c) Gi* results were marginally better than those for KDE, but for another (Newcastle theft from the person – Figure 8.16c) KDE results were marginally better than those for Gi*. Of note was that Gi* Bonferroni corrected 95% statistical significance levels for Newcastle theft from the person extended to less than 1% of the study area – it was above this point that the Gi* and KDE accuracy concentration curves separated (as shown in Figure 8.16c). This finding suggests that beyond this 1% point, the clustering of crime was not spatially significant, and hence any predictions of crime would be outside the areas that were determined as hot. Closer analysis using the 0%-1% of the study area charts again showed there to be little difference at the very small area level between Gi* and KDE hotspot analysis results. The main difference was between the better results for Gi* hotspot analysis of Camden/Islington theft from the person (Figure 8.15d) and the marginally better results for KDE hotspot analysis of Newcastle burglary dwelling than those using Gi* (Figure 8.16b).

Analysis using CPI values in the previous sections presented these CPI values using sub-sections of the accuracy concentration charts. This approach is useful when there is no means for determining when the spatial concentration of crime can be defined as *hot*. In the case of Gi* hotspot analysis, *hot* can be determined using statistical significance levels. Our interest is, therefore, towards the CPI values for study area extents determined as *hot* from Gi* hotspot analysis.

Recall that a perfect prediction is where the CPI is equal to one. Table 8.18 lists the CPI values for Newcastle burglary dwelling, theft from the person and assault with injury Gi* hotspot analysis for the areas determined as hot (using a 95% Bonferroni corrected statistical significance level), and for outputs generated using 50 m, 100 m and 150 m cell sizes (i.e., the areas shown in the Gi* maps in Figures 8.10, 8.12 and 8.14 a, b and c). The results show that for all crime types the CPI values were very high, ranging from 0.953–0.960 for burglary dwelling, 0.995–0.997 for theft from the person, and 0.992–0.996 for assault with injury. These results are also illustrated in Figure 8.17, which shows the near vertical accuracy concentration curves for each of these crime types, in particular theft from the person and assault with injury. The graphs depict the proportion of the study area by the proportion of offences relative to the area of the study area that was determined as *hot*, and using equal axes values to accurately represent the gradient of the accuracy concentration curves.

Table 8.18. CPI values for Newcastle Gi* hotspot analysis for different cell sizes, and within the extent of areas that were statistically determined as *hot*

Crime type	**Cell size (m)**	**CPI values for significance levels**		
		95%	**99%**	**99.9%**
Burglary dwelling	50	0.960	0.960	0.960
	100	0.953	0.959	0.959
	150	0.955	0.955	0.963
Theft from the person	50	0.997	0.997	0.997
	100	0.996	0.996	0.996
	150	0.995	0.995	0.995
Assault with injury	50	0.995	0.995	0.996
	100	0.993	0.993	0.993
	150	0.992	0.992	0.993

The CPI values for Gi* (shown in Table 8.18) were higher than any KDE CPI values generated in research study 3. While this research has shown that, following a complete metric assessment that has compared Gi* results to KDE results, that the prediction performance of the two techniques does not differ substantially, the identification of hotspots using KDE is subjective and ambiguous. This means that the prediction performance of KDE hotspot analysis can be varied because different researchers may choose different sized areas in their determination of what is *hot*. For example, a researcher may choose the quantile range thematic classification method rather than the equal ranges method, or use the top 10% of KDE values rather than the top 5% of KDE values to define his/her hotspot. Gi* hotspot analysis helps remove much of this ambiguity in determining areas that are *hot*, with the result being the identification of areas that are selected based on statistical grounds and that perform extremely well in identifying where crime is likely to occur in the future.

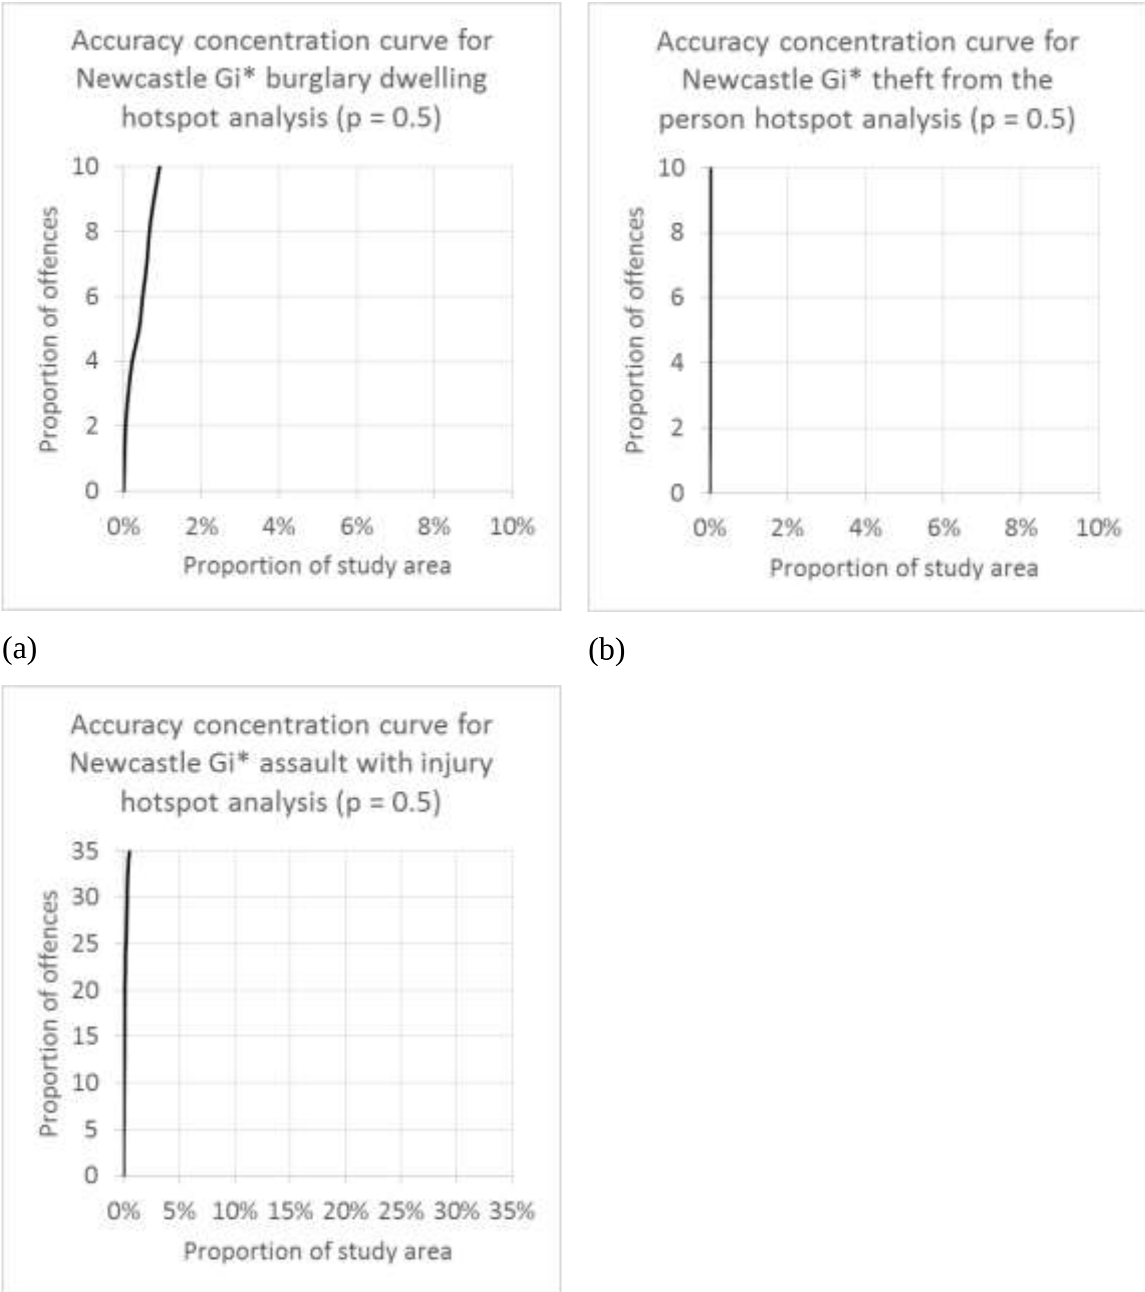

(a) (b)

(c)

Figure 8.17. Examples of accuracy concentration curves for Newcastle Gi* hotspot analysis (generated using a cell size of 50 m in each case), with axes controlled to be equal and extend to the coverage area that was statistically determined to be *hot*: (a) burglary dwelling, (b) theft from the person and (c) assault with injury

8.6. Interpretation and conclusions from research study 4

This research study's detailed assessment of the Gi* statistic for predicting spatial patterns of crime has resulted in six main findings:

- Although KDE produces hotspot maps that perform well in predicting spatial patterns of crime, the selection of areas that are *hot* is subjective. The Gi* mapping output identifies similar areas to KDE, and helps remove the ambiguity in defining areas that are *hot* by using the statistical principles of significance testing. However, due to the analyst being required to choose their preference in the multiple testing correction procedure (e.g., Bonferroni rather than false discovery rate) and the statistical significance threshold to apply (e.g., 99% rather than 95%) there remains some subjectivity in determining the areas that are hot.

- Gi* mapping output can be used to identify only those areas that are statistically significant, rather than produce a geographical representation showing all areas where crime has been committed. This can help in visually targeting attention to the areas that matter most – the hotspots
- A weakness of KDE is that it can smooth out hotspots in areas where the crime concentration is very spatially compact, or can exaggerate these when the numerical density value that is used to define *hot* (the upper thematic range) is relatively low (e.g., when a quantile range method is used). Gi* mapping output offers a better means of identifying small compact concentrations of crime but not to the extent of visually exaggerating the area of influence these compact concentrations have at the local level
- The prediction performance of Gi* mapping output initially appeared to be equivalent in performance to KDE hotspot mapping output. This was confirmed through the use of the full range of prediction performance measures – the PAI, accuracy concentration curves, the area under the curve and the CPI
- The prediction performance of Gi* mapping output improved as the statistical significance level was set higher. That is, Gi* mapping output that used a 99.9% statistical significance level to determine the areas that were *hot* were more accurate in predicting crime than when a 99% or 95% statistical significance level was used
- The initial comparison between Gi* and KDE mapping output was based on examining the prediction performance of the two techniques for different sub-sections of the accuracy concentration curves. This allowed for like-for-like comparisons between the prediction performance of Gi* and KDE outputs for a range of areal coverages of the study area. This approach is useful when there is no means for determining when the spatial concentration of crime can be defined as *hot*, but in the case of Gi* hotspot analysis, *hot* is determined using statistical significance levels. Analysis then revealed that many of the areal coverages that were used to compare Gi* to KDE included a proportion of the area that the Gi* statistic had identified to not be statistically significant. When only the areas that had been determined by the

Gi* statistic to be statistically significant were measured for their prediction performance, this analysis produced CPI values above 0.95 for burglary and above 0.99 for theft from the person and for assaults with injury. If we recall that a perfect prediction is 1, these results suggest that Gi* mapping output performs extremely well in predicting spatial patterns of crime, and offers an improvement on the equivalent CPI values that were generated for KDE. That is, the Gi* statistic appears to offer more accurate spatial predictions of crime than KDE.

With reference to the central hypothesis for research study 4 (hypothesis 4) – *spatial significance mapping methods provide an improved means of predicting where crime is likely to occur, and remove the ambiguity of defining areas that are hot* – the comprehensive range of experiments have resulted in providing evidence that supports both parts of this statement. Additionally, the ability to remove much of the subjectivity in identifying areas that are hotspots by using the Gi* statistic also makes it easier to then choose these defined hotspot areas and subject them to further analysis. This analysis can include examining the geographical characteristics of these hotspots to help determine why crime concentrates to high levels at these locations. When using KDE, because of the lack of clarity in defining the area that is *hot*, where the size of the hotspots would vary according to the thematic classification choice of the analyst, this in turn makes it difficult to clearly define the hotspot areas that should be subjected to further analysis.

The results have also shown that while cell size, and hence lag distance, typically only has a marginal influence on the prediction performance of Gi* hotspot mapping output, smaller cell sizes are preferable for their slightly higher levels of prediction performance. Also noticeable in the results, particularly when only the statistically defined hotspot areas were analysed for their prediction performance, was the extremely high CPI values. Not only do these results suggest that the Gi* statistic can be extremely effective in determining in a much focused (conservative) manner where police and crime prevention resources could be targeted and have significant impact, but that also there appears to be great value in using retrospective crime data for predicting spatial patterns of crime.

The results from this research study set an impressive benchmark for spatial crime prediction performance using the Gi* statistic, against which other techniques for predicting spatial patterns of crime can be compared. In research study 6, the prediction performance of Gi* hotspot analysis is compared to one of these other predictive mapping techniques - prospective mapping. A first task, however, is to analyse the stability of hotspots by examining the extent to which patterns of crime from the past can predict patterns of crime in the future. This is the focus of research study 5.

9. Research study 5: Examining the temporal stability of hotspots

9.1. Introduction

Research study 1 identified that an examination of retrospective crime data is required in order to establish if hotspots are present in crime data. Research studies 2 – 4 have then examined a number of hotspot analysis techniques for spatially representing hotspots, with attention also towards examining the potential in mapping where crime has concentrated previously for predicting where crime is likely to occur in the future. An examination of the common hotspot analysis techniques revealed KDE to perform best in predicting spatial patterns of crime. The Gi* statistic was then found to improve upon KDE by being less ambiguous in determining areas that are *hot* by using statistical significance principles and by performing better in predicting where crime is likely to occur. The research study that is the focus of this chapter aims to build upon these previous four research studies by examining in further statistical detail the temporal stability of hotspots.

Recent research suggests that hotspots tend to move around (Johnson and Bowers, 2004b; Johnson et al., 2008a). These findings into the fluid spatial patterns of crime reflect the foraging behaviour of offenders who operate in spates before moving on to other areas (Johnson et al., 2009). This mobility of offenders means that if areas are identified as hotspots based on a recent short-lived period of high crime, any targeting of policing or crime prevention activity to these areas may arrive after the crime issue has since moved on. The theoretical principles of offender foraging behaviour have been shown to be evident in patterns of repeat and near repeat victimisation (Fielding and Jones, 2012; Johnson et al., 2009) with these findings informing the prospective mapping approach to predicting spatial patterns of crime (Bowers et al., 2004). This research that has resulted in the development of prospective mapping has though exclusively focused on burglary dwelling[16] rather than examining whether these spatially fluid patterns of crime are evident in other crime types.

Results so far in this PhD research suggest, however, that the analysis of retrospective crime data provides a powerful means of predicting where crime is likely to occur. The

[16] Repeat victimisation and near repeat victimisation research has been conducted on crime types other than burglary dwelling, but the research on fluid spatial patterns and the foraging behaviour of offenders is more exclusive to just burglary dwelling.

current research, therefore, suggests that hotspots tend to be stable. This finding is consistent with other research that has shown the persistent longevity of crime at certain places (Brantingham and Brantingham 1981; Groff et al., 2008; Shaw and McKay, 1942; Weisburd et al., 2004). These findings on the stability of spatial patterns of crime, therefore, would appear to contradict the other findings about offender mobility and foraging behaviour and how spatial patterns of crime can be predicted.

There is, therefore, a need to examine the extent to which hotspots of crime are stable – both in terms of the stability of retrospective crime concentration and the stability of the predicted crime concentration. In relation to the research study's hypotheses (hypotheses 5) - *areas that are identified as hotspots of crime are places where the concentration of crime has been endured consistently for at least one year, and where the concentration of crime is likely to continue to persist into the future.*

It is argued that the stability of crime hotspots may be influenced by the volume of retrospective recorded crime data that is used to identify these spatial clusters. That is, does crime data for a short retrospective period of time identify different hotspots to those identified when a much larger retrospective period of crime data is used? Secondly, there is also a need to examine whether the hotspots identified using different retrospective periods of crime data identify different areas where crime is predicted to cluster in the future. For example, it is not known whether crime hotspots that occur in the future are more accurately identified using large or small (and more recent) retrospective periods of crime data. If crime hotspots are highly stable, regardless of the retrospective period of crime data that is used, this would make the selection of recorded crime data for hotspot analysis to be more straightforward. Also, if crime hotspots are highly stable it would mean that the targeting of police and public safety resources to these hotspots can be done in the confidence that this is where crime is most likely to occur in the future.

9.2. Chapter aims and structure

This research study aims to establish if the areas identified as hotspots of crime are places where the concentration of crime has been endured consistently over the retrospective past (i.e., the same hotspots are identified for different periods of retrospective crime data). The research study then examines if areas identified as hotspots are where the concentration of crime is likely to continue to persist in the future.

The chapter begins by describing the method that was used for measuring the stability of hotspots. This includes describing the data that were used and introduces the temporal stability index. Section 9.4 presents the results from the experiments that tested the temporal stability of hotspots under a variety of conditions. These findings are then interpreted to inform how they may influence the subsequent research studies in this thesis.

9.3. Method

Research study 4 found the Gi* statistic to be an effective means of identifying the spatial area of hotspots and for predicting spatial patterns of crime. These Gi*-identified *hot* areas can then be tested to determine the consistency in the volume of crime levels in these areas.

The stability of hotspots can be measured using a homogeneity index, a measure that is commonly used in social science research to summarise the distribution of data across nominal categories (Blau, 1977; Chainey and Ratcliffe, 2005; Gibbs and Martin, 1962). The homogeneity index (which, from this point, for crime hotspot analysis will be referred to as the temporal stability index) takes crime data that have been segmented into equal temporal periods and measures whether crime in the identified area is as a result of crime occurring over a short period, or whether the level of crime has been stable over a longer period. The temporal stability index (TSI) is computed using the formula $TSI = 1 - \sum p_i^{\,2}$ where p_i is the proportion of crime within each time category of the total number of crimes in the hotspot for the temporal period that is being analysed. The TSI is bounded by a minimum value of 0 and a maximum of $1 - 1/n_i$ where n_i is the total number of observed categories. A value close to the maximum indicates the data are heterogeneous, or that the level of crime is equally dispersed across the observed categories. A value of 0 indicates complete homogeneity or that the crimes are entirely concentrated within just one category (Haberman and Ratcliffe, 2012).

Table 9.1 illustrates an example of the TSI for two crime hotspots over a twenty-four-week period. The twenty-four-week period is divided into six equal four-week periods. The number of crimes in each hotspot for each four-week period is calculated. The volume of crime for each four-week period is expressed as a fraction of the total volume of crime for the entire period of interest (e.g., the number of crimes in the hotspot for the full twenty-four week period). This value for each four-week period is then squared, and these values summed. One minus the sum of this value represents the TSI value for crime in this hotspot over the twenty-four week period. In this example, the maximum TSI is 0.833 (1-(1/6)). The example in Table 9.1 shows there were 18 crimes in hotspot 1 and 18 crimes in hotspot 2 over the twenty-four week period. In hotspot 1, the eighteen crimes were spread equally over the six four-week periods, with three crimes being experienced in each period. Once the calculations described above have been made, this shows that the TSI for this hotspot is 0.833, equal to the maximum TSI value. This result illustrates that crime levels in this hotspot over the half-year period have been perfectly stable. In contrast, hotspot 2 has a TSI of 0 (following the same calculations as above). This is because all 18 crimes occurred in one of the six four-week periods.

Table 9.1. An example of the Temporal Stability Index for two hotspots

		Four-week periods							
		1	**2**	**3**	**4**	**5**	**6**	**Sum**	**TSI**
Hotspot 1	**n crimes in hotspot = 18**	3	3	3	3	3	3		0.8333
	Fraction of crime	0.167	0.167	0.167	0.167	0.167	0.167		
	Fraction of crime^2	0.0278	0.0278	0.0278	0.0278	0.0278	0.0278	0.1667	
Hotspot 2	**n crimes in hotspot = 18**	0	0	0	0	18	0		0.000
	Fraction of crime	0	0	0	0	1	0		
	Fraction of crime^2	0	0	0	0	1	0	1	

Experiments into the temporal stability of hotspots were conducted using Newcastle burglary dwelling, theft from the person and assault with injury data. These crime data were selected because they would provide the opportunity to examine differences in the temporal stability of hotspots for three distinct crime types. It was considered unnecessary to repeat the analysis for the Camden/Islington study area because results to date between the two study areas had been consistent, and instead a focus towards analysis of these three crime types in Newcastle would offer findings sufficient for the aims of this research study. Gi* hotspot analysis was used to identify the areas that could be defined as *hot*. A cell size of 150 m was used to produce Gi* hotspot maps in all experiments. Gi* hotspots were identified using a Bonferroni corrected 95% statistical significance threshold. If a hot cell was adjacent to one or more other hot cells, the cells were combined to form a single hotspot. The number of hotspots generated and the area of each were noted.

Gi* hotspot maps were generated using five input data periods:

- 12 months (the period covering the full data set, from 1st October 2009 to 30th September 2010
- 6 months (the six months prior to the 1st April 2010 measurement date)
- 3 months (the three months prior to the 1st April 2010 measurement date)
- 1 month (the month prior to the 1st April 2010 measurement date)
- The retrospective period for each crime type at which the clustering of crime was first statistically significant to 95%. For example, for burglary dwelling this was 2 weeks 4 days (from research study 1).

These separate temporal input periods were used to examine if the volume of crime data influenced the temporal stability of the hotspots they identified. Experiments were conducted by testing the temporal stability of hotspots over a one-year period, using thirteen four-week data periods. For example, Gi* hotspot maps generated using six months of input data were tested to examine the stability of crime levels in each identified hotspot across the thirteen four-week periods for the six months that were used to create the hotspot map and for the six months that followed. Similarly, Gi* hotspot maps generated using one month of input data were tested by using the thirteen four-week periods for the five months preceding the one month input data period, the one month period for which the hotspot map was generated, and for the six months that followed. For each Gi* hotspot map produced (using different retrospective periods of crime data), and for each crime type, it was expected that a number of hotspots would be identified. The TSI was calculated for each hotspot identified using the Gi* statistic. Table 9.2 provides an example of the TSI calculation process, showing the number of crimes recorded in a single identified hotspot for each four-week data period.

These experiments, therefore, would enable the stability of hotspots to be tested across a range of input data periods. This means that different retrospective input periods could be examined to determine if the hotspots identified from each different input period were places where crime had been endured at high levels for each of the four-week periods over the previous half year. Similarly, the retrospective input data could be used to determine if the hotspots these data identified were where crime continued to persist beyond the measurement date of the 1st April 2010. In turn, this methodological approach examined if there were differences in the stability of crime levels based on the retrospective input period of data that were used to create the hotspots. That is, the research study tested whether hotspots that were identified were stable and persistent, or whether they tended to be more fleeting and were based on short periodic increases in crime.

Table 9.2. Examples showing the arrangement of crime data for calculating the Temporal Stability Index of hotspots. Hotspots were identified using the defined period of input data, from which the volume of crime for thirteen four-week periods were then calculated.

Input data: 6 months		**4 week periods and count of crime**												
Hotspot	**n crimes in hotspot**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**
1	678	59	61	75	56	44	53	49	58	43	36	48	51	45
Input data: 1 month		**4 week periods and count of crime**												
Hotspot	**n crimes in hotspot**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**
1	636	56	55	70	53	39	48	49	56	43	35	44	47	41

TSI values were also calculated for each crime type for the whole Newcastle study area in order to compare the stability of crime levels in each individual hotspot to the stability of crime across the study area.

9.4. Results

9.4.1. Temporal stability of burglary dwelling hotspots

Table 9.3 shows the number of Newcastle Gi* burglary dwelling hotspots identified from each data input period and the proportion of crime that was identified in these hotspots. The number of burglary hotspots identified ranged from sixteen using a 12 month data input period, to nine hotspots using 1 month and nine hotspots using 2 weeks and 4 days of input data (2 weeks and 4 days was the retrospective point from the measurement date at which clustering was significant to 95%). The proportion of burglary dwellings in the hotspots (of all burglary dwellings in Newcastle) ranged from 25% using 12 months of input data to 10% using 1 month and 11% using 2 weeks and 4 days. These findings suggest that, using the example of Newcastle, larger retrospective periods of input data are better at identifying hotspots where the highest level of crime concentrates.

On occasion, the experiments identified small hotspots that were made up of single cells or a small number of cells. If a hotspot contained six or fewer crimes over the thirteen four-week data periods these were noted (as shown in Table 9.3) but were not included in the TSI analysis. The fact that an area with a small number of crimes was identified as a hotspot is of interest and will be discussed further in section 9.5. These *low volume* hotspots were not included in the TSI analysis because, following the recommendation of Haberman and Ratcliffe (2012), at least seven crimes would be required for an analysis of the temporal stability of crime over thirteen four-week temporal periods. For all except one burglary dwelling input data period, a number of these low volume hotspots were identified. These ranged from seven of the sixteen hotspots identified using 12 months of input data, to none identified using 2 weeks and 4 days of input data.

Table 9.3. The number of Gi* burglary dwelling hotspots identified for a range of input data periods, and the proportion of burglaries (compared to all burglary in Newcastle) contained within them. The number of hotspots with six or fewer offences across the analysis period is also listed.

Gi* hotspot input data period	**12 months**	**6months**	**3months**	**1month**	**2 weeks 4 days**
n of hotspots identified	16	12	14	9	9
Hotspots with <= 6 offences across analysis period	7	2	5	2	0
Proportion of crime in all hotspots	24.9%	22.0%	17.8%	10.4%	11.1%

The maximum possible TSI for a thirteen four-week analysis period was 0.923 (for all crime types analysed). Therefore, the closer a hotspot TSI value was to this maximum value, the more stable the hotspot, suggesting the hotspot experienced a consistently high level of crime. The minimum TSI was 0. The burglary dwelling TSI for the whole Newcastle study area was 0.920, suggesting that the volume of burglary dwelling offences across Newcastle had been very stable over the thirteen four-week analysis periods.

Table 9.4 lists the TSI results for Newcastle Gi* burglary dwelling hotspots for the range of data input periods that were used to create hotspot maps. Each hotspot that was identified is listed in Table 9.4 with a numerical identifier (e.g., the nine hotspots identified using 12 months of input data are listed in Table 9.4 from 1 to 9). Firstly, these results show that the hotspots identified using each of the input data periods were very stable. For example, the majority of hotspots had TSI values greater than 0.85. The highest TSI values tended to be for the largest hotspots that had been identified using at least 3 months of input data. For example, the TSI values for the largest hotspot identified from 12 months, 6 months and 3 months of input data ranged between 0.897 – 0.920.

Table 9.4. TSI values for Newcastle Gi* burglary dwelling hotspots using input data for (a) 12 months, (b) 6 months, (c) 3 months, (d) 1 month, and (e) 2 weeks and 4 days. The highest TSI value for each input period is shown in bold and the lowest TSI value is shown in italics. The number of burglary dwelling offences for the thirteen four-week analysis period was 1302.

(a) Input data period: 12 months

Hotspot #	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**
Hotspot TSI	**0.920**	0.879	0.880	0.861	0.862	0.874	0.857	0.872	*0.800*
n crimes in hotspot	114	43	23	29	21	25	21	24	10

(b) Input data period: 6 months

Hotspot #	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Hotspot TSI	**0.897**	0.839	0.865	0.867	0.836	0.870	0.837	0.865	0.805	*0.727*
n crimes in hotspot	102	22	17	30	15	20	19	28	13	16

(c) Input data period: 3 months

Hotspot #	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**
Hotspot TSI	**0.904**	0.867	0.840	0.858	0.827	0.857	0.875	0.826	*0.816*
n crimes in hotspot	64	26	26	25	14	21	19	11	7

(d) Input data period: 1 month

Hotspot #	**1**	**2**	**3**	**4**	**5**	**6**	**7**
Hotspot TSI	0.853	**0.871**	0.860	0.832	0.828	0.862	*0.735*
n crimes in hotspot	19	33	10	23	13	28	7

(e) Input data period: 2 weeks 4 days

Hotspot #	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**
Hotspot TSI	0.839	0.839	**0.875**	0.844	0.865	0.851	0.851	0.776	*0.760*
n crimes in hotspot	23	21	12	24	17	24	7	7	10

TSI values for hotspots identified using shorter periods of input data (as shown in Table 9.4) were also high (ranging between 0.735 – 0.875), but did not contain the high levels of crime experienced in the hotspots identified using longer periods of input data. For example, while the TSI value for the main hotspot (hotspot # 1) that was identified using 2 weeks and 4 days of input data was 0.839, only 26 crimes took place in this hotspot over the thirteen four-week periods. This compared to 114 crimes that were experienced in the main hotspot identified using 12 months of input data, over the thirteen four-week periods. The lowest TSI values were consistently for hotspots that contained the lowest number of offences over the thirteen four-week analysis period.

The results into the temporal stability of burglary dwelling hotspots suggests that if at least three months of burglary dwelling data are used, burglary dwelling hotspots that experience persistent levels of crime can be effectively identified. Shorter periods of retrospective burglary dwelling data are less likely to identify the full extent of the main areas where high levels of crime are likely to concentrate in the future, although the temporal distribution of crime in these identified hotspots are also highly stable.

9.4.2. Temporal stability of theft from the person hotspots

Table 9.5 shows the number of Newcastle Gi* theft from the person hotspots identified from each data input period and the proportion of crime that was identified in these hotspots. Only one hotspot was identified using each of the data input periods. The proportion of theft from the person offences identified in this single hotspot (of all theft from the person offences in Newcastle) ranged from 52% using 12 months of input data to 49.2% using 2 weeks and 5 days of input data (2 weeks and 5 days was the retrospective point from the measurement date at which clustering was statistically significant to 95%). This suggests, using the example of Newcastle, there was little difference across the input data periods in identifying hotspots where the highest level of theft from the person offences concentrated, albeit the higher values were for the hotspot identified using the longer retrospective periods of input data.

Table 9.6 lists the TSI results for Newcastle Gi* theft from the person hotspots for the range of data input periods. The study area theft from the person TSI was 0.911. This suggests that the volume of theft from the person offences across Newcastle was very stable over the thirteen four-week analysis period. The TSI value for the single hotspot was 0.913 for all of the input data periods. The volume of crime in the hotspot area that each data input period identified also ranged very little (between 385 to 405 offences), with the highest values being for the hotspots identified using 6 months and 12 months of input data. These results suggest the single theft from the person hotspot that the Gi* analysis identified is very stable, and that even a short period of retrospective data was effective in identifying where this type of crime has consistently concentrated.

Table 9.5. The number of Gi* theft from the person hotspots identified for a range of input data periods, and the proportion of theft from the person offences (compared to all theft from the person offences in Newcastle) contained within them.. The number of hotspots with 6 or fewer offences across the analysis period are also listed.

Gi* hotspot input data period	**12 months**	**6months**	**3months**	**1month**	**2 weeks 5 days**
n of hotspots identified	1	1	1	1	1
Hotspots with <= 6 offences across analysis period	0	0	0	0	0
Proportion of crime in all hotspots	51.9%	51.7%	49.3%	49.6%	49.2%

Table 9.6. TSI values for Newcastle Gi* theft from the person hotspots using (a) 12 months, (b) 6 months, (c) 3 months, (d) 1 month, and (e) 2 weeks and 4 days of input data. The number of theft from the person offences for the thirteen four-week analysis period was 781.

(a) Input data period: 12 months

Hotspot #	**1**
Hotspot TSI	0.913
n of crimes in hotspot	405

(b) Input data period: 6 months

Hotspot #	**1**
Hotspot TSI	0.913
n of crimes in hotspot	404

(c) Input data period: 3 months

Hotspot #	**1**
Hotspot TSI	0.913
n of crimes in hotspot	385

(d) Input data period: 1 month

Hotspot #	**1**
Hotspot TSI	0.913
n of crimes in hotspot	387

(e) Input data period: 2 weeks 5 days

Hotspot #	**1**
Hotspot TSI	0.913
n of crimes in hotspot	385

9.4.3. Temporal stability of assault with injury hotspots

Table 9.7 shows the number of Newcastle Gi* assault with injury hotspots identified from each data input period and the proportion of assaults (of all assaults in Newcastle) that was identified in these hotspots. Only one hotspot was identified using data for the 12 month and 6 month input periods, compared to two, four, and six hotspots using 3 months, 1 month, and 1 week and 5 days of input data respectively (1 week and 5 days was the retrospective point from the measurement date at which clustering was statistically significant to 95%). However, two of the four hotspots identified using 1 month of input data and four of the six hotspots identified using 1 week and 5 days of input data contained six or fewer offences over the thirteen four-week analysis period. The proportion of crime identified in the hotspots ranged from 37% using 12 months, 6 months and 3 months of input data to 30% using 1 week and 5 days of input data. These results suggest, using the example of Newcastle, that the main hotspots were identified using each of the input data periods, albeit the higher values were for the hotspot identified using the longer periods of input data.

Table 9.7. The number of Gi* assault with injury hotspots identified for a range of data input periods, and the proportion of assault with injury offences (compared to all assault with injury offences in Newcastle) contained within them. The number of hotspots with 6 or fewer offences across the analysis period are also listed.

Gi* hotspot input data period	**12 months**	**6months**	**3months**	**1month**	**1 week 5 days**
n of hotspots identified	1	1	2	4	6
Hotspots with <= 6 offences across analysis period	0	0	0	2	4
Proportion of crime in all hotspots	36.9%	37.3%	36.9%	35.4%	29.5%

Table 9.8 lists the TSI results for Newcastle Gi* assault with injury hotspots for the range of data input periods. The study area assault with injury TSI was 0.921. This suggests that the volume of assault with injury offences across Newcastle has been very stable over the thirteen four-week analysis period. The TSI value for the main hotspot identified by each of the data input periods ranged from 0.919 to 0.921. The volume of crime identified in this main hotspot was more varied, ranging from 511 identified offences in the hotspot identified using 1 week and 5 days of input data to 686 offences in the hotspot area

identified using 6 months of input data. These results suggest that this hotspot was very stable, with crime persisting across the full thirteen four-week analysis period. The results also suggest that only a short period of retrospective data were required to identify the main area where these assaults were most persistent, but that the area identified using a larger period of input data was more effective in identifying the full extent of the area where high levels of assaults were most experienced.

Table 9.8. TSI values for Newcastle Gi* assault with injury hotspots using (a) 12 months, (b) 6 months, (c) 3 months, (d) 1 month, and (e) 1 week and 5 days of input data. The highest TSI value for each input period is shown in bold and the lowest TSI value is shown in italics. The number of theft from the person offences for the thirteen four-week analysis period was 1838.

(a) Input data period: 12 months

Hotspot #	**1**
Hotspot TSI	0.920
n of crimes in hotspot	678

(b) Input data period: 6 months

Hotspot #	**1**
Hotspot TSI	0.920
n of crimes in hotspot	686

(c) Input data period: 3 months

Hotspot #	**1**	**2**
Hotspot TSI	**0.920**	*0.741*
n of crimes in hotspot	670	9

(d) Input data period: 1 months

Hotspot #	**1**	**2**
Hotspot TSI	**0.921**	*0.776*
n of crimes in hotspot	636	7

(e) Input data period: 1 week 5 days

Hotspot #	**1**	**2**
Hotspot TSI	**0.919**	*0.852*
n of crimes in hotspot	511	13

9.5. Interpretation and conclusions from research study 5

This research study examined the temporal stability of hotspots by analysing whether areas identified using hotspot analysis are the places where crime has previously persisted at high levels and are where crime is likely to persist at high levels in the future. The research also compared whether there was a difference in the stability of crime levels in hotspots when hotspots were identified using short and longer retrospective periods of crime data.

The research findings suggest that hotspots identified using the Gi* statistic display high levels of temporal stability. The research findings did, though, show some differences between crime types, with hotspots of burglary dwelling tending to vary most in their temporal stability, whereas hotspots of thefts from the person and hotspots of assault with injury were highly stable. Another difference that distinguished these crime types were the number of hotspots that were identified – many more hotspots of burglary dwelling were identified than those for theft from the person and assaults. The larger number of burglary dwelling hotspots and the lower stability in burglary levels in these hotspots would then suggest that spatial patterns of burglary tend to be more fluid than those for theft from the person and assaults. This finding offers some support to the research from Johnson and Bowers (2004b) and Johnson et al. (2008a) where they describe burglary hotspots as often moving location, reflecting the foraging behaviour of offenders. However, when at least three months of retrospective crime data were used, burglary dwelling hotspots did display high levels of temporal stability.

The results for theft from the person and assaults with injury indicated that a short retrospective period of crime data, such as data for the temporal period from which clustering is statistically significant, can be sufficient in identifying the places where high levels of these types of crime are likely to continue to persist. Findings for theft from the person and assaults in this research study (and in the previous research studies of this thesis) also showed that the geographic concentration of these two crime types were spatially compact to a small number of hotspots. This may then suggest that when crime patterns exhibit a high level of spatial concentration, the levels of crime in these areas are more likely to be highly stable. In turn, this suggests that rather than the commission of offences being explained by offender foraging behaviour (that is used to explain why hotspots may move location), a different type of offending behaviour is occurring in temporally stable hotspots.

In the next research study the theoretical principles for explaining the spatial behaviour of offenders and the differences in the temporal stability of spatial crime patterns are examined further. The focus of this next study is to compare the spatial prediction performance of the prospective mapping approach and the use of recent individual incidents of crime, to hotspot analysis.

10. Research study 6: Examining the influence that recent incidents of crime have on predicting different future periods of crime

10.1. Introduction

This research study tests the hypothesis (hypothesis 6) that recent incidents of crime provide an effective means of accurately predicting the immediate future, but the accuracy in these predictions reduces for longer periods of the future.

Study 5 of the research has illustrated the stable nature of hotspots, yet other commentators (e.g., Johnson and Bowers, 2004b; Johnson et al., 2008a) have suggested that hotspots are more *slippery* in nature, moving to other (albeit nearby) locations. The findings that suggest hotspots to be unstable draw on the theoretical arguments of the boost account and foraging nature of offenders: where offenders operate in spates, boosted from the successful commission of a very recent offence to quickly commit additional offences, before moving on to other areas. The boost account and foraging behaviour have also been shown to be the principal theoretical reasons for explaining the spatial patterns of repeat and near repeat victimisation (Bowers and Johnson, 2004; Johnson et al., 2008a).

The empirical findings into the patterns of repeat and near repeat victimisation have shown previous incidents to be very good predictors of where crime is likely to occur in the future (Johnson and Bowers, 2004a). This prediction quality is now often used in policing for targeting operational resources to prevent more burglaries occurring at recently burgled and neighbouring properties by minimising this predicted heightened risk (Chainey, 2012b; Fielding and Jones, 2012). The prediction quality recognised from the patterns of repeats and near repeats also forms the foundation of the prospective mapping technique of spatial crime prediction (Bowers et al., 2004). The prospective mapping approach creates a risk surface based on the spatial distribution of recent incidents to predict where crime is likely to occur. It has been suggested that this prospective mapping approach performs better in predicting spatial patterns of crime than hotspot analysis, however, the only comparison to date has been against KDE (Johnson et al., 2008b). The research by Johnson et al. (2008b) also only compared prospective mapping to KDE for burglary dwelling, and only for differences in predicting crime for the seven days following the date of the production of the mapping outputs. Indeed, to date, most attention has been placed on using patterns of repeats and near repeats to predict where crime is likely to occur very soon (i.e., over the next few days) rather than for predictions for any further point in the future (i.e., over the next month). No consideration was given either in Johnson et al.'s (2008b) study into whether the volume of retrospective crime data used in the production of the prospective mapping and KDE mapping outputs would have an influence in the prediction performance of the mapping outputs.

Study 1 of this PhD research showed that in many cases a number of weeks of retrospective crime data are required before statistical evidence of clustering is present. Furthermore, several other parts of this PhD research have shown that if retrospective crime data are used for hotspot analysis before the point that clustering has been evident, then any resulting map is unlikely to be useful in predicting where future offences are likely to occur. That is, to identify where crime is likely to concentrate in the future, hotspots need to be identified in the first place from the available retrospective data. If only very recent incidents are used, this may be at the point before clustering is statistically evident in the retrospective data and, hence, any attempts to produce a meaningful hotspot map would be futile. Studies 2 to 5 of this research have shown that once hotspots are identified using retrospective crime data, the hotspots are very effective in predicting spatial patterns of crime.

The collective findings from both previous and the current research, therefore, suggest the following: recent incidents that are modelled following repeat and near repeat patterning principles (i.e., prospective mapping) appear to be very good in predicting where crime is likely to occur over the next few days; and hotspot analysis using the Gi* statistic appears to be very good in predicting where high crime levels are likely to persist. What is not known, however, is whether the predictive quality of prospective mapping using recent incidents is consistently more or less accurate than Gi* hotspot analysis for a range of temporal prediction periods (i.e., for the next few days, the next week, the next month and other periods beyond).

10.2. Chapter aims and structure

This research study examines the prediction performance of the prospective mapping technique in comparison to the Gi* statistic. In this study, the approach to measuring prediction performance examines the differences between the techniques for predicting the *immediate* future (i.e., for the next 7 days) in comparison to the *near* future (the period beyond the next seven days). The study also focuses on using only recent incidents of crime to generate mapping output. This is with the hypothesis of determining whether recent incidents of crime not only provide an effective means of accurately predicting the immediate future, but also whether the predictions using recent incidents are equally as accurate beyond this immediate future. The findings from research part 4 on the prediction performance of Gi* hotspot analysis provide the benchmark against which the predictions using a prospective mapping approach will be made. Additional Gi* hotspot analysis is provided where necessary in order to compare against the spatial prediction performance of prospective mapping.

The chapter begins by describing the method that was used for producing prospective mapping output and the metrics that were used for measuring prediction performance. The results section then presents these findings, comparing them against Gi* hotspot analysis. The results are then interpreted to inform how they influence subsequent research parts.

10.3. Method

The Vigilance Modeller prospective mapping tool was used in this study to generate predictive mapping output. The Vigilance Modeller is based on an algorithm that incorporates the patterning principles of repeats and near repeats. The Vigilance Modeller uses data on crimes that have taken place over a one week period to generate mapping output that identifies places at the highest risk of crime for each week thereafter.

Prospective mapping output was generated (using the Vigilance Modeller) from Newcastle burglary dwelling, theft from the person and assaults with injury data for the input period of 25-31 March 2010 (i.e., using a measurement date of the 1 April 2010). The output from the Vigilance Modeller is a series of cell-based risk surfaces, with each cell scored in relation to the level of crime risk at each location. Each cell is 100 m x 100 m in size. There is no published account of why cells of this size are used, nor a published account of analysis that has tested different cell sizes. The higher the risk value in each cell, the more likely it is that crime will occur in this area. These risk values for cells are generated from the location of recent incidents, with the influence of each incident decaying with distance. This means that when only a short retrospective period of crime data is used (and thus the volume of crime data is likely to be low), many areas across the study area will not be populated with cells, and therefore zero risk of future crime is associated with these areas.

Each prospective risk surface for each crime type were measured for their ability to predict crime that took place in the 7 days, 4 weeks, 12 weeks and 24 weeks following the measurement dates of the 1 April 2010. Prediction performance was calculated using a hit rate, accuracy concentration curves, the area under the curve and the Crime Prediction Index. These results were compared against equivalent Gi* prediction performance results.

10.4. Results

Figure 10.1 shows the prospective mapping output that was generated from Newcastle burglary dwelling, theft from the person and assault with injury data for the input period of 25-31 March 2010. These maps show the areas of risk, based on where incidents of crime occurred during the 25-31 March 2010 period. Table 10.1 lists the sizes of areas covered by each prospective mapping output and the number and proportion of offences that were committed in the prospective mapping risk areas, by output period. The area of the prospective mapping outputs, as shown in Figure 10.1, varied in size, from 2% of the entire study area for theft from the person to 12% for burglary dwelling and theft from the person.

The results in Table 10.1 show that for each crime type, the prediction performance of the prospective mapping approach was highest for the 7 days following the input data period. The proportion of offences that were committed in the area identified using prospective mapping then reduced for each longer output period. For example, the prospective mapping area identified 62% of the burglaries that were committed in the 7 days immediately following the input data period. This then reduced to 43%, 38% and 38% for the 4 weeks, 12 weeks and 24 weeks output periods respectively. One exception to this declining prediction performance was assault with injury, where the 24 weeks output period recorded the second highest proportion of offences.

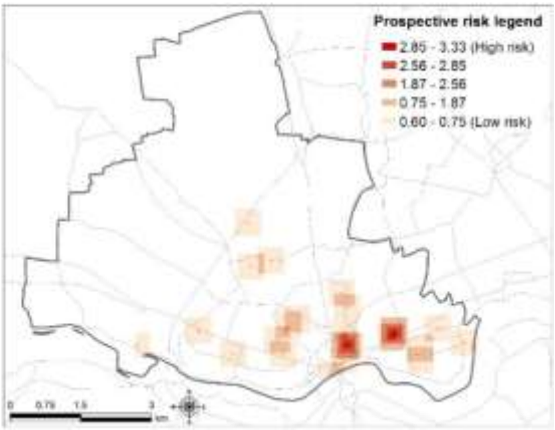

(a) Burglary dwelling

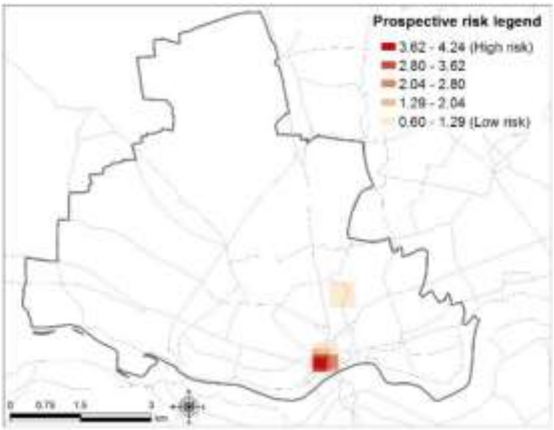

(b) Theft from the person

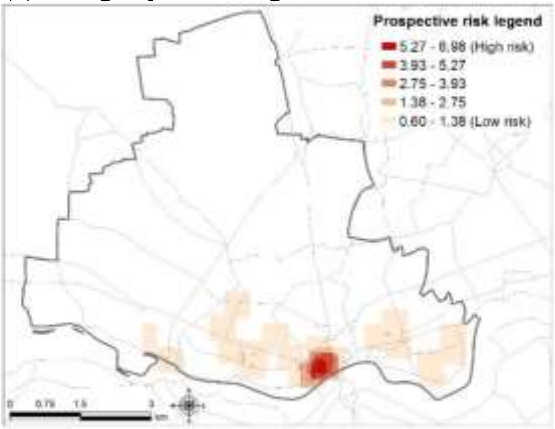

(c) Assault with injury

Figure 10.1. Prospective mapping output for Newcastle (a) burglary dwelling, (b) theft from the person, and (c) assault with injury

Table 10.1. Sizes of area covered by each prospective mapping output and the number and proportion of offences committed in the prospective mapping risk areas, by output period. Values in bold relate to the highest value for the proportion of offences in prospective mapping areas for each crime type.

Crime type	Prospective mapping % of study area	Output period	n of offences during output period	n offences in prospective mapping area	% of offences in prospective mapping area
Burglary dwelling	11.9%	7 days	26	16	**61.5%**
		4 weeks	138	59	42.8%
		12 weeks	329	123	37.5%
		24 weeks	691	259	37.5%
Theft from the person	1.6%	7 days	19	14	**73.7%**
		4 weeks	60	37	61.7%
		13 weeks	187	88	47.1%
		24 weeks	324	152	46.9%
Assault with injury	11.7%	7 days	37	25	**67.6%**
		4 weeks	154	95	61.7%
		12 weeks	459	274	59.7%
		24 weeks	835	527	63.1%

Figure 10.2 shows accuracy concentration curves for Newcastle prospective mapping outputs of burglary dwelling, theft from the person and assaults with injury. The curves trend towards random variation from the point representing the extent of the spatial coverage of the prospective mapping output. For example, for burglary dwelling (Figure 10.2a), from 12% of the study area (which represents the spatial coverage of the prospective risk surface), the accuracy concentration curve follows a pattern representing spatially random variation. Sub-sections of the accuracy concentration curves for the study area coverage of between 0% to 5% allow for a more detailed examination of the differences by crime type and output period.

The charts in Figure 10.2 consistently show that the steepest curve gradients were for the 7 days output period. It was also noticeable that the curves for theft from the person and assault with injury had almost vertical gradients, reflecting the high prediction performance of the prospective mapping technique. In addition, the curve for the 7 days output period for burglary dwelling was much steeper up to the 12% of offences point, than curves for other periods. This indicates that prospective mapping output is sensitive to repeats and near repeats. That is, the highest risk scores are generated for those areas where a burglary has recently occurred, and where, as in this example, subsequent burglaries have occurred in these high-risk areas in the week following the initial incidents.

Table 10.2 lists CPI values for sub-sections of the accuracy concentration curves for prospective mapping output generated for burglary dwelling, theft from the person and assaults with injury in Newcastle. These results show that the CPI values were consistently highest for the 7 days output period and tended to be lowest for the longer output periods.

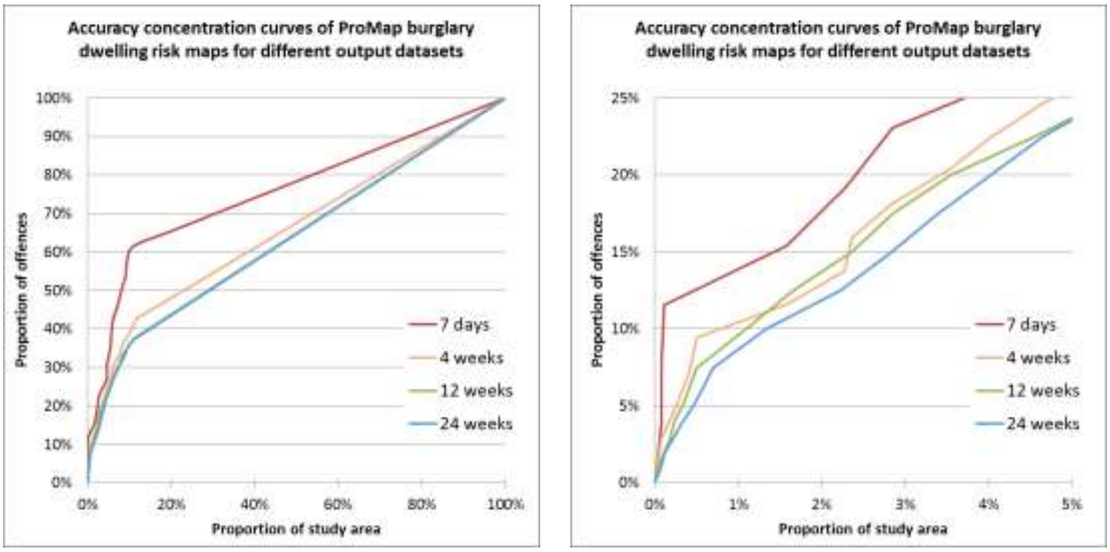

(a) Burglary dwelling

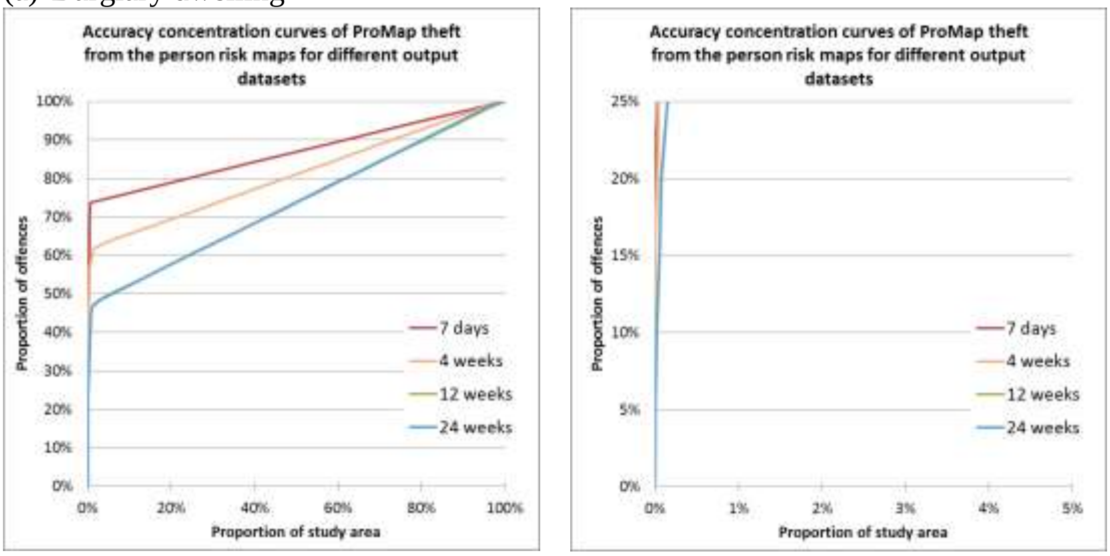

(b) Theft from the person

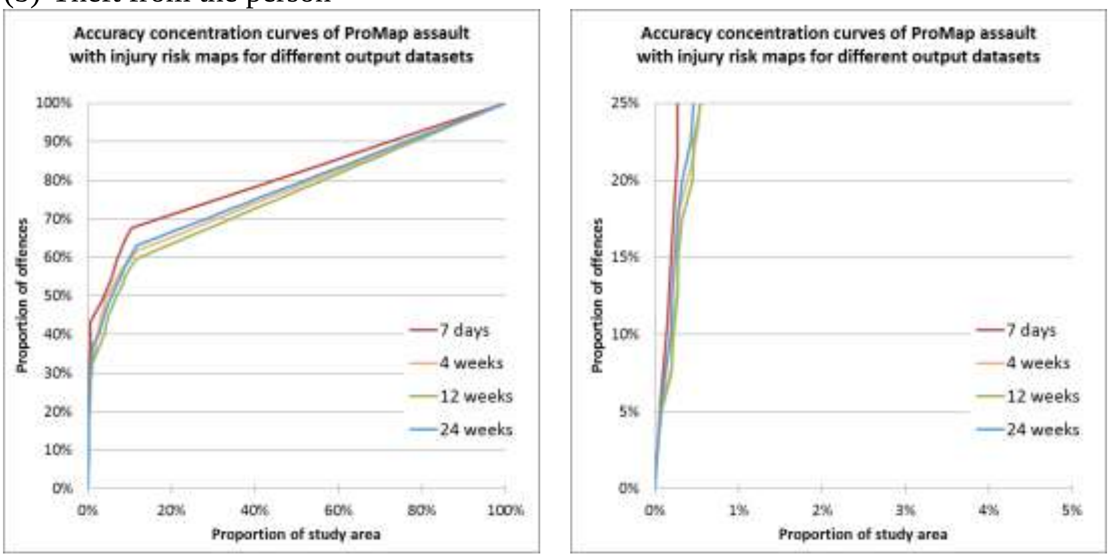

(c) Assault with injury

Figure 10.2. Accuracy concentration curves for 100% and 5% of the proportion of the study area of prospective mapping output of Newcastle (a) burglary dwelling, (b) theft

from the person, and (c) assault with injury, for data outputs periods of 7 days, 4 weeks, 12 weeks and 24 weeks

Table 10.2. Crime Prediction Index values for different sub-sections of accuracy concentration curves, for 7 days, 4 weeks, 12 weeks and 24 weeks out periods and for (a) burglary dwelling, (b) theft from the person and (c) assault with injury. Values in bold represent the largest CPI values and those in italics the smallest CPI values.

(a) Burglary dwelling

Output period	0.5% x 5% CPI	1% x 10% CPI	5% x 25% CPI	10% x 50% CPI	20% x 80% CPI	100% x 100% CPI
7 days	**0.902**	**0.935**	**0.772**	**0.658**	**0.605**	**0.758**
4 weeks	0.810	0.733	0.635	0.490	0.432	0.661
12 weeks	0.710	0.631	0.608	0.455	0.392	0.638
24 weeks	*0.602*	*0.528*	*0.553*	*0.440*	*0.389*	*0.637*

(b) Theft from the person

Output period	0.5% x 5% CPI	1% x 10% CPI	5% x 25% CPI	10% x 50% CPI	20% x 80% CPI	100% x 100% CPI
7 days	**0.991**	**0.993**	**0.998**	**0.994**	**0.947**	**0.868**
4 weeks	0.986	0.992	0.997	0.993	0.812	0.810
12 weeks	*0.985*	0.992	*0.987*	*0.938*	*0.627*	*0.736*
24 weeks	0.990	*0.990*	0.991	0.959	0.646	0.737

(c) Assault with injury

Output period	0.5% x 5% CPI	1% x 10% CPI	5% x 25% CPI	10% x 50% CPI	20% x 80% CPI	100% x 100% CPI
7 days	**0.941**	**0.934**	**0.969**	**0.943**	**0.779**	**0.806**
4 weeks	0.940	0.908	0.952	0.887	0.688	0.771
12 weeks	*0.933*	*0.897*	*0.948*	*0.836*	*0.651*	*0.758*
24 weeks	0.940	0.923	0.958	0.881	0.690	0.777

Similar to how statistical significance thresholds define the extent of hotspots using the Gi* statistic, prospective mapping generates output that has a limited areal coverage representing risk – the gird cells with values greater than 0 are where crime is predicted to occur (as shown in Figure 10.1). As reported in Table 10.1, the extent of the prospective risk mapping surfaces was 11.9% for burglary dwelling, 1.6% for theft from the person, and 11.7% for assault with injury. CPI values for these prospective mapping coverage areas were calculated in order to offer direct comparison to CPI values for Gi* output. A Bonferroni corrected 95% significance threshold was applied to Gi* hotspot mapping output for calculating CPI values for burglary dwelling, theft from the person and assault with injury. Three months of input data were used to create the Gi* hotspot maps, and CPI values were generated for each of the 7 days, 4 weeks, 12 weeks and 24 weeks output periods.

Table 10.3. CPI values for Gi* (based on three months of input data with hotspots defined using a Bonferroni corrected p=0.05 threshold) and for prospective mapping (based on seven days of input data), for output periods of 7 days, 4 weeks, 12 weeks and 24 weeks, and for burglary dwelling, theft from the person and assault with injury. Values in bold represent the largest CPI values.

Output period	**Crime type**	**Gi* CPI**	**Prospective mapping CPI**
7 days	**Burglary dwelling**	0.598	**0.641**
	Theft from the person	**0.995**	**0.995**
	Assault with injury	**0.991**	0.942
4 weeks	**Burglary dwelling**	**0.942**	0.493
	Theft from the person	**0.995**	0.992
	Assault with injury	**0.991**	0.874
12 weeks	**Burglary dwelling**	**0.950**	0.456
	Theft from the person	**0.996**	0.991
	Assault with injury	**0.993**	0.812
24 weeks	**Burglary dwelling**	**0.963**	0.429
	Theft from the person	**0.997**	0.990
	Assault with injury	**0.995**	0.857

Table 10.3 shows that in comparison to Gi* hotspot analysis CPI values, prospective mapping CPI values for the 7 day output periods were higher than the comparable Gi* CPI values for burglary dwelling, were the same for theft from the person, but were lower than the comparable Gi* CPI values for assault with injury. None of the prospective mapping CPI values for output periods of 4, 12 or 24 weeks for burglary dwelling, and

assaults with injury were higher than the Gi* CPI equivalent values, but for theft from the person the Gi* and prospective mapping CPI values were very similar. These results suggest that the predictive performance of the prospective mapping technique is best for predicting the immediate future rather than predicting where crime is likely to occur in the more distant future. While the prospective mapping technique was very good in some situations for predicting where crime may occur for more distant periods in the future, these predictions were not consistently better than comparable spatial predictions of crime generated using Gi* hotspot analysis.

10.5. Interpretation and conclusions from research study 6

In this research study, the hypothesis that was tested (hypothesis 6) was whether recent incidents of crime provide an effective means of accurately predicting the immediate future, but the accuracy in these predictions reduces for longer periods of the future. To test this hypothesis involved examining the prediction performance of prospective mapping for different output periods – the immediate future (predicting where crime was likely to occur over the next 7 days), and for periods thereafter (4 weeks, 12 weeks and 24 weeks). The prediction performance of the prospective mapping output was then compared to the prediction performance of the Gi* statistic.

The results from this research show that the prospective mapping approach, using the example of crime data for Newcastle, was effective at predicting the immediate future (i.e., within the next 7 days), but was less effective at predicting where crimes occurred for more distant periods of the future. That is, in reference to the research hypothesis, recent incidents of crime provide an effective means of accurately predicting the immediate future, but the strength in these predictions reduces for longer periods of the future. For two of the three crime types (burglary dwelling and theft from the person), the prospective mapping approach performed better than the Gi* statistic for predicting where crime occurred within the next seven days. However, for longer periods beyond seven days, the prediction performance of the Gi* statistic was just as good for theft from the person, but better for burglary dwelling and assaults with injury than that for prospective mapping.

The prospective mapping approach is based on the frequently observed finding of repeat and near repeat patterns in crime data for many different crime types. This approach has also become a popular technique that is used in practice for predicting spatial patterns of crime and determining where to target operational police resources. The findings from this research study suggest that the use of recent incidents is better than Gi* hotspot analysis for predicting where crime happens in the immediate future for offences such as burglary dwelling, but no better than Gi* hotspot analysis for predicting theft from the person and assault with injury in this same immediate timeframe. In addition, Gi* hotspot analysis was found to be consistently more effective at predicting where crime is likely to occur over longer periods. These findings have practice implications for targeting policing and crime prevention resources that will be discussed further in chapter 12. The theoretical implications from the findings in these differences will also be considered in chapter 12.

11. Research study 7: Examining the use of geographically weighted regression for helping to explain why hotspots exist, and for informing spatial predictions of crime

11.1. Introduction

The research across studies 1 to 6 has shown that the presence of crime hotspots in crime data can be determined by the application of a simple statistical process (the nearest neighbour index), that the Gi* statistic provides an effective means of helping remove much of the ambiguity in defining the spatial areas that are hotspots (separating statistically what is *hot* from what is *not hot*) and that the areas it identifies produce effective predictions of where crime is likely to concentrate in the future. The research has also shown that crime hotspots display temporally stable patterns, and that while recent events can have a bigger influence in determining where crime is likely to occur in the immediate future for certain crime types, the performance of these predictions can quickly reduce for more distant temporal periods of the future. All the analytical processes behind these findings have used a single source of information to guide these predictions – recorded crime data.

A sound theoretical framework has been developed to help explain the spatial distribution of crime, and was described in chapter 2. In summary, the routine activity approach provides a model to predict if a crime has the necessary components to occur, involving the presence of a likely offender and a suitable target, and the absence of a capable guardian, meeting in time and space. The rational choice perspective enables us to determine some of the thinking behind an offender's ultimate decision to commit a crime, by understanding how offenders weigh up the risks, efforts and rewards during the crime commission process. In the case of property offences, offenders are drawn to targets and products that meet appealing CRAVED qualities. Crime pattern theory then helps to explain the spatial and temporal patterning of crime, drawing together the concepts of awareness space, opportunity space and the least effort principle. Together, these theories help explain why the spatial distribution of crime tends not to be uniform or random. Added to this theoretical framework are the concepts of offender foraging and the boost account: combined, these two theories state that short periodic spates of crime that are committed by offenders are due to the boost in confidence they receive to return to locations where they have previously been successful in the commission of crime (i.e., crimes committed in the last few days). These boost and foraging behavioural principles also help explain how offenders take advantage of similar opportunities nearby before moving on in order to avoid capture, and/or because the opportunity resource has been exhausted. This PhD research has also shown that hotspots experience very stable patterns of crime, with these areas over time being the places to which many offenders are likely to be drawn because of the enduring crime opportunity characteristics these areas tend to retain. To some degree these crime opportunity characteristics can be explained by the background norms of these areas, drawing from the original social disorganisation concepts of the Chicago School, and from the flag hypothesis that suggests certain targets signal an opportunity for crime due to some enduring quality that attracts offenders to these targets.

In terms of the latter, the characteristics of an area, determined by its background norms and crime opportunity space, can present a rich and attractive range of targets. Recognising how the characteristics of an area may influence its crime levels then suggests there is value in using information, other than retrospective crime data, to help explain why these places are where crime tends to concentrate. That is, while hotspot analysis effectively identifies where crime is likely to take place in the future, it does not identify the physical, social or economic conditions that give rise to why these concentrations occur at these specific locations. If the conditions that explain the presence of crime hotspots can be identified using data modelling, then these conditions could be used alongside or in replacement of retrospective crime data to produce predictions on where crime is likely to occur.

Chapter 2 introduced geographically weighted regression (GWR) as an analytical framework for helping to determine the conditions that explain why crime concentrates in certain locations, and if these explanations vary spatially. The ability to be able to identify and explain the spatial variation in relationships between crime and other characteristics of the area may indeed help identify that the reasons for the presence of one hotspot in a study area may be different to the reasons for another in the same study area. Using this analytical framework, GWR could provide an effective means of developing statistical insight into determining those conditions (hereafter referred to as variables) that explain why crime concentrates in some areas and not in others. This knowledge can then inform a more data-rich approach to hotspot analysis through the inclusion of these variables alongside, or in replacement of retrospective crime data. This PhD research has already shown the value of retrospective crime data for predicting spatial patterns of crime. Therefore, the inclusion of any additional variables requires careful scrutiny to ensure they provide the relevant level of precise spatial scale to support hotspot analysis, are data that are easy to access and consistent in format for use by crime researchers, and can be effectively embedded into a hotspot modelling process that improves spatial crime predictions. In relation to the hypotheses proposed in chapter 3, this research study tests the hypothesis (hypothesis 7) that GWR provides an effective means of determining at the local level the reasons why hotspots exist, and why these explanatory variables vary between hotspots.

An additional or alternative analytical framework that uses GWR involves examining how a change in explanatory variables can influence a change in future crime levels in order to support long-term predictions of crime. This alternative analytical framework would involve using explanatory variables to test *what if* scenarios and direct strategic policy in crime reduction. For example, this might involve predicting how crime may change based on a change in an explanatory variable. In relation to the hypotheses previously proposed, this research study also tests whether GWR analysis can be effectively used for supporting long-term predictions of crime by examining how a change in explanatory variables can influence a change in future crime levels (hypothesis 8).

11.2. Chapter aims and structure

This research study examines whether GWR provides an analytical framework for helping to determine why crime hotspots exist, and if the reasons for the presence of hotspots across a study area vary. This research involves scrutinising the application of GWR using crime data, and at the spatial scale that hotspots are identified. The research study also examines if GWR provides a further means to support spatial crime prediction by examining how a change in explanatory variables may influence a change in crime.

The chapter begins by describing the methodological application of GWR and the necessary conditions to consider in the treatment of spatial data for regression analysis. This methodological process includes a description of Ordinary Least Squares (OLS) regression, Gaussian and Poisson approaches to linear regression and GWR, the spatial scale of data required for GWR, and the types of transformations and standardisation methods that may be required to prepare spatial data for spatial regression analysis.

Outlined in the method section is the step-by-step process involved in calibrating a GWR model, beginning with the analysis of an OLS model and a number of statistical diagnostic tests to determine how well the OLS model performs. The results from the OLS regression analysis and the statistical diagnostic tests also determine whether the model is suitable for GWR analysis. The GWR analysis process is then applied to two examples of crime data to test whether a GWR analysis can help determine why hotspots exist. Using these two examples, the research examines if the reasons for explaining hotspots within the same study area vary, if these findings can support a more data rich hotspot modelling framework (using explanatory variables alongside crime data), and/or if the results can be used for supporting long-term predictions of crime by testing *what if* scenarios.

11.3. Method

11.3.1. Modelling approaches and data

Ordinary Least Squares (OLS) regression is a widely used technique for identifying relationships between a dependent variable and one or more explanatory variables. For example, it might be used to explore the relationship between crime and variables such as poverty, deprivation and unemployment. An OLS regression produces a global model, explaining the relationships (and their significance) between the dependent and explanatory variables. However, in reality, these variables are likely to vary spatially, with the relationship between the variables in one area possibly being stronger (or weaker) than the relationship in another area. This, therefore, suggests that any analysis that aims to examine relationships between different types of spatial data should be sensitive to exploring whether the relationships vary spatially.

An assumption in the application of OLS regression is that the error terms are independent of one another. However, with most forms of spatial data this assumption is violated, because (in following with Tobler's First Law of Geography) observations close together will tend to be similar than those further apart (Tobler, 1970). This relationship between geographical observations is effectively the positive spatial autocorrelation that is observed in many forms of geographic data, or in relation to this research, the observation that crime tends to cluster into hotspots. The effect of spatial autocorrelation in an OLS regression model, therefore, results in the violation of the assumption of independently distributed errors of the model's variables (Haining, 1990), the underestimation of the standard errors when positive spatial autocorrelation is present in the residuals, and the potential inflation of Type 1 errors (i.e., the incorrect rejection of the null hypothesis) (Legendre, 1993). These problems result in unreliable statistical inferences and can lead to conclusions that a supposed effect or relationship exists between variables when in fact they do not (Anselin and Griffith, 1988). These issues do not mean that we cannot apply regression analysis to spatial data, but their application does require these spatial effects to be explicitly incorporated into the specification of the regression model, with the model used being appropriate for spatial data (Fotheringham et al., 2002; Voss et al. 2006). In practice, this requires the application of several statistical diagnostic tests for model significance and model bias (e.g., tests for spatial stationarity and heteroscedasticity).

In chapter 2, GWR and a number of other spatial regression techniques were introduced. These other spatial regression techniques included spatial lag and generalised additive models, with the results from numerous studies (Bini et al., 2009; Nakaya et al., 2005; Wang et al., 2005; Zhang et al., 2005, Zhang et al., 2009) concluding that GWR consistently performed as well or better than other alternatives for modelling spatial non-stationarity and for accounting for spatial autocorrelation in parameter residuals. GWR is also available to practitioners in ArcGIS and as a free standalone software application. The findings from other research into the modelling performance of GWR (as described in chapter 2) and its availability in software were the main reasons for its selection in this research for modelling spatially varying relationships between crime and other variables. There has also been mention by some commentators on the use of GWR for prediction. Harris et al.'s (2010) discussion of GWR for spatial prediction referred to the use of the technique as a spatial interpolation method (estimating values for unsampled locations based on the values at sampled locations), while Zhang et al. (2009) referred to the use of GWR for producing accurate predictions of the response variable from the explanatory variable inputs.

GWR is not without its flaws, with some critics advocating other techniques as being superior to GWR. Most recently, this has included several researchers advocating Bayesian spatially varying coefficient models by comparing the processes and outputs from these models to GWR (Waller et al., 2007; Wheeler and Calder, 2007; Wheeler and Waller, 2009). Outcomes from the research by others suggest that Bayesian modelling produces more accurate inferences than GWR, and that GWR output can result in strong dependence between estimated parameter surfaces (even when these parameters are associated to be independent of each other in the modelling generation process).

However, these Bayesian modelling approaches are significantly more expensive in computer processing time, and require a greater level of sophistication in determining inputs to the modelling procedure (Waller et al., 2007; Wheeler and Waller, 2009). While GWR may, therefore, not provide a perfect spatial regression modelling process, it is considered to be sufficiently robust for crime analysis practitioners to use, and, hence the decision to focus attention on this technique in this research study. Possible advantages in using Bayesian spatially varying coefficient models and others will be discussed further in chapter 12 following the results from this chapter.

Many of the previous studies in this PhD research have examined techniques for their spatial crime prediction performance. The aim of this research study was not necessarily to measure the prediction performance of the model results, but rather to identify an analytical process that would help explain why hotspots exist, explore whether these results could potentially inform a more data-rich hotspot analysis process, and identify if GWR could be used to test *what if* strategic prediction scenarios using variables that have been determined to spatially correlate with crime patterns.

This research study involved using GWR to develop two types of models. The first used an exploratory approach to examine the spatial relationships between burglary dwelling and several explanatory variables. The second used a hypothesis testing approach to examine the spatial relationships between violent assaults and several variables that were theoretically considered to be associated with hotspots of this type of crime. The crime data in both types of models were treated as the dependent variables. The study area for both models was Newcastle because of the availability of data for these two contrasting crime types and because of the availability of data that could be used as explanatory variables. Analysis using burglary dwelling and assault with injury data, under two different modelling processes, was considered to be sufficient for examining if a GWR modelling processes could help explain why hotspots exists and inform spatial predictions of crime. The entire one year period of crime data for Newcastle was used (October 2009 – September 2010). The full Newcastle crime data set was used in order to help address any issues of low counts for small areas. Data for use as the explanatory variables were collected from a number of agencies including the Office for National Statistics (data from the 2011 Census and the 2007 Index of Deprivation), Newcastle City Council (data on licensed premises), and Northumbria Police (data on the home addresses of offenders). Table 11.1 lists the twenty-five data variables that were sourced from the Office for

National Statistics and the Newcastle agencies. The table includes a description of each variable and the spatial scale for which they were available. GWR in ESRI ArcGIS v10.1 and GWR v4.0 (developed by Nakaya et al., 2012) were used to perform the OLS and GWR analyses. The ArcGIS GWR application was used because it provides more diagnostic statistical results than GWR v4.0 to aid model selection and allows for direct interaction between the GWR modelling process and the mapping of results. However, ArcGIS only contains the functionality to apply Gaussian GWR modelling, with any Poisson GWR modelling requiring the application of the standalone GWR software.

Table 11.1. Explanatory variables used in the regression analysis of burglary dwelling and assault with injury in Newcastle

Data variable	**Source**	**Description**	**Spatial scale**
Household density	Census 2011	Households per km²	Output areas (OAs)
Population density	Census 2011	Population per km²	OAs
Low level of education	Census 2011	Percentage of population with no qualifications or have only achieved Level 1 status (1+'O' level passes, 1+CSE/GCSE any grades, NVQ level 1, Foundation GNVQ)	OAs
Asian population	Census 2011	Percentage of population that are Asian	OAs
Black population	Census 2011	Percentage of population that are Black	OAs
Chinese population	Census 2011	Percentage of population that are Black	OAs
Student population	Census 2011	Percentage of the normally resident population that are students (attending a higher education establishment)	OAs
Owned housing	Census 2011	Percentage of households where the household head owns the property (with or without a mortgage)	OAs
Socially rented housing	Census 2011	Percentage of households rented from the local Council or Housing Association	OAs
Private rented housing	Census 2011	Percentage of households privately rented	OAs
Population age 20 to 29	Census 2011	Percentage of the normally resident population aged 20-29	OAs
Population age 30 to 39	Census 2011	Percentage of the normally resident population aged 30-39	OAs
Population age 40 to 49	Census 2011	Percentage of the normally resident population aged 40-49	OAs

Population age 50 to 59	Census 2011	Percentage of the normally resident population aged 50-59	OAs
Population age 60 to 69	Census 2011	Percentage of the normally resident population aged 60-69	OAs
Population age 70 and over	Census 2011	Percentage of the normally resident population aged 70 and over	OAs
Population born in the UK	Census 2011	Percentage of the population that were born in the UK	OAs
Population that arrived from outside the UK between 2001-2003	Census 2011	Percentage of the population that arrived in the UK between 2001 - 2003	OAs
Population that arrived from outside the UK between 2004-2006	Census 2011	Percentage of the population that arrived in the UK between 2004 - 2006	OAs
Population that arrived from outside the UK between 2007-2009	Census 2011	Percentage of the population that arrived in the UK between 2007 - 2009	OAs
Serious Acquisitive Crime (SAC) offenders	Northumbria Police (2010)	The number of SAC offenders per 1000 population. SAC offences include burglary dwelling, theft of vehicles, theft from vehicles, and robbery (from people and businesses)	Home address
Income Deprivation	Index of Deprivation (2010)	The percentage of children that live in families that are income deprived (i.e., in receipt of Income Support, Income based Jobseeker's Allowance, Working Families' Tax Credit or Disabled Person's Tax Credit below a given threshold)	Lower super output areas (LSOAs)
Education and skills deprivation	Index of Deprivation (2010)	A score relating to the education and skills attained by children and young people	LSOAs
Living environment deprivation	Index of Deprivation (2010)	A score relating to quality of the indoors living environment of housing	LSOAs
Licensed premises	Newcastle City Council (2010)	Registered licensed premises, including nightclubs, bars, pubs and off-licenses	Address of premise

The research was organised into four parts:

- An assessment of the spatial scale of crime data and the data variables against which these data were to be modelled
- An assessment of the data variables to determine whether Gaussian and Poisson GWR modelling was required, and if these data needed to be transformed to make them suitable for analysis. This also included a critical assessment of these modelling processes to crime data
- An analysis of OLS and diagnostic statistical results to determine whether the variables selected for a model were suitable for GWR analysis
- A GWR analysis of dependent and explanatory variables, including an assessment of the spatial bandwidth for modelling spatially varying relationships.

11.3.2. Assessment of spatial scale

The first part of the regression analysis process involved an assessment of the spatial scale of crime data and the explanatory data variables against which crime data were to be analysed. Throughout the research study, regression analysis was to be applied to the entire study area, albeit with the aim of conducting analysis at a level of spatial precision that would allow hotspots to be distinguished from each other and from those areas that were not hotspots. Hotspots were identified using the Gi* statistic following the method that was used in research study 4 (chapter 8). Gi* hotspots were identified using a 95% significance level to maximise the areal size of those areas statistically determined as hotspots. The assessment of spatial scale involved an examination of the size of hotspots and whether the explanatory data variables were of a geographic scale that was compatible for exploring relationships with these hotspots.

11.3.3. Gaussian GWR, Poisson GWR and data transformations

Consideration was given in each modelling process to determine whether a Gaussian or Poisson GWR modelling approach would be applied and if any transformation of variables were required. Gaussian GWR assumes both dependent and explanatory variables to be normally distributed, and is applied to continuous data. Poisson GWR is suitable for count data that follow a Poisson distribution (i.e., where there may be many observations with low counts, and fewer with high counts). A natural logarithm transformation was applied to data where necessary to help address any non-normal data distributions and where certain statistical diagnostic results advised on transforming data in this way. A full description of the statistical diagnostic tests is provided in section 11.3.3.

In the first case, a Gaussian approach was applied to the GWR modelling of burglary dwelling data. This followed the conversion of burglary dwelling count data to burglary rates (burglaries per 1000 households per annum), and an assessment of whether a transformation of these rates was required. The conversion of the burglary dwelling point data to a rate also meant these data had to be aggregated to Census output areas for which household data were available. An analysis of the explanatory variables was then performed to determine if each followed a normal distribution. Where they did not, the variable was also transformed using a natural log approach. All explanatory variables were then used in the first part of the modelling process to explore which were statistically significant. This selection was then refined through an iterative process to determine which model was most suitable for GWR analysis. A full description of this iterative process is described in the results section.

Assault with injury crime data cannot be suitably converted to a rate because an area's population does not accurately represent the on-street population where many of these types of violent incidents occur. Census workplace statistics are also not a substitute measure for the on-street population as they do not account for trips such as those for shopping or trips associated with people travelling to town centres to enjoy the area's night-time economy (Chainey and Desyllas, 2008). Assault with injury data therefore remained in count format.

The GWR modelling of assault with injury data followed a hypothesis testing approach. This approach involved choosing variables based on sound theoretical principles and supported with empirical research. In the first instance, a model was created using only licensed premises data as an explanatory variable. This was for two reasons. Firstly, licensed premises data were available as point data, meaning that counts of assaults and counts of licensed premises could be aggregated to a precise user defined set of grid cells rather than constrained to geographic administrative areas of varying sizes and shapes (e.g., output areas). Secondly, there is much evidence that suggests a relationship between alcohol consumption and violent assaults, and that such assaults are concentrated in the same places where pubs, bars and nightclubs (collectively referred to hereafter as licensed premises) are similarly located (for example, see Babor et al., 2003; Graham and

Homel, 2008; Maguire and Hopkins, 2003). The relationship between assaults and licensed premises is both associated with violence inside the premises and the violence that occurs outside these premises. In terms of the later, the contemporary view is that violent incidents that occur on the streets outside licensed premises is directly correlated with drinking alcohol purchased from licensed premises (Institute of Alcohol Studies, 2013). That is, if it were not for the presence of these licensed premises, assaults would not take place at the volumes and in the location that is observed. The licensed premises data supplied by Newcastle City Council contained both on-licenses (pubs, bars, restaurants, nightclubs) and off-licenses (shops selling alcohol). A subset of these data was created containing only night-time economy related off-licenses (i.e., pubs, bars and nightclubs) to allow for further analysis between assaults and this specific variable.

Four different grid cell sizes (150 m, 300 m, 500 m and 1000 m) were used in the analysis of assaults and licensed premises to determine how many cells contained zero assaults. At present, there is little guidance on the size of cells to use for GWR analysis, and instead the selection of suitable cell size requires some experience of the study area and consideration of the interplay between the possible spatial relationships of variables. This consideration of the interplay between variables requires an assessment that ensures that the cell size is large enough for cells to display a range of values from low to high, but not too large to negate an examination of local spatial relationships. The application of GWR, therefore, requires the researcher to make an assessment of the interaction between the variables that are to be analysed, drawing from their experience to do so, rather than relying on an analytic solution to determine the optimal size of geographic unit. If a large number of cells contained zero values, this would be problematic for the modelling process. A zero-inflated modelling approach could be applied, but to be fully robust required specialist software. Additionally, the research focus was to explore how functionality available to a wide audience of crime analysts (e.g., GWR in ArcGIS or the free to download GWR v4.0 software) could be used for modelling spatially varying relationships. The zero-inflated approach is, though, discussed further in the results section and in chapter 12.

Both assaults with injury and licensed premises data were in count form. This meant that a Poisson GWR model was applied in the first instance. The results of the Poisson GWR model then informed whether other explanatory variables were required for the model, and hence whether the aggregation of data could remain in grid cell format or needed to be aggregated to Census geography to make it compatible with the spatial scale of these other explanatory variables. The results on whether to apply grid cells or Census geographic units then determined whether a Gaussian GWR modelling approach would be more suitable following the necessary transformation of the assault data and any explanatory variables.

11.3.4. OLS regression and diagnostic statistical tests

The first stage of GWR modelling required an OLS regression analysis to test if the variables chosen provided a suitable model for explaining the presence of the dependent variable. Where variables were not significant, they were excluded from the model until only those variables that were significantly correlated remained. Another part of this initial process involved performing several statistical diagnostic tests to determine if the model was suitable for GWR, or if there was a need to revise the variables that were included in the model (and if there was a need to revise the hypothesis that explained the presence of the dependent variable). The tests were grouped into three categories: model performance, model significance, and model bias.

Tests for model performance determined if there was, in a statistical sense, a relationship between the dependent variable and the explanatory variables. The test results of interest were the adjusted R^2, the coefficient for the explanatory variable, the Koenker statistic, and the probability or robust probability measures of relationship significance. When the model consisted of multiple variables, the variation inflation factor (VIF) also required assessment. Each of these test statistics were used as follows:

- Adjusted R^2: this statistic helps describe the goodness of fit of a model and the percentage of the explained variance. The adjusted R^2 adjusts for the number of explanatory terms in a model. Unlike R^2, the adjusted R^2 increases only if a new variable that has been added improves the model more than would be expected by chance. In the case of using multiple variables to explain the dependent variable, the adjusted R^2 statistic is a better measure to use than R^2.
- Coefficient of the explanatory variable: this represents the strength and type of relationship between the explanatory variable and the dependent variable. For instance, it can be used to suggest that for every one unit change in the explanatory variable, the dependent variable increases by the amount equivalent to the coefficient.
- Koenker statistic: this is primarily used as a model significance measure (explained in the following section), but its use for model performance is to determine whether

probability or robust probability measures should be used. If the Koenker statistic is significant then the robust probability results should be used rather than the standard probability results.

- Probability/Robust probability: this is a measure of the significance of the relationship between the dependent variable and each explanatory variable. An explanatory variable's coefficient means little if the relationship is not statistically significant.
- Variation Inflation Factor (VIF): the VIF quantifies the severity of multicollinearity in the model. If the modelling begins by including a number of explanatory variables, it is possible that probability results show that several of the explanatory variables were not statistically significant. This means that consideration should now be given to removing these variables from the OLS model that do not appear to be contributing in any way in helping to explain the variation in the dependent variable. To help with this process of selecting variables to remove from the model, the VIF for each variable indicates if the variable is not contributing to the model. If the VIF value for a variable is greater than 7.5 this suggests the variable is redundant in the model and should be removed. The same procedure applies when an explanatory variable is significant, but has a VIF value that is greater than 7.5. This is likely to be because this variable offers a similar explanation to another (statistically significant) variable, and therefore one (the variable with the highest VIF) should be removed.

Tests for model significance help determine if the model is a good candidate for GWR. For this purpose, interest is in the results of three tests: the Koenker statistic, the Joint F statistic, and the Joint Wald statistic:

- Koenker statistic: this is a measure of stationarity in the model. That is, it determines whether the explanatory variables in the model have a consistent relationship with the dependent variable. There are two ways to consider the behaviour of these relationships. In spatial terms, the spatial processes represented by the explanatory variables behave the same everywhere. In data terms, it also means that the variation in the relationship between the dependent variable and each explanatory variable does not change with changes in the magnitude of the explanatory variable (i.e., there is no heteroscedasticity in the model). For example, the relationship between the dependent variable and the explanatory variables is just as accurate in areas where observations for the explanatory variable are low and where these observations are high. If the Koenker statistic is significant, it suggests there is heteroscedasticity and/or non-stationarity in the model, meaning the relationship between the dependent and explanatory variables vary across space, and that the model is therefore a good candidate for GWR.

- Joint F statistic and Joint Wald statistic: these statistics are measures of overall model statistical significance. If the Koenker statistic for the model is statistically significant, the Joint F statistic's result is unreliable, and therefore the Joint Wald statistic should be used instead.

The third group of tests are for model bias. Firstly, this requires an analysis of the residuals (the difference between the observed value of the dependent variable and the estimated value - based on the statistical relationship with the explanatory variables) to determine if they are normally distributed. Secondly, the residuals are then tested to determine if they are spatially clustered. The two tests for these are the Jarque-Bera statistic and Moran's I:

- Jarque-Bera statistic: if this statistic is significant, it suggests the model is biased. This would be, for example, where the model performs well for low values but does not perform well for high values. If the model is biased then there is the option to transform the variables or remove influential outliers (Rosenshein et al., 2011), or include other explanatory variables in the model
- Moran's I: this second test for model bias identifies whether the residuals are spatially clustered. If they are then this suggests that a key variable is missing from the model (to explain the observed spatial patterns in the dependent variable). If the Moran's I result suggests evidence of clustering in the residuals, then mapping the residuals thematically may help identify the explanatory variable that is missing. That is, where residual values are high, and clustered, what else is present in this area that could help explain the presence of the phenomenon, in addition to or in replacement of the existing explanatory variables?

The initial OLS regression and statistical diagnostic process identifies if the model is suitable for GWR. If not, for example, the Jarque-Bera statistic is significant, a further analysis of explanatory variables is required before progressing to developing a GWR model. Once a suitable OLS model has been identified, the explanatory variables from this model can be applied to a GWR model. This involves a straightforward process of using the statistically significant explanatory variables from the final OLS regression model for inclusion in the GWR model. At this stage it is also useful to record the Corrected Akaike Information Criterion coefficient (AICc) value for the OLS regression model that, for most software applications of GWR, is calculated as part of this initial diagnostic process and is also calculated for subsequent GWR models. The AICc will indicate if the GWR model performs better than the OLS model. The lower the AICc, the better the model.

11.3.5. GWR modelling and spatial bandwidths

The GWR model requires the specification of a bandwidth that defines the size of the neighbourhood over which relationships between the dependent variable and the explanatory variables will be explored. A Cross Validation method and the AICc are considered to be useful bandwidth optimisation methods (both built into the two GWR software options), although Fotheringham et al. (2002) recommend the use of the AICc to determine optimal bandwidth size for GWR models. The GWR modelling process results in the automated calculation of the bandwidth size, determined by minimising the AICc value. An alternative approach is for the user to specify a bandwidth size, with the user then evaluating the performance of this GWR model by comparing AICc measures. A fixed or adaptive bandwidth can also be applied. The fixed approach determines an optimal single bandwidth size for exploring the relationship between variables across the whole study area. The adaptive approach determines an optimal number of neighbours to include, therefore, allowing the metric size of the bandwidth to vary across space (i.e., smaller in metric size for small geographic units of analysis and larger in metric size for larger geographic units of analysis). For both the burglary dwelling and assault with injury GWR models an AICc bandwidth optimisation process was used, applying both fixed and adaptive bandwidth processes. The GWR models that produced the lowest AICc values and the highest adjusted R^2 values were the ones selected for further analysis.

The GWR modelling process generates a number of outputs for mapping. The first of these is a *local* R^2 value that indicates the relationship between the explanatory variables and the dependent variable and how this relationship varies across the study area. The local R^2 value may indicate that the relationship between these variables is stronger in some areas than others. In the areas where the relationship is weak, other variables, other than the explanatory variables used in the model, are likely to explain the patterning of the dependent variable in these areas. A second value generated is a condition number. This indicates if the results are unstable due to local multicollinearity. A value greater than 30 for any of the geographic areas suggests that the result for this cell is unreliable. In addition to this value, the Coefficient of Standard Error that is also generated from the

GWR model for each geographic area helps determine if the local R^2 values are reliable. Where the Coefficient of Standard Error value is small, confidence can be placed in the relationship between the dependent and explanatory variables. Maps of GWR output were generated for both the burglary dwelling and assault with injury models.

11.4. Results

11.4.1. Using GWR to determine at the local level the reasons why hotspots exist, and that explanatory variables vary between hotspots

In research study 4 (Chapter 8), Gi* hotspot maps were produced for burglary dwelling and assault with injury for the Newcastle study area using six months of input data. For research study 7 (the current chapter), the analysis of the spatial relationships between crime and the explanatory variables used twelve months of data in order to help minimise the problem of many small areas containing zero or low counts of crime. In the first instance, Gi* hotspot maps of burglary dwelling and assault with injury were produced using six months of data and twelve months of data to compare if the areas identified were different. Each of the Gi* hotspot maps were generated using a cell size of 150 m. The maps in Figure 11.1 show there to be no major differences between the areas identified as hotspots.

Twenty hotspots of burglary dwelling were identified using twelve months of input data. The size of these hotspots ranged from 0.023 km^2 to 0.855 km^2, and had a mean size of 0.162 km^2 (see Table 11.2). Only one hotspot of assault with injury was identified in the study area, and was 1.148 km^2 in size.

To perform regression analysis of burglary dwelling, these crime data had to be aggregated to the smallest geographic area for which the explanatory variables were available. While many of these data variables were available at the Census output area level (OAs), deprivation data were only available at the lower super output area level (LSOAs). This meant that in order to make direct comparisons between the crime data and the explanatory variables, all data had to be aggregated to LSOAs. In addition, an analysis of twelve months of crime data aggregated to OAs revealed 36% of OAs contained no burglaries and 44% contained no assaults with injury. If OAs had been used as the geographic unit of analysis, the high proportion of OAs with zero counts would likely to have caused difficulties.

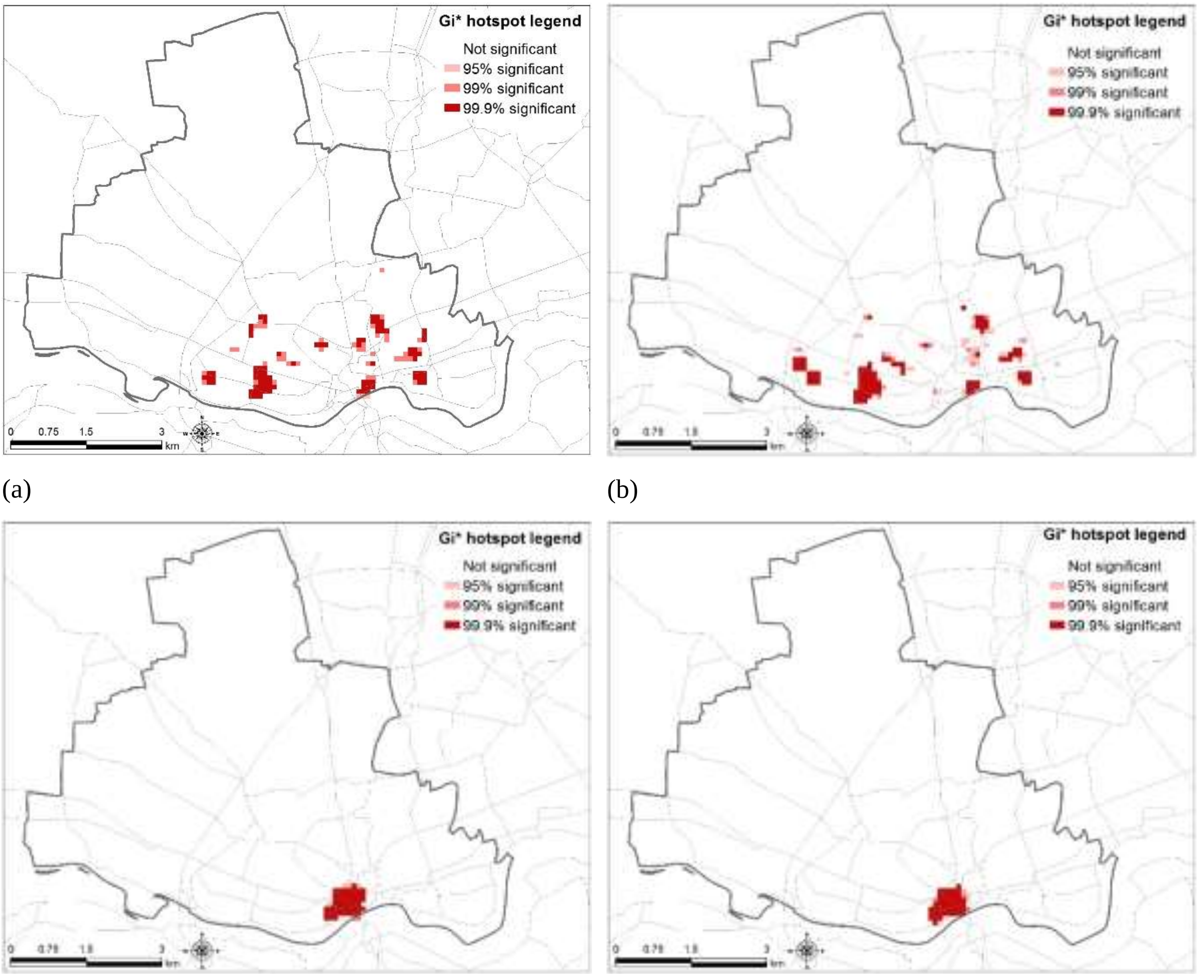

(a) (b) (c) (d)

Figure 11.1. Gi* hotspot maps of burglary dwelling generated using (a) six months of input data, (b) twelve months of input data, and of assault with injury generated using (c) six months of input data, (d) twelve months of input data

Table 11.2. Mean, minimum, maximum and standard deviation of the size of Gi* hotspots for burglary dwelling and assault with injury in Newcastle

Crime type	Mean size of hotspots (km²)	Minimum hotspot size (km²)	Maximum hotspot size (km²)	Standard deviation of hotspot size (km²)
Burglary dwelling	0.162	0.023	0.855	0.197
Assault with injury	1.148	1.148	1.148	-

Table 11.3 lists the mean, minimum, maximum and standard deviation of the size of the Newcastle LSOAs. The figures show that these areas were of a much larger spatial size than the areas that were identified as hotspots. For example, the mean size of the LSOAs was over four times larger than the mean size of the burglary dwelling hotspots.

Additionally, Figure 11.2 shows a number of the burglary dwelling hotspots and their respective LSOAs. The figure illustrates that many of the hotspots were of a much smaller size than the geographic units of the explanatory variables. Figure 11.2 also shows that many of the Gi* hotspots had a spatial coverage that overlapped several LSOAs, meaning that any analysis that explored the relationship between burglary dwelling and the explanatory variables would relate to the whole of each LSOA that overlapped the hotspots, rather than the specific areas that had been identified as hotspots. These findings mean that a comparison between the specific causes of these spatially precise burglary dwelling hotspots was not possible due to this mismatch between spatial scales of the geographic units of analysis.

Table 11.3. Mean, minimum, maximum and standard deviation of the size of Newcastle LSOAs

Crime type	**Mean size (km^2)**	**Minimum size (km^2)**	**Maximum size (km^2)**	**Standard deviation of size (km^2)**
LSOAs	0.663	0.071	15.545	1.594

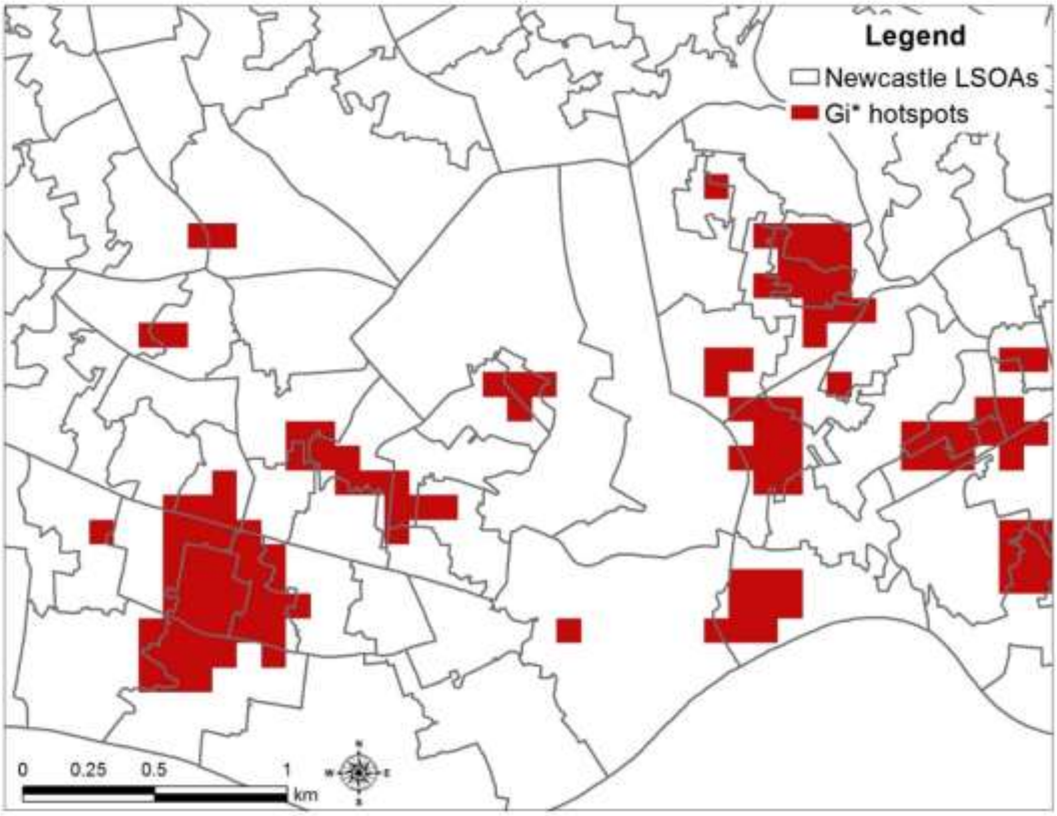

Figure 11.2. Burglary dwelling Gi* hotspots in Newcastle compared to their respective LSOAs

Assault with injury regression analysis followed a hypothesis testing approach rather than an exploratory approach. The initial hypothesis was that the spatial distribution of assault with injury could be explained in relation to the distribution of licensed premises. Licensed premises data were available in point form, meaning that data could be calibrated for analysis using grid cells rather than being constrained to LSOAs. The availability of point format data for both data sets, therefore, offered promise that the analysis into the relationship between assaults and licensed premises could be explained using these more specific grid cells, thus allowing for analysis of whether the influence of licensed premises varied among hotspots.

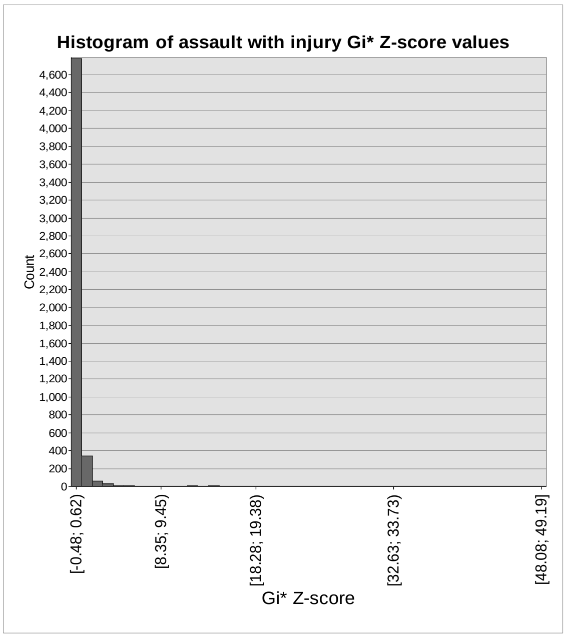

(a)

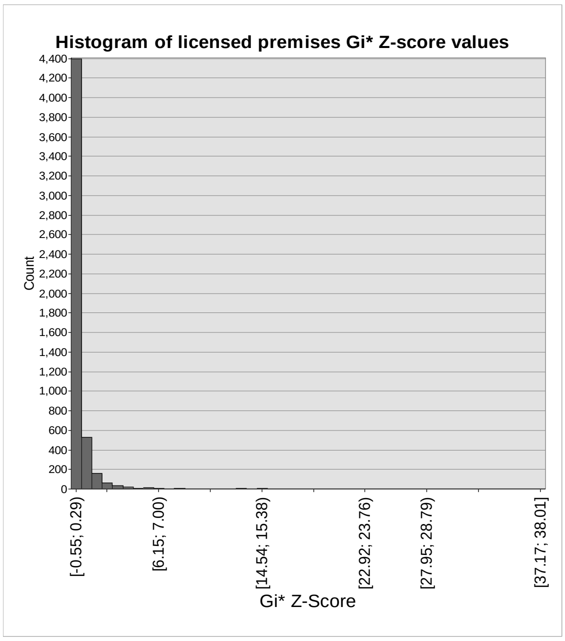

(b)

Figure 11.3. Histograms of Gi* values for (a) assault with injury and (b) licensed premises in Newcastle

In the first instance, the use of Gi* values that represented the distribution of assault with injury and Gi* values representing the distribution of licensed premises, using the same grid cells for each, was considered the optimal basis for comparing the relationship between these two variables. This is not only because this method is the means by which hotspots were identified, but also because the approach resulted in each cell within the study area containing a value relating to the spatial distribution of crime and the spatial distribution of licensed premises. An approach using Gi* values for both data variables would also avoid any problems created from many geographic units containing zero values. Figure 11.3 shows histograms of the distribution of the Gi* values for assault with injury and licensed premises in Newcastle. This figure shows the data to be highly skewed, following a Poisson distribution. This would then suggest a Poisson GWR model would be most suitable for these data. However, this presented three difficulties.

Firstly, Poisson regression analysis cannot be applied to variables that have negative values (as in the case of Gi* values). An option then would be to apply a Gaussian regression analysis approach after a log transformation of these Gi* values, but transformation of negative values is not possible. An option to address this is to add a constant to each value before applying the log transformation.

Secondly, Poisson regression assumes that the occurrence of an event (e.g., crime) will not make another event more or less likely. However, co-dependence among events is an established feature of crime, as evident in the patterns of repeat victimisation and near repeat victimisation as explained and empirically evidenced by the boost account and offender foraging theories in chapters 2 and 10.

Third is the problem of overdispersion. An alternative to using the Gi* values was to use the count of the number of assaults and the count of the number of licensed premise in each grid cell, and apply a Poisson regression against these variables. Poisson regression assumes the mean is equal to the variance in the variables. Where the variance is greater than the mean, overdispersion is present. The presence of overdispersion could cause an underestimation of the standard errors in the estimated variable, meaning that a variable may appear to be a significant predictor of the dependent variable, when in fact it is not. This issue of overdispersion is a common feature in spatially distributed crime data, particularly for crimes which are highly clustered (such as assault with injury offences), because the data displays some locations where the volume of crime is high, and many places where it is zero.

The initial methodological proposal was to explore the use of different cell sizes for spatial regression. In the first instance this approach would involve choosing a suitable cell size that minimised the number of cells with zero counts and balanced this with examining spatially varying relationships at the scale relevant to hotspot analysis. The Gi* hotspot maps presented in Figure 11.1 used a cell size of 150 m. If this cell size was used as the basis for exploring the relationship between assaults and licensed premises, it would be extremely problematic since 85% of cells covering the study area contained a zero count. Table 11.4 shows the problem of zero counts extended to cell sizes of 300 m, 500 m and 1000 m. In the latter case, a 1000 m cell size was representative of the entire size of the assault with injury Gi* hotspot, but even at this spatial scale, 43% of cells covering the Newcastle study area contained zero counts.

Table 11.4. The proportion of cells containing zero counts of assault with injury offences for different cell sizes

	150m	300m	500m	1000m
Proportion (and n) of cells with zero counts	85% (4509)	70% (968)	56% (296)	43% (65)

To overcome the problem of overdispersion a number of solutions are available. One involves addressing the large number of zeros in the data by producing a zero-inflated model (i.e., by adding a value of one to all observed incidents). However, for the analysis of crime data it is important to distinguish areas where there has been no crime from those where there has been crime, particularly for rare crimes (e.g., robbery, sexual assaults). Some zero-inflated models specify separating the analysis of the cells that originally contained a value of zero from those that contained a value greater than zero. However, the analysis and interpretation of spatial data that are split in this way can be cumbersome without specialised software for treating zero-inflated models. A second solution that is designed to help address the issue of overdispersion involves applying a Poisson Gamma regression model. However, this modelling approach was not available in the software used in this analysis. A third solution considered was to apply a log transformation to the assaults and licensed premises data, and apply a Gaussian regression analysis. However, even with large grid cells of 1000 m in size, a log transformation would not be possible for 43% of the data (the cells that contained a zero count), therefore, this approach was not pursued. A fourth option was to reduce the number of geographic units that contained zeros in the assault with injury data by increasing the size of the geographic units that were used for analysis.

It was decided that three approaches would be applied to progress the analysis of the relationships between assault with injury offences and licensed premises. In the first instance, an attempt would be made to apply a Poisson GWR regression analysis to the 1000 m grid cells using the GWR software and examine the results. The second approach involved removing the majority of 1000 m cells that contained zero values, but to still produce a study area that was spatially contiguous, and apply Poisson GWR to these data. This resulted in creating a *City* grid lattice of 82 cells covering the southern areas of the Newcastle district study area (see Figure 11.4a). Only four of these 82 cells contained zero values. The distribution of the number of assaults in each cell still followed a Poisson distribution (see Figure 11.4b). The third approach involved aggregating assaults and licensed premises data to LSOAs, applying a log transformation to these two variables, and performing a Gaussian GWR analysis using ArcGIS. Aggregating assaults data in this way resulted in only nine out of 173 LSOAs containing zero counts (representing 5% of the geographic units in the study area). As previously noted, an assaults rate was not used because suitable denominator data, such as an on-street pedestrian count, was not available. Application of the Gaussian GWR model in ArcGIS would also allow a more comprehensive statistical diagnostic assessment of the relationship between assaults and licensed premises (the standalone GWR software does not include the same extensive range of measures for model performance, significance and bias described above). Additionally, this approach using LSOAs would allow an assessment of whether licensed premises alone were enough to explain the spatial distribution of assaults with injury and if model performance improved with the addition of other variables. In turn, this would permit a more straightforward interactive regression modelling approach in ArcGIS by refining hypotheses and selecting other variables alongside, or in replacement of licensed premises data.

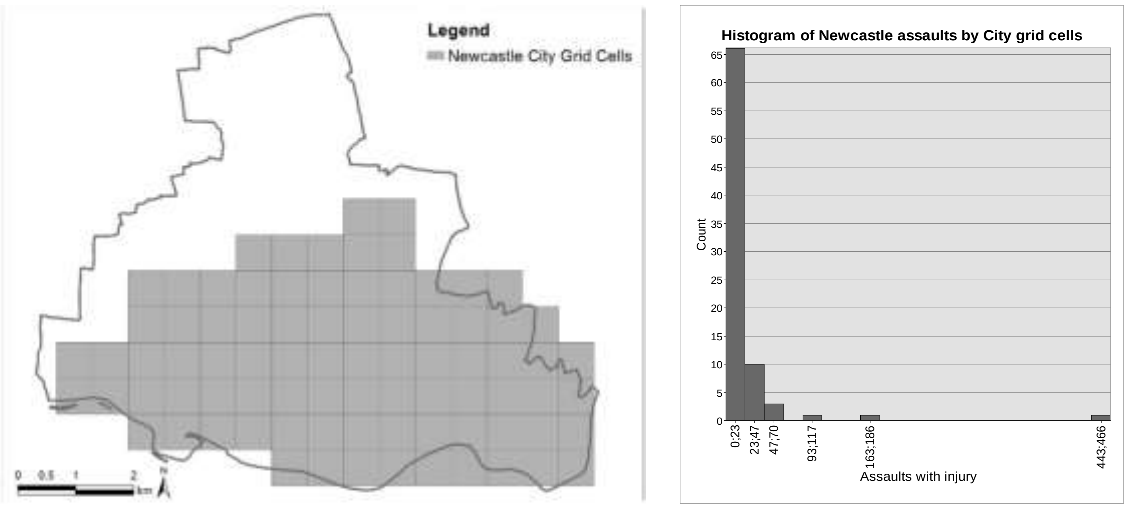

(a) (b)

Figure 11.4. (a) City grid and (b) histogram of assaults by City grid cells

Returning to the first of the two hypotheses for this chapter's empirical study, the research involved examining whether GWR could provide an effective means of determining at the local level the reasons why hotspots exist, and whether the explanatory variables that were identified to be significantly related to crime, varied between hotspots. The results described in this section have shown that local level analysis of spatially varying relationships at the spatial scale at which hotspots are identified is problematic for three main reasons. The first problem relates to a mismatch in the spatial resolution of crime data and the spatial resolution of data against which relationships are explored: data that are used as explanatory variables are typically not available at the spatial scale for areas that are identified as crime hotspots, and instead requires crime data to be aggregated to larger spatial scales (e.g., LSOAs for which data on explanatory variables are available) for relationships to be explored. Secondly, where explanatory variables are available at the spatial scale for which hotspots are identified, the very feature of the clustering of crime data into hotspots creates the problem of many other areas containing zero counts, causing the crime data to be overdispersed. The issue of overdispersion then results in difficulties in the application of standard Poisson regression modelling. The third problem then relates to compromising the examination of spatially-varying relationships at the precise spatial level at which hotspots are identified, due to the requirement to aggregate crime data to large geographic units: the tendency of crime to spatially cluster and the requirement to minimise the number of geographic units that contain zero counts results in the need for crime data to be aggregated to large geographic units for the purposes of performing spatial regression analysis. This requirement to aggregate crime data to large geographic units then reduces the ability to examine at the local spatial scale at which hotspots are identified the conditions that help explain the presence of these hotspots.

11.4.2. GWR modelling of burglary dwelling

The analysis of burglary data (aggregated to LSOAs) followed an exploratory approach. This initially involved including all of the twenty-five explanatory variables listed in Table 11.1 in a single preliminary exploratory OLS regression model and a statistical diagnostic analysis process. The burglary dwelling rate was transformed using a natural log approach in order to orient the distribution of the dependent variable towards a normal distribution.

Table 11.5 lists the main findings. These show that only four variables were significantly correlated to burglary dwelling. These variables were the number of Serious Acquisitive Crime (SAC) offenders per 1000 households, deprivation of education and skills, the 50-59 year old age group, and residents born in the UK. These variables were selected for Model 1, and the OLS and diagnostic analysis process was re-run. The analysis processes would then determine, if together, these variables were significant, and calculate key characteristics of the model in relation to its performance, significance and bias. This analysis process was re-run seven times using different combinations of explanatory variables until the model passed the necessary tests for it to be suitable for GWR. Due to the large number of variables included in the preliminary model, it was expected that a high degree of multicollinearity would be present. This was the case, with many variables that were identified as not significant having VIF values greater than 7.5. For each model that then followed, the VIF for each explanatory variable was less than 7.5 indicating the issue of multicollinearity had been addressed by including only those variables that were statistically significant. The decision-making behind the choice of explanatory variables was as follows:

- Model 1: This model included only those variables (from all variables listed in Table 11.1) that were significantly related to burglary dwelling: SAC offenders per 1000 households, deprivation of education and skills, 50-59 year old age group and residents born in the UK. From the preliminary model, both SAC offenders and deprivation were positively associated with burglary dwelling, whereas 50-59 year olds and UK born residents were negatively associated with burglary dwelling. Following the recalibration of the model using just these four variables, all variables were significant except for deprivation of education and skills. SAC offenders remained positively associated with burglary dwelling, and the 50-59 year old age group and UK born residents remained negatively associated with burglary dwelling. The Jarque-Bera statistic was significant, suggesting the model was biased because the residuals were not normally distributed.
- Model 2: This model included only those variables that were significant in Model 1. This resulted in the combination of SAC offenders, the 50-59 year old age group and residents born in the UK explaining 29% of the dependent variable. The Jarque-Bera statistic, however, remained significant.
- Model 3: This model used the natural log transformations of the explanatory variables used in Model 2 to attempt to pass the Jarque-Bera test. The Jarque-Bera statistic remained significant.

- Model 4: As mentioned in the method section (section 11.3), mapping a model's residuals thematically can help identify explanatory variables that could contribute to the model. An analysis of the distribution of Model 2 residuals was conducted, followed by a thematic spatial analysis of each of the twenty-one explanatory variables that had been removed following the analysis of the preliminary model. The thematic spatial analysis of these twenty-one explanatory variables suggested that the proportion of the population that was Asian visually matched with the mapped distribution of the Model 2 residuals and was, therefore, considered to be a variable that could contribute to the three variables used in Model 2. The processing of this model (Model 4) resulted in SAC offenders, the 50-59 year old age group and residents born in UK continuing to be significantly correlated to burglary dwelling, but the Asian population was not significantly correlated to burglary dwelling.
- Model 5: Following the thematic spatial analysis of each explanatory variable that was included in Model 4, an analysis of the Model 4 residuals, and further thematic analysis of all of the other initial explanatory variables, the proportion of the population that were students were included in Model 5 alongside the three variables from Model 2. The student population was included because the distribution of this variable visually matched with the mapped distribution of the Model 2 residuals. The analysis of Model 5 revealed the student population to be significantly positively associated to burglary dwelling, but at the cost of the proportion of the 50-59 age group no longer being significant. Model 5 did however lead to an improvement in the model as shown by a small increase in the adjusted R^2 (from 0.29 in Model 2 to 0.30 in Model 5) and a small reduction in the AICc (from 372 in Model 2 to 369 in Model 5). However, the Jarque-Bera statistic remained significant.
- Model 6: This model used only those variables that were significant in Model 5: student population, SAC offenders and UK born population. All explanatory variables were significant, but the Jarque-Bera statistic remained significant suggesting the model was still biased.
- Model 7: Following analysis of the distribution of Model 6 residuals and a thematic spatial analysis of each of the twenty-five explanatory variables, the proportion of the population that was Asian was included in Model 7 alongside the three variables in Model 6. This resulted in both the Asian population and the UK born population not being significantly correlated to burglary dwelling. The Jarque-Bera statistic remained significant.

- Model 8: The results of Model 7 showed the Asian population was only marginally not significant (p=0.07) compared to the UK born population (p=0.18). The Asian population was, therefore, included alongside the SAC offenders and the student population as the explanatory variables in Model 8. This resulted in all three being significant and positively associated to burglary dwelling, an adjusted R^2 that matched the highest from all previous models, the lowest AICc compared to previous models, and a non-significant Jarque-Bera statistic. A Moran's I test was performed on the model's residuals and these were determined not to be clustered.

Table 11.5. Main results from OLS models, identifying explanatory variables that were significant, model performance, significance and bias. LN = natural log transformation.

ID	Model variables	Coefficient (significance)	Adjusted R^2 (model significance)	AICc	Jarque-Bera significance (p=0.05)
1	• SAC offenders rate • Deprivation education/skills • 50-59 year old age group • Residents born in UK	• 0.14 (p=0.001) • -0.004 (p=0.19) • -7.05 (p=0.001) • -1.80 (p=0.01)	0.29 (p=0.001)	372	Significant
2	• SAC offenders rate • 50-59 year old age group • Residents born in UK	• 0.10 (p=0.001) • -6.69 (p=0.001) • -1.96 (p=0.001)	0.29 (p=0.001)	372	Significant
3	• SAC offenders rate (LN) • 50-59 year old age group (LN) • Log residents born in UK (LN)	• 0.26 (p=0.001) • -0.66 (p=0.001) • -1.41 (p=0.01)	0.26 (p=0.001)	379	Significant
4	• SAC offenders rate • 50-59 year old age group • Residents born in UK • Asian population	• 0.10 (p=0.001) • -6.76 (p=0.01) • -1.79 (p=0.05) • 0.002 (p=0.74)	0.28 (p=0.001)	372	Significant
5	• SAC offenders rate • 50-59 year old age group • Residents born in UK • Student population	• 0.12 (p=0.001) • -3.73 (p=0.17) • -1.64 (p=0.01) • 0.01 (p=0.05)	0.30 (p=0.001)	369	Significant
6	• SAC offenders rate • Residents born in UK • Student population	• 0.13 (p=0.001) • -1.87 (p=0.001) • 0.02 (p=0.001)	0.30 (p=0.001)	370	Significant
7	• SAC offenders rate • Residents born in UK • Student population • Asian population	• 0.14 (p=0.001) • -0.93 (p=0.19) • 0.02 (p=0.001) • 0.01 (p=0.08)	0.30 (p=0.001)	370	Significant
8	• SAC offenders rate • Student population • Asian population	• 0.16 (p=0.001) • 0.02 (p=0.001) • 0.01 (p=0.001)	0.30 (p=0.001)	369	Not significant

A number of additional exploratory models were calibrated in an attempt to improve Model 8. This involved adding the following variables in turn to the three explanatory variables from Model 8: 30-39 age group population, 40-49 age group population, 50-59 age group population, 60-69 age group population, 70 years and older age group population, population that arrived and became resident in 2007-2009, population that arrived and became resident in 2004-2006, social rented housing, owned housing, living environment deprivation, income deprivation, and the deprivation of education and skills. None of these other combinations improved on Model 8.

Model 8, consisting of the student population, the Asian population and SAC offenders offered the best combination of explanatory variables to produce a model with the highest level of performance, was significant, and was unbiased. This model resulted in an adjusted R^2 of 0.3 and an AICc of 369. Figure 11.5 shows the distribution of each of these explanatory variables across Newcastle. The maps show the student population to concentrate around the city centre, the Asian population to be greatest to the west of the city centre, and the SAC offenders' rate to be highest across the southern border of the district and in the central district area. Figure 11.5d shows the burglary rate to be highest in the southern and central regions of the district of Newcastle. Each of the three variables were found to be positively correlated with the log of the burglary dwelling rate, with the following coefficient values: SAC offenders rate 0.16; proportion of the population that were students 0.02; and the proportion of the population that were Asian 0.01. The modelling applied to burglary dwelling was an exploratory approach. This approach directs the analyst to identify variables that result in a good global model (such as Model 8), and a model that can then be applied using GWR. Following the application of a GWR model using the exploratory approach, the onus is then on the analyst to interpret the results based on plausible logic or some theoretical basis. The interpretation of the relationship between SAC offenders, the student population and the Asian population follows the GWR analysis of these variables.

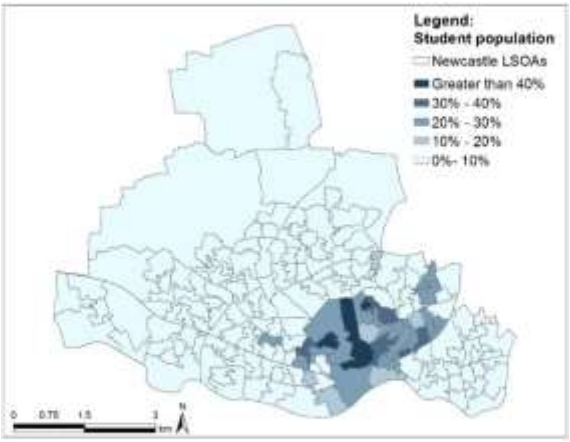

(a)

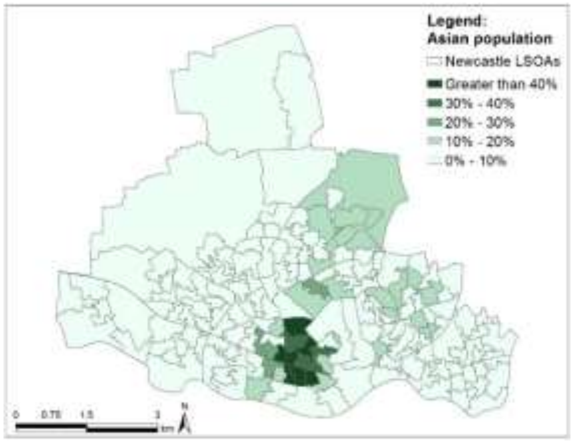

(b)

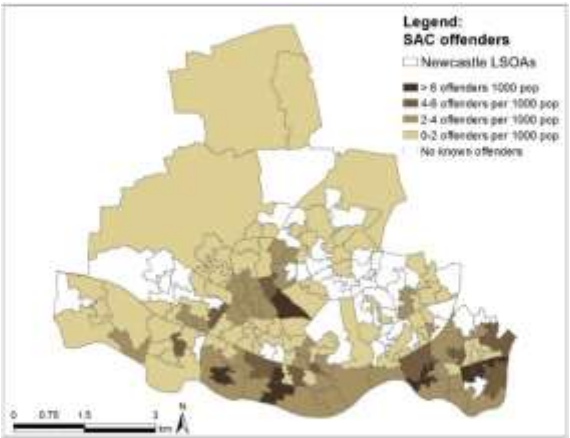

(c)

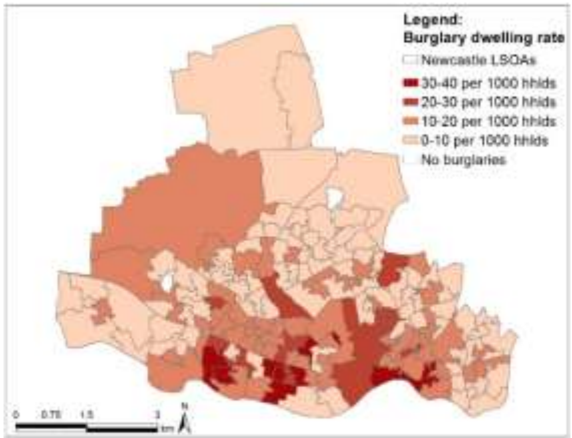

(d)

Figure 11.5. The distribution of (a) the student population, (b) the Asian population, (c) SAC offenders and (d) the burglary rate across the district of Newcastle

The three variables from Model 8 (SAC offenders, the student population and the Asian population) were used in two GWR models: one that applied a fixed bandwidth selection process and a second that applied an adaptive bandwidth selection process. The optimal fixed bandwidth was calculated to be 2421 m and 44 neighbouring LSOAs were determined as optimal for the adaptive bandwidth approach.

Table 11.6. Adjusted R^2 and AICc results comparing the OLS model for burglary dwelling against the Asian population, the student population, and SAC offenders, in relation to fixed and adaptive bandwidth GWR models

OLS adjusted R^2	**OLS AICc**	**GWR Fixed adjusted R^2**	**GWR Fixed AICc**	**GWR Adaptive adjusted R^2**	**GWR Adaptive AICc**
0.30	369	0.39	353	0.42	348

Table 11.6 lists the adjusted R^2 and AICc values for the OLS model and the two GWR models. The table shows that both GWR models improved upon the OLS model, and that the adaptive GWR model performed the best of all three – with an adjusted R^2 of 0.42 and an AICc of 348. These results show that rather than there being a stationary global relationship between burglary and the three explanatory variables (SAC offenders, student population and the Asian population), the relationship with these variables varied spatially, as indicated by the model improvements using GWR. Using an adaptive bandwidth, the student population, Asian population and the SAC offenders' rate variables were able to explain 42% of the distribution of burglary dwelling, with an improved AICc of 348, compared to the OLS model results of an adjusted R^2 of 0.30 and an AICc of 369.

The GWR results from the adaptive bandwidth model were used for further analysis. A number of outputs were provided from the GWR adaptive model for each geographic unit:

- Condition number: if this value is greater than 30 for any geographic unit, this suggests the presence of strong collinearity, and the results can be unreliable. None of the LSOAs had a condition number greater than 30. The highest was 8.3.
- Local R^2: these values describe how well the model fits with the observed values. The values range from zero to one, and are reflective of the model's global R^2 value. Low values indicate where the local model does not perform well, whereas higher values show where it performs best. Figure 11.6a shows the mapped Local R^2 values for the GWR adaptive model. This Local R^2 map shows that in the southern and south eastern parts of Newcastle (the more urban parts of the district) the model performed best, whereas in the more suburban and rural areas of the district, and in Newcastle city centre, the model did not perform as well. These results suggest that certain conditions relating to burglary dwelling and its relationship with the student population, the Asian population and the distribution of SAC offenders differed between urban residential areas, rural parts of the district of Newcastle, and Newcastle city centre.
- Standardised residuals: the standardised residuals have a mean of zero and a standard deviation of 1. Thematically mapping the distribution of the standardised residuals shows where the predicted value for the dependent variable (based on the relationship with the explanatory variables) has been over-estimated or under-estimated, relative

to the standard deviation of the residuals. Only nine of the 173 LSOAs had residuals greater than 2 standard deviations of the residuals' mean, with all these being less than 2.5 standard deviations from this mean (see Figure 11.6b).

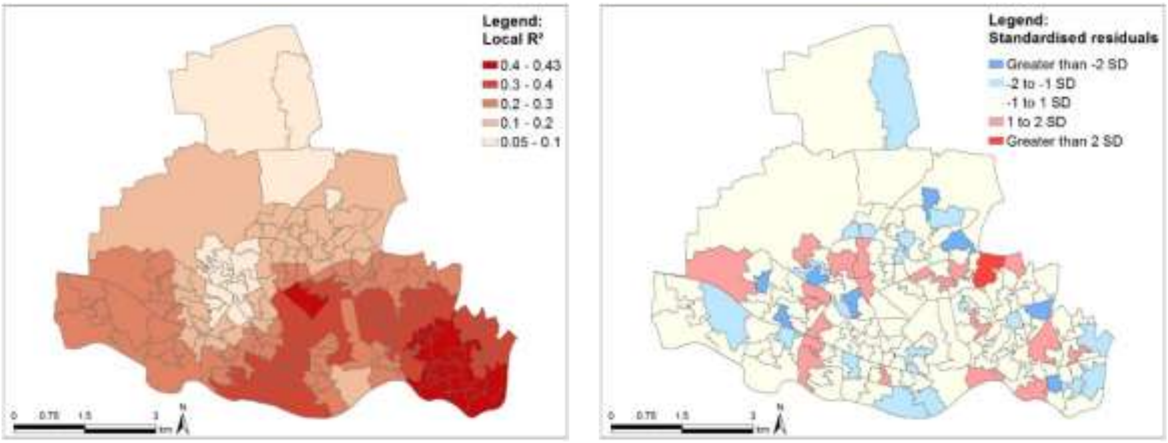

(a) (b)

Figure 11.6. (a) Local R^2 and (b) standardised residuals maps of the GWR model for burglary dwelling (dependent variable) and the student population, the Asian population, and SAC offenders

Figure 11.7 shows the mapped coefficient values and mapped standard errors of each of the three explanatory variables. Each coefficient illustrates the spatially varying relationship between the variable and burglary dwelling. For example, Figure 11.7a shows there was a stronger positive relationship between burglary and the student population in the central and eastern regions of the district, and that in the western region this relationship was negative. Similarly, the maps showing the relationship between the Asian population and burglary dwelling (Figure 11.7c), and the SAC offender rate and burglary dwelling (Figure 11.7e), show the relationship between these variables varied spatially. These results illustrate the value of analysing and determining if the relationships between the dependent variable and explanatory variables vary across space, rather than assuming the relationship to be stationary. The maps of standard errors (Figures 11.7b, d and f) show the areas where there was a greater level of reliability in the results for each explanatory variable – where these errors were small, more confidence can be placed in the relationship between the dependent variable (burglary dwelling) and each of the explanatory variables.

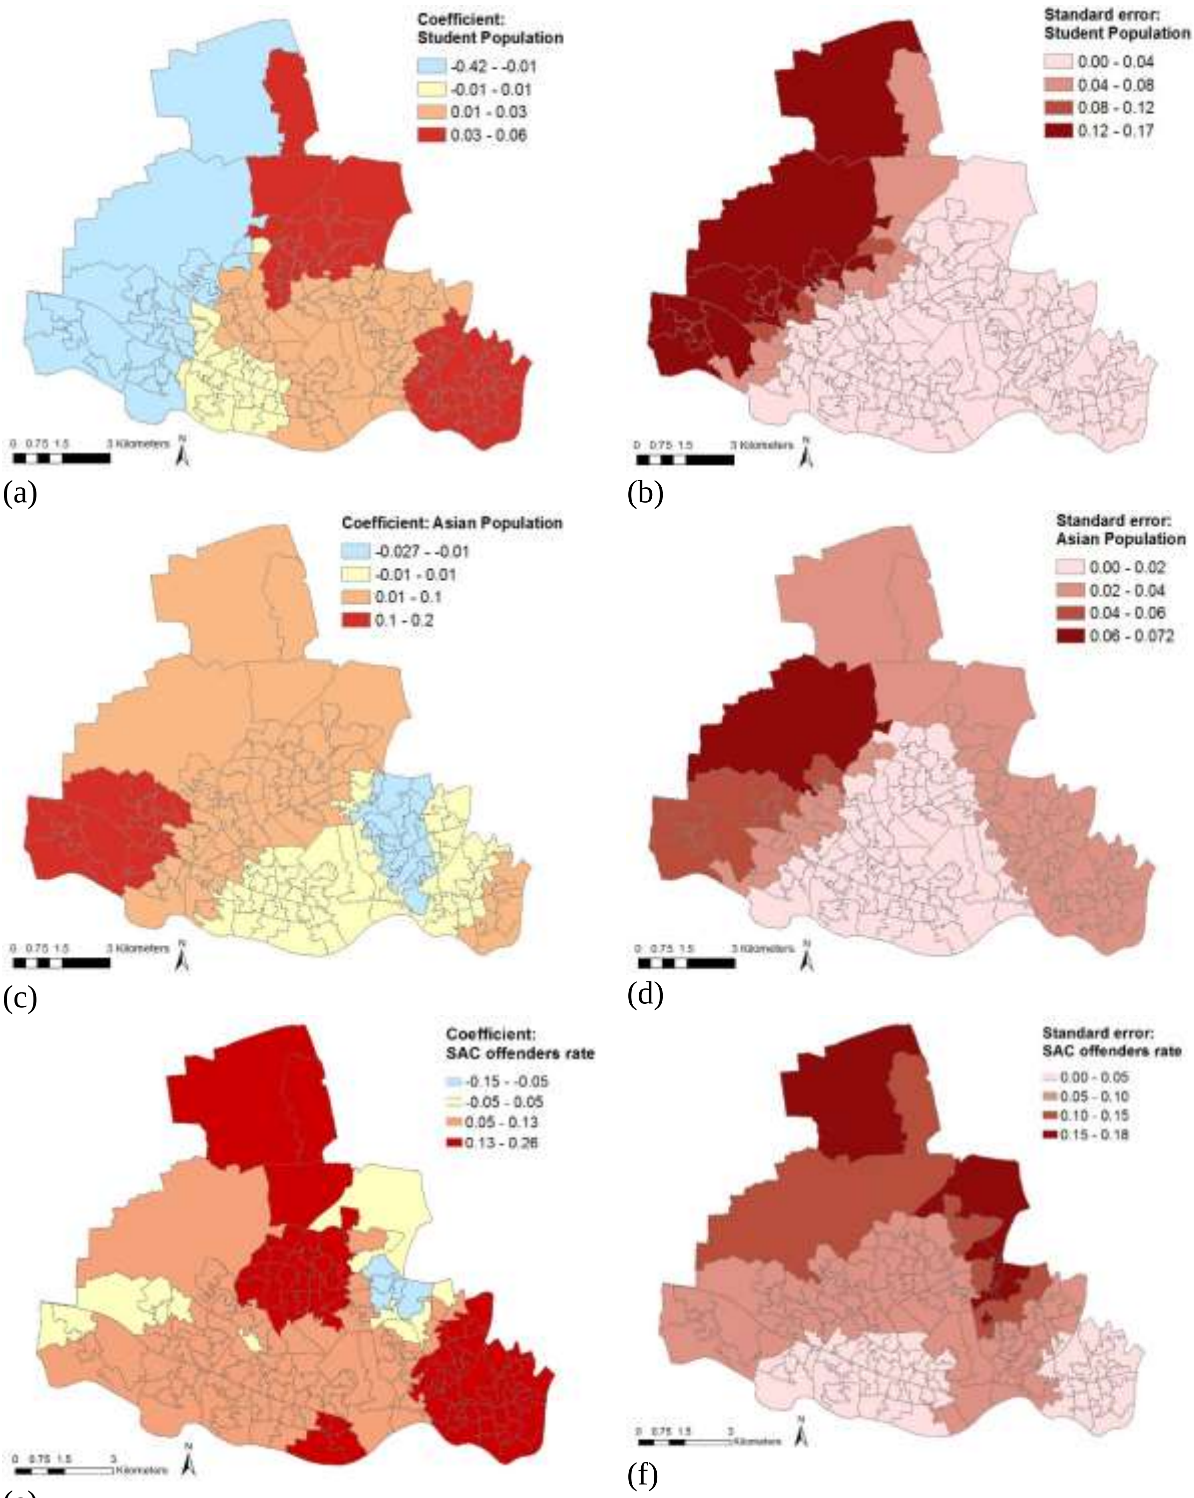

Figure 11.7. GWR model outputs of the student population, the Asian population and the SAC offender rate in relation to burglary dwelling: (a) coefficient and (b) standard error of student population; (c) coefficient and (d) standard error of Asian population; (e) coefficient and (f) standard error of SAC offender rate.

The results of the OLS regression analyses and GWR modelling for burglary dwelling illustrate the value in examining whether relationships between dependent and explanatory variables vary across space. In this example, an exploratory approach identified three variables that were significantly correlated with the distribution of burglary dwelling, and that the relationships between these explanatory variables and burglary dwelling varied spatially. The exploratory approach to a GWR analysis then requires the analyst to attempt to interpret the results based on plausible logic or some theoretical basis.

The positive relationship between SAC offenders and the burglary dwelling rate is supported by research into offender journey to crime patterns – offenders who commit burglary are likely to not travel far and live locally to the locations where they commit these offences (Rossmo, 2000, contains an extensive review of the journey to crime literature). The positive relationship between the distribution of SAC offenders and the burglary rate did though vary, with the relationship being stronger in some areas than it was in others. In one area, the relationship between SAC offenders and burglary levels was found to be negative (in the central eastern part of the district). This could be as a result of local offenders that reside in the district of Newcastle not offending in this area, or could be associated with offenders that live close to this area but over the district border in North Tyneside being responsible for the high levels of burglary in this area. Or, the explanation for the high level of burglary in this area could be for reasons not identified from this analysis (recall the best performing model explained 42% of the distribution of burglary dwelling) and remains unexplained. The positive relationship between the student population and Asian population can most likely be explained by the higher level of vulnerability to burglary these two demographic groups typically endure. For example, with students, multiple-occupancy living and the ownership of many electronic items, combined with a general relaxed attitude towards personal security presents them as attractive targets to offenders (McCreith and Parkinson, 2004). Similarly, Asian households tend to be attractive targets to offenders because of the high level of gold ownership associated with this demographic group, and the likelihood that this gold is stored within the house with limited in-built security to protect these valuable possessions (Gray, 2000). These interpretations could then inform policing and community safety initiatives that help reduce the positive and significant relationships between SAC offenders, the student population and the Asian population, with burglary dwelling in Newcastle.

While the spatial scale of analysis did not permit an analysis that could determine differences in the causes of the different hotspots, the GWR analysis has identified variables that appear to have a spatially varying influence on burglary dwelling levels across Newcastle. These findings suggest there may be potential in using results from a GWR analysis to inform the direction of crime prevention policy. For example, the results from the GWR analysis for burglary and the student population indicates that any measures designed to help tackle the high levels of burglary victimisation experienced by students could be best targeted to those areas where the relationship between the student population and burglary is strongly positive. In turn, planned strategic activity to reduce vulnerability to burglary amongst this highly victimised group could be measured in terms of the predicted crime reduction that could come from a targeted strategic activity. These strategic predictions of crime would be based on measuring the value an initiative may have in changing the explanatory variables (or the relationship with burglary) in some way. The possibility of GWR informing spatial predictions of crime is discussed further in section 11.4.4 following an analysis of assault with injury offences using GWR.

11.4.3. GWR modelling of assault with injury

The use of GWR for modelling assault with injury offences followed a hypothesis testing approach. The hypothesis testing approach involved selecting explanatory variables that were determined to theoretically explain the variation in assault levels. In the first instance, only licensed premises data were selected as the single explanatory variable. This was because of the theoretically determined strong relationship between the presence of licensed premises and violent assaults. In addition, licensed premises data were one of the few explanatory variables that were available at the point level and thus allowing for a grid based GWR analysis rather than one that was constrained to geographic administrative units such as super output areas.

In the first instance the standalone GWR software was used to conduct a GWR Poisson (GWPR) analysis of the 1000 m x 1000 m grid cells covering the full Newcastle study area, using a fixed bandwidth approach and an adaptive bandwidth approach. The GWR software simultaneously conducts an OLS analysis, albeit with a less comprehensive range of statistical diagnostic tests compared to the ArcGIS OLS modelling function. For a Poisson regression analysis, the proportion of the deviance explained rather than a R^2 value is generated. Table 11.7 lists the OLS and GWPR modelling results. The application of the GWPR models significantly improved the proportion of the deviance explained, from 0.55 using an OLS regression to 0.86 and 0.87 using GWPR fixed bandwidth and GWPR adaptive bandwidth modelling respectively. In both models using GWPR, the AICc also significantly improved, from 3404 for the OLS model to 1097 (GWPR fixed bandwidth) and 1028 (GWPR adaptive bandwidth). These results showed the adaptive GWPR model performed the best. What was also of note was the large proportion of deviance (0.87) that was explained from the GWPR adaptive bandwidth model using the single explanatory variable of licensed premises.

Table 11.7. OLS, GWPR fixed bandwidth and GWPR adaptive bandwidth model results of assault with injury (dependent variable) and licensed premises (explanatory variable) in Newcastle

	Proportion of deviance explained	**AICc**	**Bandwidth**	**Summary statistics for local coefficients**	
				Mean	**SD**
OLS analysis	0.55	3404	-	-	-
GWPR fixed	0.86	1097	1803 m	0.073	0.044
GWPR adaptive	0.87	1028	44 grid cells	0.261	0.263

Figure 11.8 shows the variation in the spatial distribution of the explained local deviance, the estimated local coefficient values, the standard errors, and t-values for the relationship between licensed premises and the distribution of assault with injury offences. These maps show that the strength and significance in the relationship between licenced premises and assault with injury was greatest around the city centre of Newcastle. The t-value results also show that the relationship between licensed premises and assaults was not significant in the north part of the district and towards the west. However, in both GWPR models the large number of cells containing zero values resulted in a problem of overdispersion. The presence of overdispersion meant that the standard errors between assaults and licensed premises could be underestimated, meaning that licensed premises may appear to be a significant predictor of assaults, when in fact they are not.

The most practical solution for helping to address the problem of overdispersion was to remove the majority of cells that contained zero values. This involved applying GWPR using the city grid lattice of 82 cells covering the southern areas of the district. Table 11.8 lists the results following the application of an OLS, fixed bandwidth and adaptive bandwidth GWPR analyses to the city grid cells. Removing the majority of grids that contained a zero count of assault with injury immediately resulted in an improvement in the proportion of the deviance explained under an OLS regression analysis from 0.55 to 0.71. In addition, the removal of these cells and a city-wide focus for GWPR analysis resulted in neither model showing evidence of overdispersion. Both GWPR models (using fixed and adaptive bandwidths) improved on the OLS model, with the fixed model performing best (AICc 696 and 0.83 of the deviance explained). Figure 11.9 shows the variation in the distribution of the explained local deviance, the estimated local coefficient values, the standard errors, and t-values between licensed premises and the distribution of assault with injury offences. These maps show similar results to those for the full district grid cell coverage shown in Figure 11.8, albeit results that are more reliable: the strength and significance in the relationship between licensed premises and assault with injury was greatest around the city centre of Newcastle.

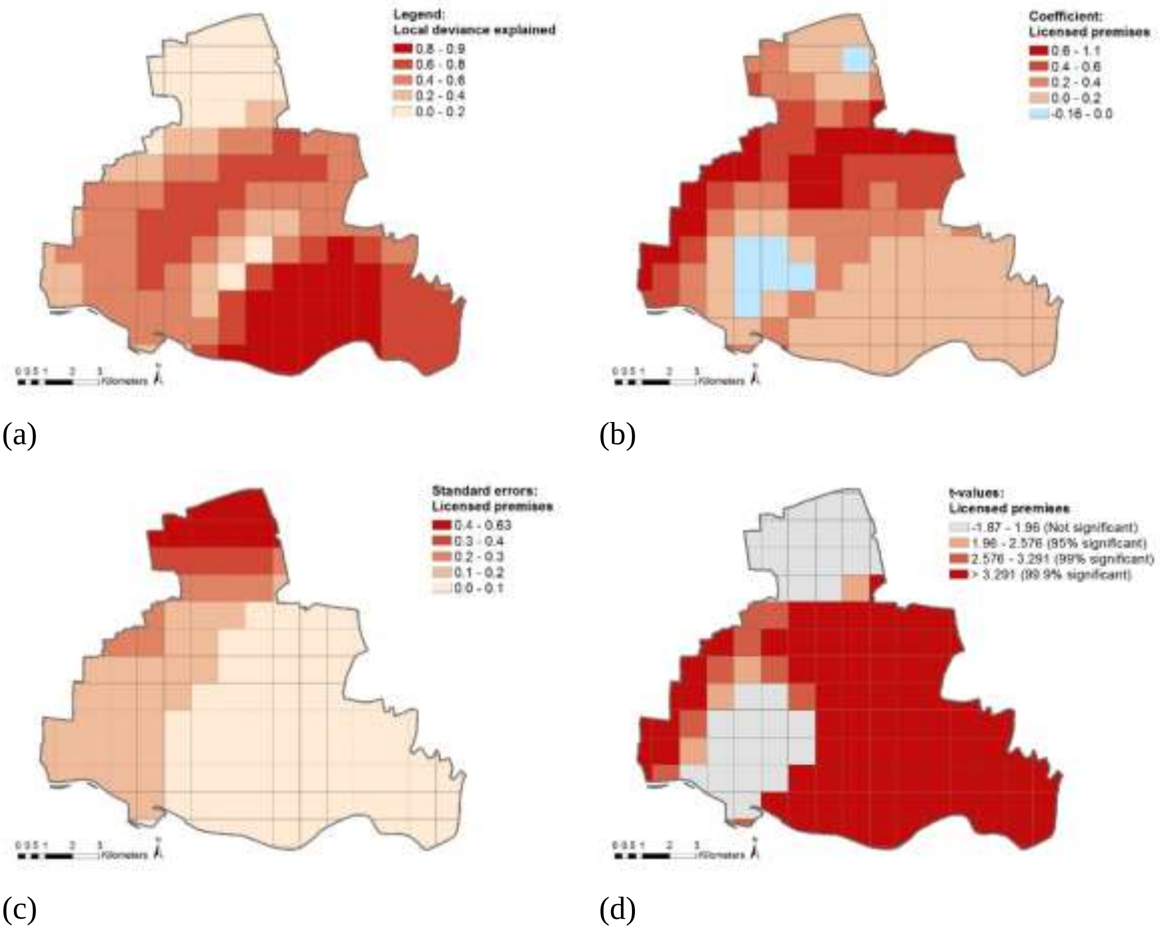

Figure 11.8. Adaptive bandwidth GWPR results showing (a) the local deviance in assaults explained by licensed premises, (b) the variation in the coefficients, (c) the standard errors, and (d) local t-values between licensed premises and assaults with injury

Table 11.8. OLS, GWPR fixed bandwidth and GWPR adaptive bandwidth model results of assault with injury and licensed premises for the city grids Newcastle study area

	Proportion of deviance explained	AICc	Bandwidth	Summary statistics for local coefficients	
				Mean	SD
OLS analysis	0.71	1174	-	-	-
GWPR fixed	0.83	696	1803 m	0.031	0.048
GWPR adaptive	0.81	753	45 grid cells	0.041	0.087

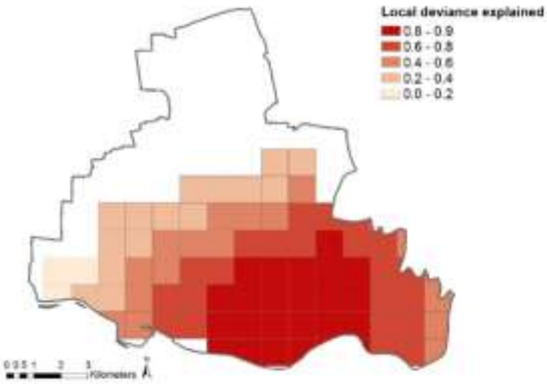

(a)

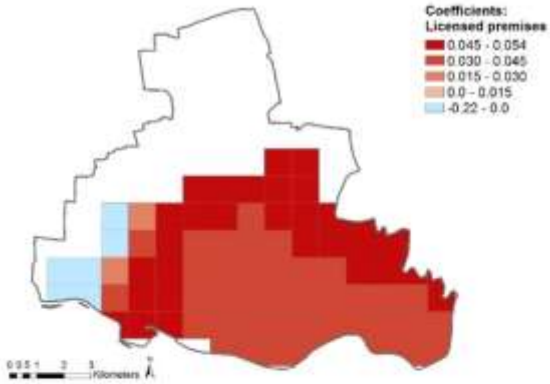

(b)

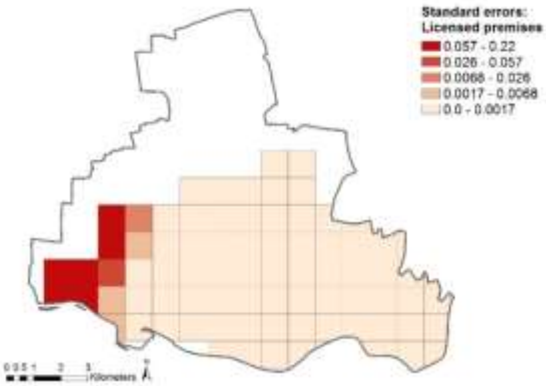

(c)

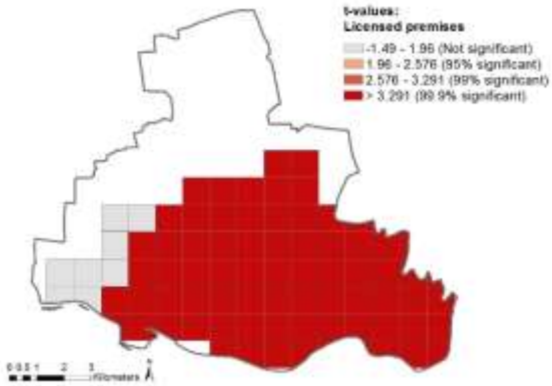

(d)

Figure 11.9. Fixed bandwidth GWPR results of the city grids study area showing (a) the local deviance in assaults explained by licensed premises, (b) the variation in the coefficients, (c) the standard errors, and (d) local t-values between licensed premises and assault with injury offences

The results using grid cells show that the relationship between assaults and licensed premises was significant, that the distribution of licensed premises explained a significant proportion of the deviance in assaults, and that this relationship spatially varied.

However, even using grid cells of 1000 m, the spatial scale of analysis was not fine enough to determine if there were differences in the relationship between assaults and licensed premises within and immediately around the single assault with injury hotspot that was identified using Gi* (recall that the Gi* assault with injury hotspot was only marginally greater in size than a single 1000 m grid cell). This inability to determine if there were differences in the relationship between assaults and licensed premises within and immediately around the single assault with injury hotspot was due to two main reasons. Firstly, aggregating assault and licensed premises data to the spatial scale of 1000 m grid cells was required in order to avoid the problem of a large number of geographic units containing zero counts and in turn causing a problem of overdispersion. Secondly, the need to use large bandwidths (that were determined as optimal) in the GWPR modelling process for examining the relationship between assaults and licensed premises (the size of which were also a reflection of the size of the geographic unit of study). For example, for the city grids fixed GWPR model, a bandwidth of 1803 m was determined as optimal – twice the size of the Newcastle city centre assault with injury hotspot that was identified using Gi*.

The analysis was re-run, this time using data for LSOAs and applying a Gaussian GWR approach on log transformed variables of assault with injury and licensed premises. The purpose of re-running the analysis was to determine if there were differences in the results when compared to the GWPR analysis, and identify if other variables (available only for census districts), rather than just the single licensed premises variable, improved the model. A Gaussian GWR analysis using ArcGIS also allowed for a more comprehensive examination of OLS statistical diagnostic tests (available in the ArcGIS application of GWR) and for a more interactive mapping and modelling iteration process.

Table 11.9 lists the OLS regression results from seven models that analysed the relationship between assaults, licensed premises and other explanatory variables for LSOAs in Newcastle. These other variables were chosen due to a hypothesised reasoning for their inclusion. The first model (Model 1) used only licensed premises as an explanatory variable and resulted in an unbiased model (i.e., the Jarque-Bera statistic was not significant), but an adjusted R^2 of 0.16, much lower than the explained deviance of 0.55 from the full 1000 m grid coverage OLS model of Newcastle. This suggests that changing and increasing the spatial scale of the unit of analysis from 1000 m grids to LSOAs may have resulted in the loss of some of the detail in the relationship between assaults and licensed premises. The second model included the student population, income deprivation and living conditions deprivation, alongside licensed premises data, as explanatory variables on the basis of hypothesising that each of these three new variables contributed to patterns of assaults across Newcastle – students, through their active engagement with the night-time economy, income deprivation because a number of studies have identified a link between this variable and assaults (e.g., Ching-Chi Hsieh and Pugh, 1993), and living conditions due to research findings linking this variable to assaults (e.g., Roncek, 1981). Of these variables, only income deprivation was found to be statistically significant (and positively associated to assaults). The third model used log transformations of each of these variables (licensed premises, student population, income deprivation and living conditions deprivation), but again, only licensed premises and income deprivation were identified as significant explanatory variables (and both positively associated to assaults).

The fourth model involved the inclusion of only licensed premises (log transformed) and income deprivation. Both variables were statistically significant and resulted in a large improvement in model performance (compared to Model 1) with an adjusted R^2 of 0.46 and an AICc of 390 (from Model 1 the adjusted R^2 of 0.16 and the AICc of 466). However, the Jarque-Bera statistic was significant indicating bias in the model. A log transformation of the income deprivation variable also resulted in the same conclusions (Model 5), and no improvement in model performance over Model 4. At this point it was decided to use a sub-set of only those licensed premises associated with Newcastle's night-time economy from the full licensed premises dataset. The selected premises would include only bars, pubs and nightclubs rather than off-licenses because it was hypothesised that this sub-set was more likely to be associated with assaults. Using this subset of night-time economy licensed premises and income deprivation (Model 6) resulted in an improved and unbiased model (adjusted R^2 of 0.52, AICc of 373, and a non-significant Jarque-Bera statistic). The VIF for the two variables was 1.0, indicating that both variables contributed to the model. The Koenker statistic for Model 6 was statistically significant ($p=0.01$), suggesting the relationship between the dependent and explanatory variables varied across space, and that the model was a good candidate for GWR. A re-run of Model 6 using a log transformation of income deprivation (Model 7) did not result in any model performance improvement. Therefore, at the LSOA level, the relationship between assault with injury and licensed premises was improved by including only those licensed premises associated with the night-time economy and through the inclusion of income deprivation. An analysis of the residuals from Model 6 determined that these were randomly distributed (Moran's I z-score = 0.81). A GWR analysis was then performed using night-time economy licensed premises and income deprivation as the two explanatory variables for the model.

Table 11.9. OLS regression summary statistics for models examining the relationship between explanatory variables and assaults with injury in Newcastle at the LSOA level. LN = natural log transformed.

ID	Model variables	Coefficient (significance)	Adjusted R^2 (model significance)	AICc	Jarque-Bera significance (p=0.05)
1	• Licensed premises (LN)	• 0.434 (p=0.001)	0.16 (p=0.001)	466	Not significant
2	• Licensed premises (LN) • Student population • Income deprivation • Living conditions deprivation	• 0.36 (p=0.01) • -0.003 (ns) • 3.58 (p=0.001) • 0.017 (ns)	0.47 (p=0.001)	390	Significant
3	• Licensed premises (LN) • Student population (LN) • Income deprivation (LN) • Living conditions deprivation (LN)	• 0.39 (p=0.01) • 0.054 (ns) • 0.62 (p=0.001) • -0.013 (ns)	0.44 (p=0.001)	398	Significant
4	• Licensed premises (LN) • Income deprivation	• 0.43 (p=0.001) • 3.75 (p=0.001)	0.46 (p=0.001)	390	Significant
5	• Licensed premises (LN) • Income deprivation (LN)	• 0.41 (p=0.001) • 0.59 (p=0.001)	0.45 (p=0.001)	395	Significant
6	• NTE licensed premises (LN) • Income deprivation	• 0.81 (p=0.001) • 3.80 (p=0.001)	0.52 (p=0.001)	373	Not significant
7	• NTE licensed premises (LN) • Income deprivation (LN)	• 0.80 (p=0.001) • 0.61 (p=0.001)	0.51 (p=0.001)	376	Not significant

Table 11.10 lists the GWR model performance results (using both a fixed and adaptive bandwidth) of assault with injury offences in relation to night-time economy licensed premises and income deprivation. A bandwidth of 3242 m was calculated as optimal for the fixed GWR model, and a bandwidth of 55 neighbours was calculated as optimal for the adaptive GWR model. The results show that both GWR models improved upon the performance of the OLS model, and that while both GWR models generated the same AICc value, the adaptive GWR model performed marginally better with an adjusted R^2 of 0.58.

Table 11.10. Adjusted R^2 and AICc results comparing the OLS model for assault with injury offences against night-time economy licensed premises and income deprivation, in relation to fixed and adaptive bandwidth GWR models

OLS adjusted R^2	**OLS AICc**	**GWR Fixed adjusted R^2**	**GWR Fixed AICc**	**GWR Adaptive adjusted R^2**	**GWR Adaptive AICc**
0.52	373	0.55	364	0.58	364

Figure 11.10 shows outputs from the adaptive GWR model. Figure 11.10a shows the relationship between assaults, night-time economy licensed premises and income deprivation varied spatially, with this relationship being strongest in around the city centre and in the Walker and Byker neighbourhoods to the east of the city centre. The standardised residuals map (Figure 11.10b) shows the model performed well in most areas, with the lowest levels of performance in the suburban areas around the city centre area and in some outlying rural areas of the district. The coefficient map of night-time economy licensed premises (Figure 11.19c) shows how the relationship with assaults varied across Newcastle, with this relationship being highly positive and most significant (Figure 11.19e map of t-values) in and around the city centre area, but negative (and not significant) in the west and northern-most region of the district. The supporting standard error map of the night-time economy licensed premises coefficient values shows the results were most reliable in the city centre area. Figure 11.10f shows the relationship between assaults and income deprivation also varied spatially, with all areas exhibiting a positive and significant relationship (Figure 11.10h) between these two variables, with the relationship being strongest in the Lemington and Gosforth neighbourhoods of Newcastle.

The results of the OLS regression analyses and GWR modelling for assault with injury offences illustrate the value in exploring whether relationships between dependent and explanatory variables vary across space. In this example, a hypothesis testing approach was applied, using an iterative process to include explanatory variables that were considered to be related to violent assaults, based on sound theory and empirical evidence. While the spatial scale of analysis did not permit an analysis that could determine differences within and around the single assault hotspot in Newcastle city centre, the GWR analysis did identify variables that appeared to have a spatially varying influence on assault levels across the district of Newcastle.

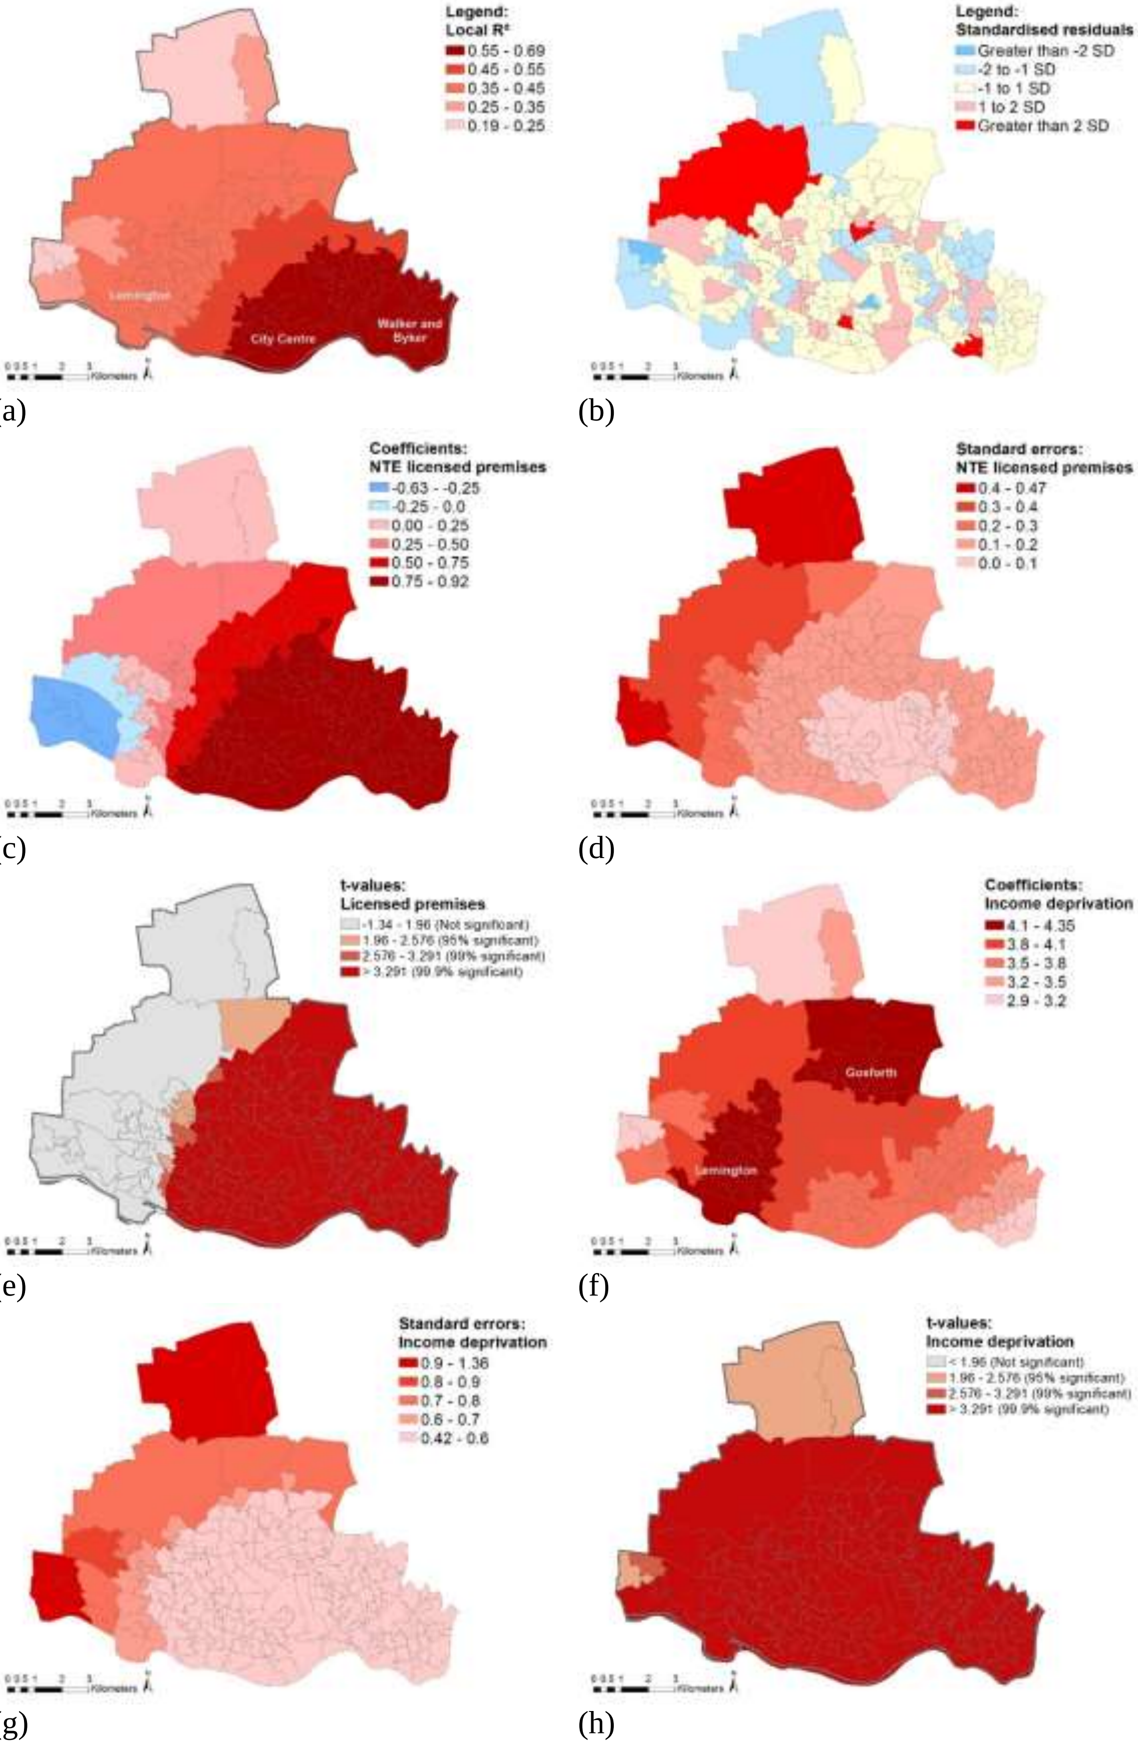

Figure 11.10. GWR model outputs of (a) Local R² and (b) standardised residuals for assault with injury (dependent variable) and night-time economy licensed premises and income deprivation, and (c) coefficient, (d) standard error and (e) local t-values of night-

time economy licensed premises; and (f) coefficient, (g) standard error and (h) local t-values of deprivation income

The analysis also employed two different types of GWR modelling: a GWR Poisson model and a GWR Gaussian model. Both models produced similar results, illustrating the strong and positive relationship between assaults and licensed premises, with the GWR Gaussian model that was applied to LSOAs having the benefit of testing whether other variables alongside licensed premises explained the spatial distribution of assaults. However, the compromise in using LSOAs rather than smaller grid cells came at the cost of losing some of the spatial detail in explaining the relationship between assaults and licensed premises. The GWPR modelling also presented the challenge of tackling overdispersion, which was solved by removing a large number of grid cells that contained zero counts and focusing instead on a sub set of the full study area which covered the main city conurbation of Newcastle. The results of the GWR modelling also further indicate (following the findings from the GWR modelling of burglary) the potential in using a GWR analysis process for informing the direction of crime prevention policy and strategic crime prediction. For example, if crime prevention activity is targeted to licensed premises in the areas where the GWR model has indicated the relationship with assaults is strongly positive, this may help reduce crime, and to a level that can be measured using the coefficient calculated for licensed premises. The potential of using GWR to inform spatial predictions of crime is discussed further in the next section.

11.4.4. Use of GWR outputs for directing policy for strategic crime prediction

The results from the two GWR case studies (burglary dwelling and assaults with injury in Newcastle) have illustrated how analysis into the relationships between variables that attempt to explain the distribution of crime can be conducted, and how these relationships vary across space. However, the findings have also shown that the spatial scale of analysis required for GWR is most likely to be greater than the geographical size of crime hotspots identified using techniques such as the Gi* statistic. Hence, it would appear from the research findings, using the case studies from Newcastle, that a GWR analytical process does not permit an analysis that determines explanatory differences between individual hotspots. In turn, this means that a more data-rich approach to hotspot analysis through the inclusion of variables alongside, or in replacement of retrospective crime data cannot be instructed from a GWR analysis.

The critiques of GWR in other fields of science (as previously discussed in Chapter 2 and section 11.3 of the current chapter) have also expressed some concerns over the extent of using GWR results for inference. However, there is a danger in statistical modelling of failing to account for the balance that is required between the technical specification of a technique and its practical application. The analysis in this chapter has illustrated that in practice, the purity of a statistical modelling process needs to be balanced by specifying the model to suit the data and the context in which these data are applied. Through the diligence of correctly specifying a model (without compromising the model's technical integrity), balanced with the constraints that data may place on the modelling process, the results from this research are illustrative of the valuable contribution that GWR modelling could offer to policing and public safety practice. The use of GWR for policing and public safety applications may not only include identifying variables that have a spatially varying relationship with crime, but also could be used to inform strategic predictions of crime. For example, as the results from one of the case studies showed evidence of a relationship between assaults, night-time economy licensed premises and income deprivation, these results could be used to help inform policy direction for long-term crime reduction strategies and offer some predictive indication of how crime may change as a result.

The manner in which GWR modelling could be used to inform strategic predictions of crime is provided in the following two examples. Both examples illustrate the potential of using GWR for strategic crime prediction based on the results from the best performing models from the analyses in sections 11.4.2 and 11.4.3. The coefficients determined from the GWR modelling results are designed to be illustrative of potential changes in crime rather than offering exact forecasts. Other events may have a long-term impact on crime (such as, the national introduction of a minimum price per unit of alcohol on violent assaults) or changes in crime may result from the strategic activity that was informed from a GWR modelling process (e.g., the geographical displacement of crime to other areas). Nevertheless, identifying variables that statistically correlate and spatially vary with patterns of crime provides practitioners with a potential opportunity for targeting strategic activity and predicting its impact.

- An interpretation of the GWR modelling of burglary dwelling (albeit with some caution to the exact extent of the statistical inference results) could be used to inform policy and a strategic reduction plan for burglary. The OLS and GWR modelling of burglary dwelling identified the student population (coefficient 0.02), the Asian population (coefficient 0.01) and the SAC offender rate (coefficient 0.16) as being significant explanatory variables. An adaptive GWR model of these variables resulted in an adjusted R^2 of 0.42, with Figures 11.6 and 11.7 showing how the relationship between burglary dwelling and these explanatory variables varied across Newcastle. While it is unlikely and unethical to expect the number of students or the Asian population to be changed in those areas where the relationship with burglary dwelling is strongest, the GWR results help identify where crime reduction programmes could be targeted to address the significantly observed relationship (and high level of victimisation) experienced by these two groups. Targeted activity towards offenders of burglary (and offenders of other serious acquisitive crimes), that would include their intensive supervision, disruption, reduction in repeat-offending, and their removal (e.g., through the serving of custodial sentences), could also result in burglary dwelling reductions, particularly in the neighbourhoods where this relationship was found to be strongest. For example, using the coefficient determined from the OLS model as an indication of the impact of this type of offender targeted activity, for every one unit reduction in the number of SAC offenders per 1000 households we could anticipate a reduction of 16% in the burglary rate (recall the burglary rate was log transformed for the OLS and GWR modelling). In the areas of Newcastle where this relationship was strongest, we could anticipate a reduction in the burglary rate of 26% for every one unit reduction in the SAC offenders' rate. While other factors may influence a change in offending behaviour, the GWR analysis is illustrative of the predicted impact a change in one variable may have on changes in crime.

- The best performing model that used LSOAs to examine the relationship between assaults with injury and a number of explanatory variables was Model 6 (see Table 11.9). Model 6 included night-time economy related licensed premises and income deprivation. Recall that the assaults and licensed premises data were log transformed while income deprivation data remained in its original indexed format. Income deprivation values for Newcastle ranged between 0.01 and 0.61, with higher values representing a higher level of income deprivation. Using the results from the OLS model of assaults aggregated to LSOAs, a strategic programme aimed at addressing income deprivation that resulted in a reduction in the income deprivation score of 0.1 could yield a reduction of assaults of the order of 38% (taken from the coefficient of 3.8 in Table 11.9). The income deprivation coefficient for this relationship with assaults was highest in the Lemington and Gosforth areas, with values on average of

4.2 (see Figure 11.10). These results suggest that a reduction in income deprivation in the Lemington and Gosforth areas could have a greater effect on reducing the number of violent assaults. Similarly, a 1% reduction in the number of licensed premises could yield a 0.8% reduction in assaults (taken from the coefficient of 0.81 – Table 11.9). There are just over 100 night-time economy related licensed premises in Newcastle, and in 2009/2010 there were 1841 recorded assaults. Therefore, for every licensed premise that closed, we could anticipate a reduction of approximately 15 assaults per year. However, this relationship with assaults was highest in Newcastle city centre, with coefficient values of 1.0 on average. This suggests that a reduction in the proportion of licensed premises in the city centre area could yield a larger reduction in assaults. Conversely, if the number of licensed premises in Newcastle was allowed to increase, we would expect an increase in violent assaults, with this increase in assaults effect being greatest if new licensed premises opened in the city centre.

Reducing the actual number of licensed premises could be problematic, so an alternative crime reduction solution that is often used instead is to implement stricter licensing conditions on pubs, bars and nightclubs in crime hotspots. A strategic policy that implemented stricter licensing conditions on pubs, bars and nightclubs in Newcastle city centre (such as, police and community safety officials working more closely with managers of these premises to reduce violent assaults) could prove fruitful in reducing the strong relationship between violent assaults and licensed premises in the city centre. A target for this strategic policy would be to reduce violent assaults in the city centre so this area did not stand out as being any more problematic than other areas of the district (e.g., reducing the average coefficient value for the city centre from 1.0 to some new, lower target). In practice, this could also involve a more detailed analysis of each licensed premise to identify which experienced (or were associated with) the highest volumes of assaults, with focused policing and community safety activity towards these. This more detailed analysis of licensed premises could also examine whether the type and size of venue, and the volume of customers during opening hours also had an influence on the levels of violent assaults.

11.5. Interpretation and conclusions from research study 7

The spatial scale at which hotspots are identified to assist the targeting of operational policing tactics and crime prevention initiatives is at a much finer level than spatial regression analysis using GWR allows. Therefore, (and in relation to hypotheses 7) GWR does not provide a means of determining why individual hotspots exist, and identifying if there are explanatory differences between these hotspots. In turn, this means that a GWR analysis is not able to inform a more data-rich approach to hotspot analysis that uses explanatory variables alongside, or in replacement of retrospective recorded crime data. However, at a broader scale of analysis, GWR can be used to identify those variables that explain crime levels, and importantly show how these explanatory relationships vary. This has included identifying those areas where crime levels are highest (but larger in area than the hotspots identified using Gi*) and the strength (and reliability) of the relationship between the distribution of crime and explanatory variables. While there remain some concerns over the extent to which GWR results can be used for determining inference, the results do at least provide some indication of the changes in crime that could be anticipated if action was targeted to address the conditions resulting in the explanatory variables being significantly correlated to crime. That is, (and in relation to hypothesis 8) the results of this type of GWR analysis do offer value in helping to inform crime prediction for strategic and policy forecasting purposes.

The findings from research study 7 also illustrate some of the technical spatial analysis challenges associated with handling spatial scale in spatial regression modelling. These challenges include recognising that different explanatory variables are likely to be measured at different levels of spatial scale (e.g., point level or census output area level), and that the scale at which the model attempts to represent the observable world (e.g., using grids of different sizes) and the relationships between variables (i.e., different bandwidth sizes) is likely to impact upon the modelling process and the results. The effect of all these issues associated with spatial scale are then compounded in a multivariate model. Furthermore, the desire to perform precise spatial examinations of the relationships between crime and explanatory variables using GWR appeared to be held back due to the large bandwidths that were determined as optimal for the modelling process. This suggests that some further examination of bandwidths for the GWR analysis of crime data would merit some attention. In addition to this, an examination of other spatial regression methods that permit a more local examination of spatially varying relationships would also be of benefit.

Correlation of course does not equal causation, therefore, there is the need to ensure that any modelled relationship, and how it spatially varies, can be interpreted on sound theoretical grounds. For example, the journey to crime research literature shows that offenders tend to travel short distances to commit crime, particularly crimes such as burglary. Therefore, a high rate of burglars in an area is likely to result in a high level of burglary in that same area. There is often also the need to examine any identified relationship further in order to accurately interpret what the patterns infer. For example, the presence of a high concentration of licensed premises in an area, may alone not necessarily fully explain the high levels of violent assaults in this area. Instead, the high levels of violent assaults could be due to a small number of particularly problematic premises being the main cause of the problem, amongst a high density of licensed premises. This, therefore, indicates the importance of additional analysis that examines and further explains correlations between explanatory variables and patterns of crime revealed from a GWR analysis before policy decisions are made.

In this research study, two approaches were used for modelling spatial relationships: an exploratory approach and a hypothesis testing approach. The exploratory approach (on burglary dwelling) began with twenty-five variables and through a series of iterations resulted in the selection of three variables for the model. The hypothesis testing approach (on assault with injury) began with one variable and resulted in a subset of this initial variable and one other variable for the model. While both approaches resulted in producing GWR models that showed how relationships spatially varied, a level of hypothesis testing was also required in the selection of variables for the iterations of the exploratory approach. The iterative process to the exploratory approach illustrated that rather than choosing variables at random to add to a model, practical efficiency directed the selection of other explanatory variables on the basis of some theoretical grounds. In addition, even when an exploratory approach identifies relationships between dependent and explanatory variables, these relationships need interpretation to ensure there is sound theoretical reasoning to explain the relationships. This interpretation is also important for helping to determine the types of policing and public safety activities that could then be designed to influence reductions in crime.

The research has also exposed the challenges in performing spatial regression analysis on crime data. This not only refers to the violation of assumptions in regression analysis that are presented when exploring spatial relationships, but the very nature of crime to cluster in space. The spatial patterning of crime into hotspots naturally results in producing statistical outliers (i.e., the hotspots) and areas where very little or no crime occurs. This then creates the challenge of ensuring that appropriate regression analysis techniques are used, and a comprehensive statistical diagnostic analysis of the dependent and explanatory variables is performed to ensure the results are not misinterpreted. This includes checking for overdispersion and the spatial clustering of residuals. The required rigour in appropriately specifying a model for analysis, from which the results can be trusted, illustrates the importance of correctly accommodating the unique nature of spatial data, and understanding the spatial qualities of the data that are to be applied to a spatial regression modelling process.

12. Discussion, implications, contribution to the field, potential new areas for research and conclusions

The research that has been conducted over seven empirical studies has examined differences in the prediction performance of hotspot mapping techniques and how other spatial analysis techniques can contribute alongside hotspot analysis in predicting where crime is likely to occur. The research has also identified the temporal stable nature of crime concentration, while further illustrating the value of using recent crime events for predicting where crime is likely to occur in the immediate future. Analysis of the conditions that are significantly related to the spatial distribution of crime has also provided some indication of how this can be used to predict the impact of strategic plans for crime reduction.

In this final chapter, the results across the seven research studies are brought together and discussed in more detail. This discussion begins by considering how the current research has answered the primary research question and tested each hypothesis that was used to frame the direction of each empirical research study. The discussion in this chapter then considers how the technical and methodological findings from the research could have an influence on analytical practice in policing and public safety, the implications of the findings on policing and public safety tactical and strategic response practice, and the policy implications of these findings. The implications of the research findings on environmental criminology theory are also discussed. The final sections of this chapter record the contributions to the field that this research offers, potential areas of new research and the primary conclusions from the research findings.

12.1. Summary of findings in relation to the primary research question and hypotheses

The primary question this PhD research aimed to answer was *to what extent can hotspot mapping be used to effectively predict where crime is likely to occur?* This PhD research has shown that extremely accurate predictions of where crime is likely to concentrate in the future can be determined by using good hotspot analysis of where crimes have concentrated in the past. The current research has also shown, however, the importance of qualifying what is meant by 'the future' (in a predictive sense), with the distinction being made between spatial crime predictions that relate to the immediate future, the near future and the more distant future. By considering the prediction of spatial patterns of crime in these three temporal terms of the future has allowed for a more considered assessment of how hotspot analysis and other spatial analysis techniques can provide accurate predictions of where crime is likely to occur. The need to consider predictions of the spatial patterns of crime for these three different temporal periods is discussed further in section 12.7 of this chapter.

Eight hypotheses were used to frame the direction of the seven empirical research studies that were completed. The results from each hypothesis that was tested are considered in the section that follows, and collectively are discussed further in this chapter in relation to how the research findings have implications on practice, policy and theory.

Hypothesis 1: Hotspots can be identified using retrospective data for a short period of time rather than requiring retrospective data for longer periods of time

The findings from the research that tested hypothesis 1 (chapter 3, research study 1) showed that hotspots of crime that are likely to exist in the future can be determined from relatively short retrospective periods of recorded crime data, albeit with differences in the retrospective period between study areas and crime types. The statistical presence of clustering in retrospective crime data can be identified using the Nearest Neighbour Index (NNI). Using the NNI, the research relating to hypothesis 1 showed that for the Camden/Islington study area, only one week of burglary dwelling, theft from the person data and theft from vehicle data for Camden/Islington was required for identifying the presence of hotspots. However, for the Newcastle study area, over one week of crime data was required for each crime type for hotspots to be evident. In the case of theft of vehicles in Newcastle, over 16 weeks of crime data was required for hotspots to be statistically evident. Analysis of the number of crime events that were required before hotspots were evident showed that this number ranged from 34 theft from the person offences in Newcastle to 66 theft from vehicle offences in Camden/Islington. Once hotspots were detected in the crime point data, all other input data for longer retrospective periods showed statistical evidence of clustering.

In practitioner terms, simply choosing a retrospective period, whether it be based on a retrospective number of days or retrospective volume of crime, and expecting hotspots to be present is not sufficient if the analyst then expects hotspots to appear on a map. The results from testing this first hypothesis illustrated the value in performing the NNI test as a preliminary stage to hotspot mapping to ensure hotspots are present in the data that are examined.

Hypothesis 2: Common hotspot mapping techniques differ on how accurately they predict spatial patterns of crime

To test hypothesis 2, a series of experiments were conducted (research study 2) that compared the spatial crime prediction performance of spatial ellipses, thematic mapping of geographic units, thematic mapping of grid cells and kernel density estimation. The results from testing hypothesis 2 showed that KDE consistently outperformed the other common hotspot mapping techniques in predicting spatial patterns of crime. The consistency in KDE outperforming the other techniques was not only across the two study areas and for different measurement dates, but also for the range of different crime types. In one example, KDE hotspot maps of theft from the person in Newcastle predicted where 68% of crimes of this type were predicted to occur, in an area representing just 3% of the study area of Newcastle.

Hypothesis 3: The technical parameters used in hotspot analysis techniques have an influence on the techniques' spatial crime prediction performance

As a result of testing hypothesis 2, KDE was identified as the commonly used hotspot analysis technique that consistently outperformed the other common hotspot techniques in predicting spatial patterns of crime. The technical parameters of KDE, namely the cell size and the bandwidth size, were then the subject of analysis relating to hypothesis 3 (research study 3) to determine if these two parameters influenced KDE hotspot mapping spatial crime prediction performance. The findings from testing hypothesis 3 showed that the bandwidth size rather than the cell size did have an influence on KDE spatial crime prediction performance. The range of cell sizes that were used for producing KDE hotspot maps were found to mainly have an impact on the resolution of the density surface, with smaller cells improving the visual appeal of the KDE hotspot map, rather than cell size having an impact on spatial crime prediction performance. The choice of bandwidth size did, though, have an impact on the prediction performance of KDE hotspot maps, with smaller bandwidths consistently producing the best prediction performance results.

Hypothesis 4: Spatial significance mapping methods provide an improved means of predicting where crime is likely to occur in comparison to common hotspot mapping techniques, and removes the ambiguity of defining areas that are hot

To test hypothesis 4, research study 4 compared KDE hotspot analysis to the Gi* statistic. The Gi* statistic can be used to identify areas where the spatial concentration of crime is statistically significant. As result of a number of experiments that compared hotspot mapping output generated using Gi* in comparison to mapping output generated using KDE, the use of the Gi* statistic was found to help remove much of the ambiguity in determining hotspot areas and consistently performed better than KDE for predicting where crime was likely to occur. The Bonferroni correction measure was also used with the Gi* statistic to further improve the determination of hotspot areas by helping to address the issues of multiple testing.

Hypothesis 5: Areas that are identified as hotspots of crime are places where the concentration of crime has been endured consistently for at least one year, and where the concentration of crime is likely to continue to persist into the future

The research results from testing hypothesis 5 (research study 5) found that hotspots identified using the Gi* statistic displayed high levels of temporal stability – both in terms of the longevity in where crime had previously been endured, and where the concentration of crime then continued to persist into the future. The research findings did, though, show some differences between crime types, with hotspots of burglary dwelling tending to vary most in their temporal stability, whereas hotspots of thefts from the person and hotspots of assault with injury were highly stable.

Hypothesis 6: Recent incidents of crime provide an effective means of accurately predicting the immediate future, but the accuracy in these predictions reduces for longer periods of the future

The results from testing hypothesis 6 (research study 6) showed that the prospective mapping approach, based on using only data from the recent past and predicting where crime was likely to occur based on the patterning principles of repeat and near repeat victimisation, was effective at predicting the immediate future (i.e., within the next 7 days), but was less effective at predicting where crimes occurred for more distant periods of the future. However, in comparison to the Gi* statistic, for only one of the three crime types that were tested (burglary dwelling) did the prospective mapping approach perform better than Gi* for predicting where crime occurred within the next seven days. For longer periods beyond seven days, the prediction performance of the Gi* statistic was better than prospective mapping for burglary dwelling and assaults with injury. The prediction performance of where theft from the person offences occurred were similar for Gi* and prospective mapping for both the immediate future and periods beyond.

Hypothesis 7: GWR provides an effective means of determining at the local level the reasons why hotspots exist, and that these explanatory variables vary between hotspots

The spatial scale at which hotspots were identified in the current research (research study 7) was found to be at a much finer level than spatial regression analysis using GWR allowed. Therefore, GWR did not provide a means of determining why individual hotspots exist, and identifying if there were explanatory differences between these hotspots. In turn, this meant that a GWR modelling process was not able to inform a more data-rich approach to hotspot analysis that used explanatory variables alongside, or in replacement of retrospective recorded crime data.

Hypothesis 8: GWR analysis can be used for supporting long-term predictions of crime by examining how a change in explanatory variables can influence a change in future crime levels

Although, following the testing of hypothesis 7, GWR was not suitable for effectively determining at the local level the reasons why hotspots exist, it was found, however, that at a broader scale of analysis, GWR was suitable for identifying those variables that explained crime levels, and how these explanatory relationships spatially varied (research study 7). This included identifying those areas where crime levels were highest, and the strength (and reliability) of the relationship between the distribution of crime and explanatory variables. While other researchers have expressed some concern about the extent to which GWR results can be used for determining inference, the results that related to testing hypothesis 8 did, though, provide some promising indications of the changes in crime that could be anticipated if crime prevention activity was targeted to address the conditions relating to the explanatory variables that were identified as being significantly correlated to crime. That is, the results from testing the application of GWR modelling for supporting long-term predictions of crime by examining how a change in explanatory variables can influence a change in future crime levels did illustrate the value of this spatial regression technique for helping to inform crime prediction for strategic and policy forecasting purposes.

The eight hypotheses that were tested to frame the direction of the empirical studies for this PhD research have resulted in a comprehensive set of research results. In the sections that follow, the results from the current research are more collectively discussed in relation to how the research findings have implications on practice, policy and theory.

12.2. Measures for spatial crime prediction

To date, as the concept of predictive policing and techniques for predicting where crime is likely to occur have developed, very little attention has been given to the use and recommended standardisation of measures for calculating and comparing the accuracy of mapping techniques for predicting spatial patterns of crime. Hit rates provide a quick and easy means of measuring how many crimes were successfully predicted, but they require the size of the prediction areas to be controlled in order to provide comparisons. The Prediction Accuracy Index (PAI) was introduced as a measure to help address this problem by taking into consideration the size of the areas that were determined to be *hot*. However, single global measures such as the PAI and hit rates restrict the examination of how effective the technique is for crime prediction at different spatial scales.

Johnson et al. (2008b, 2012) illustrated the use of accuracy concentration curves as a means of improving the measurement of spatial crime prediction performance of different mapping techniques. The accuracy concentration curves approach allows for the spatial prediction performance of the mapping technique to be compared across the full range of spatial scales of a study area by charting the number of crimes that were successfully predicted against the incremental percentile coverage of the study area. However, this process typically requires the researcher to make some visual comparison between different mapping techniques' accuracy concentration curves, rather than determining from the charts some value that provides a simpler metric comparison. In this PhD research, the use of accuracy concentration curves has been taken forward by drawing from the analogy of similar metrics used for determining the effectiveness of treatments and programmes in other scientific disciplines (e.g., medical tests and clinical trials into the effectiveness of drugs and other treatments). This involved introducing a simple, yet effective means of measuring the area under the accuracy concentration curve and standardising this to an index – the Crime Prediction Index. The CPI can be used to determine how good a mapping technique is at predicting spatial patterns of crime. The research has also shown the importance of generating these area under the curve measures and CPI values for several sub-sections of the accuracy concentration curve to determine how the mapping technique differs in its spatial prediction performance at different spatial scales.

In the research studies (chapters 5-11) that used the range of spatial crime prediction metrics described above, the results on the performance of mapping techniques for predicting spatial patterns of crime were presented across several sub-sections in each chapter. Here I introduce a measurement template that captures the key metrics for documenting the spatial prediction performance of a mapping technique's output. Figure 12.1 provides an illustration of a completed spatial crime prediction measurement template. Figure 12.1a shows burglary dwelling Gi* hotspots (using a grid cell size of 150 m) for the district of Newcastle-upon-Tyne. Several hotspots are identified, with the suggestion being that these are the areas where burglary dwelling is predicted to occur. Figure 12.1b and 12.1c show accuracy concentration curves for the Gi* output, for the full study area coverage and for the Gi* coverage that represents 1% of the study area coverage (i.e., those geographic cells that represent the top 1% of Gi* values). In Figure 12.1b, the more vertical the accuracy concentration curve (especially for small proportions of the study area), the better the spatial crime prediction. Figure 12.1c provides an indication of the proportion of crime that could be prevented if policing resources were allocated to the area representing the top 1% of Gi* values. The measure of curve gradient on the accuracy concentration curve graphs can be determined from the Crime Prediction Index, where a value of 1 indicates a perfect prediction. Figure 12.1d lists the CPI values for the spatial crime predictions determined using the Gi* statistic (for three statistical significance thresholds). A completed measurement template could then be used to compare against the spatial crime prediction performance of output generated from other mapping techniques.

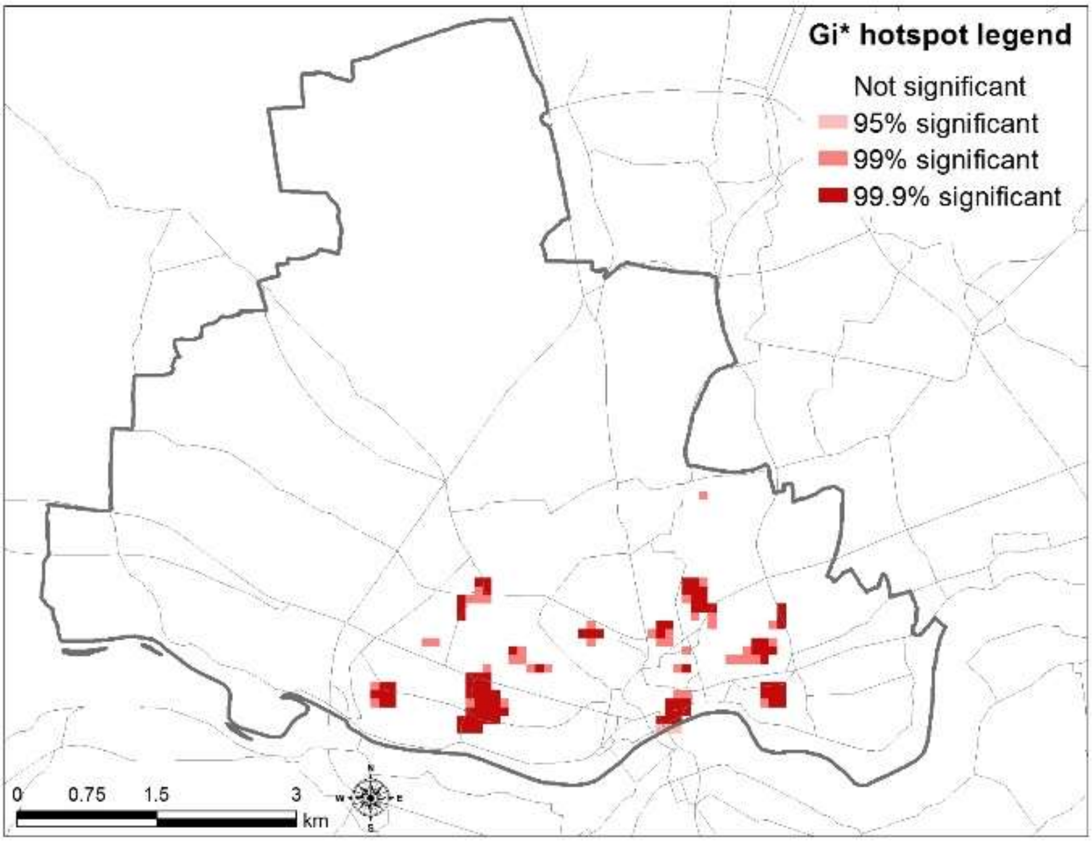

(a)

(b)

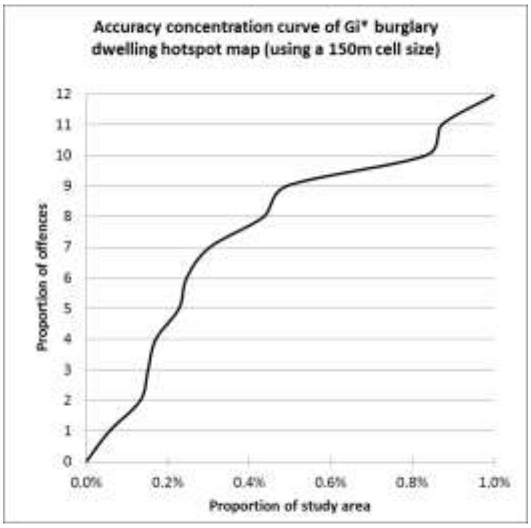

(c)

	Gi* statistical significance levels		
	99.9%	**99%**	**95%**
Proportion of study area	1.78%	2.52%	3.17%
Crime Prediction Index	0.963	0.955	0.955

(d)

Figure 12.1. Newcastle burglary dwelling (a) Gi* hotspot map, accuracy concentration curves for (b) the full coverage and (c) 1% of the study area coverage, and (d) CPI values for Gi* statistical significance threshold levels

The ability to use the spatial crime prediction measurement template to compare between mapping techniques is illustrated using a second example shown in Figure 12.2. Figure 12.2 shows the results for mapping output generated using kernel density estimation (using a cell size of 30 m and a bandwidth of 300 m). In this example from Figure 12.2, because KDE does not statistically determine the areas that are *hot*, 0.5%, 1%, and 5% of the study area were chosen to illustrate CPI values. Both the Gi* and KDE results from Figures 12.1 and 12.2 show they are very good at predicting spatial patterns of burglary dwelling at small spatial scales, illustrated by the near vertical gradient of the accuracy concentration curves for very small proportions of the study area and the CPI values (as shown in both Figures 12.1 and 12.2) being close to 1. The CPI values provide a useful means of quantifying the spatial prediction performance of different hotspot mapping techniques. This example shows that those areas where the concentration of crime was determined by the Gi* statistic to be significant at the $p<0.001$ level had a CPI value of 0.963 compared to a CPI of 0.846 for 0.5% of the coverage area produced from the KDE hotspot map. The results from this comparison, as illustrated using the spatial crime predication measurement template, suggests the Gi* hotspot map of burglary dwelling is a better predictor of spatial patterns of crime than the KDE hotspot map.

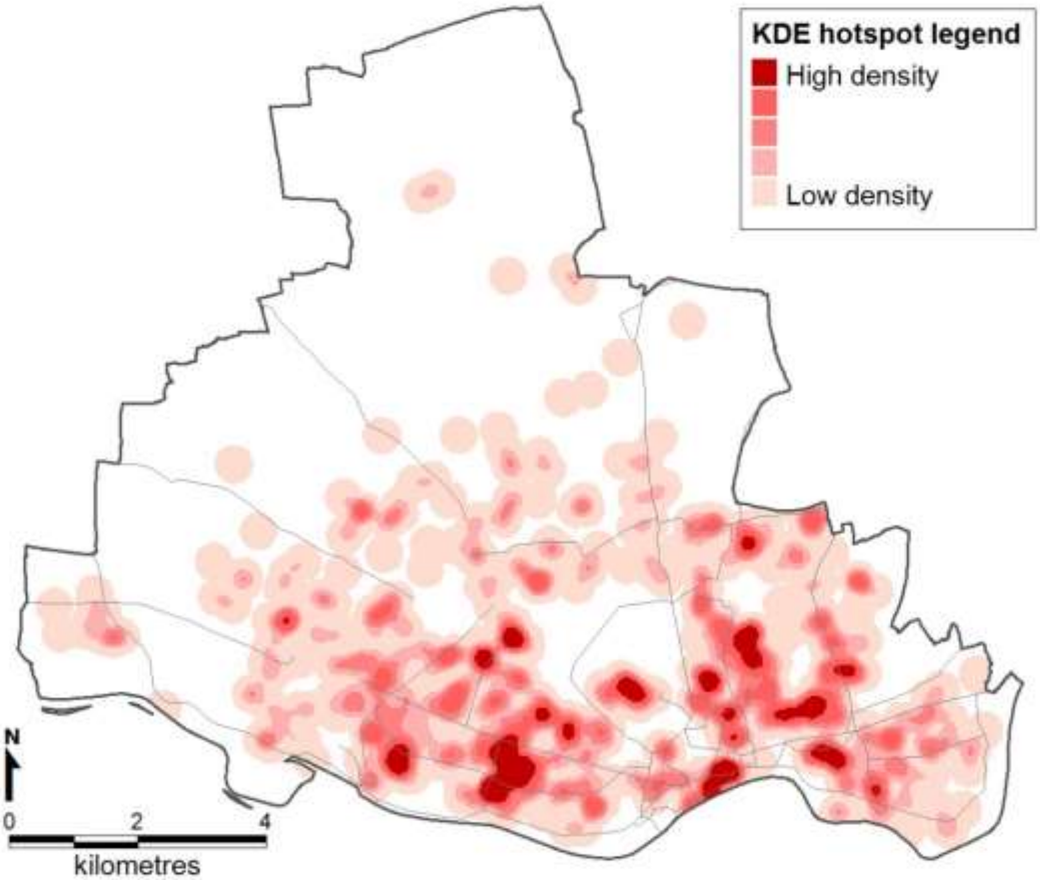

(a)

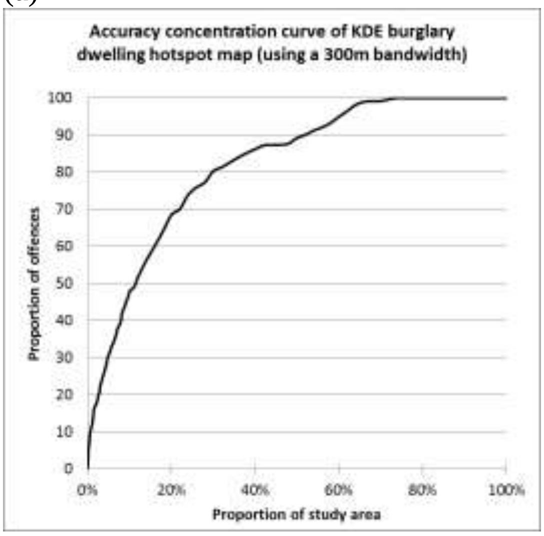

(b)

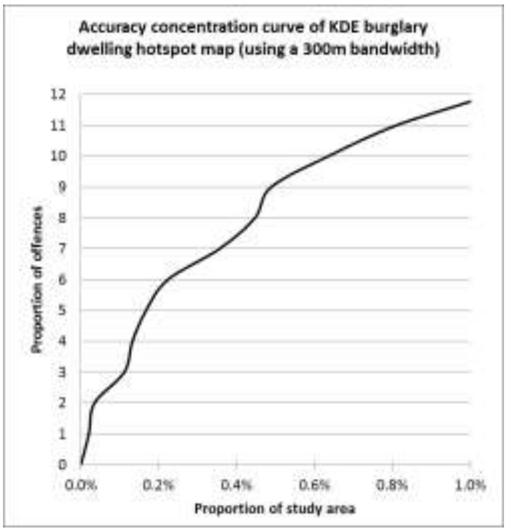

(c)

	Proportion of study area		
	0.5%	**1%**	**5%**
Crime Prediction Index	0.846	0.770	0.738

(d)

Figure 12.2. Newcastle burglary dwelling (a) KDE hotspot map, accuracy concentration curves for (b) the full coverage and (c) 1% of the study area coverage, and (d) CPI values for 0.5%, 1% and 5% of the study area coverage

12.3. Hotspot analysis: implications of the research findings

Hotspot analysis has become a common feature of crime analysis. Hotspot analysis uses data from the past to identify where crime has concentrated, and in turn uses this mapping output to help determine where police and public safety resources can be targeted. However, to date, a rigorous examination has not been conducted of the differences in the spatial prediction performance of hotspot mapping techniques, the influence of technical parameters and the influence of different retrospective periods of input data on hotspot analysis results. Since the commencement of this PhD research, developments in predictive policing have gathered pace, with new techniques being introduced that claim to offer more accurate means for determining where crime may occur in the future. While the purpose of this PhD research has not been to evaluate these new methods, it has aimed to provide a comprehensive metric examination of hotspot analysis, and in so doing establish a set of benchmark results against which other techniques can be compared. This detailed metric examination has also involved examining the technical features of hotspot analysis and certain statistical procedures that are valuable in helping to ensure that hotspot map production follows some good practice principles.

An often overlooked preliminary process of hotspot mapping involves testing (statistically) whether hotspots are evident in the crime data being examined. If there is no evidence of spatial clustering, any subsequent attempts at hotspot mapping would be futile and potentially misleading. With the increasing international adoption of and improvements in electronically recorded crime records and the geocoding of these records to the exact location where each crime event occurred, the Nearest Neighbour Index has become the preferred measure for determining if there is statistically significant evidence of hotspots being present in the data that are examined. Where crime data are only available as aggregated counts to geographic units, spatial autocorrelation measures such as Moran's I and Geary's C should be used to determine if there is statistical evidence of hotspots. The research has shown there is no single rule for how many geocoded crime records are required for clustering to be evident because the number can vary between crime types and for different study areas. Therefore, when an analyst is confronted with the question of whether they have enough data for hotspot mapping, the simple procedure of testing for clustering using the NNI would determine whether the application of hotspot mapping is worthwhile.

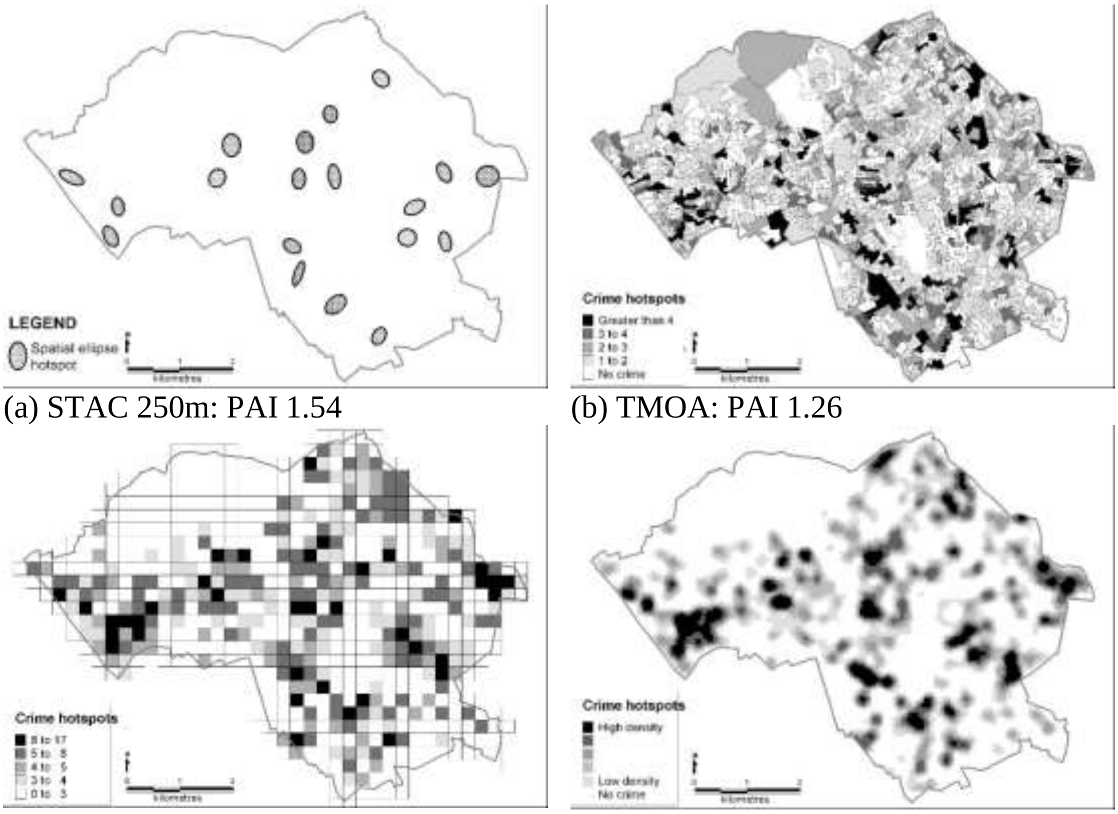

(a) STAC 250m: PAI 1.54 (b) TMOA: PAI 1.26

(c) GTM 250m: PAI 1.64 (d) KDE: PAI: 2.18

Figure 12.3. Hotspot maps generated from three months of burglary dwelling input data for the Camden/Islington study area using (a) spatial ellipses, (b) thematic mapping of Output Areas, (c) grid thematic mapping, and (d) kernel density estimation. Each map is shown with its PAI value, based on one month of output data.

There are a number of techniques that are commonly available in GIS to analysts and researchers for mapping hotspots. Kernel density estimation has increasingly becoming the hotspot mapping technique of choice, partly based on the findings from previous reviews that have profiled the technique's ability to outperform others in accurately identifying the location, size, orientation and spatial distribution of the underlying point data, and the visual appeal in the output that the KDE technique generates (Chainey and Ratcliffe, 2005, Eck et al., 2005). This PhD research has shown that KDE is also the best of the commonly used hotspot mapping techniques for predicting spatial patterns of crime. Grid thematic mapping proved to be slightly better than thematic mapping of geographic administrative units (using output areas), while standard deviation spatial ellipses were the worst at predicting spatial patterns of crime. As an example, Figure 12.3 shows hotspot maps generated for each of these techniques from three months of burglary dwelling input data for Camden and Islington when the measurement date was the 1st January 2010. The figures show that each technique identified similar areas, but in terms of the ability to predict future spatial patterns of burglary dwelling over the next month, KDE was better at predicting where burglaries dwellings did occur. While different techniques may be more suitable for certain scenarios (e.g., thematic mapping of geographic administrative units is suitable when wishing to compare changes in crime for jurisdictional performance assessment purposes), these results suggest that KDE should be the analyst's technique of choice when assisting in determining the targeting of resources.

KDE, like many mapping methods, requires the user to determine certain technical parameters as inputs to the spatial calculations: the cell size and bandwidth size. The results from this PhD research (research study 3, chapter 7) show that KDE hotspot maps generated using different cell sizes have little impact on the mapping outputs ability to predict spatial patterns of crime, but that different bandwidth sizes do have an impact. Cell size mainly impacts on the visual appeal of the KDE mapping output, with higher resolutions producing maps that avoid the *blocky* pixilation of outputs generated using larger cell sizes. For example, the maps shown in Figure 12.4 are equally as good as each other for predicting where crime may occur in the future. However, Figure 61a, which uses a much smaller cell size, is preferable due to its more appealing visual representation of the density distribution of crime. While smaller cell sizes require greater computer processing due to the larger number of calculations that are required. This extra length of processing was not a significant impairment in the experiments that were conducted.

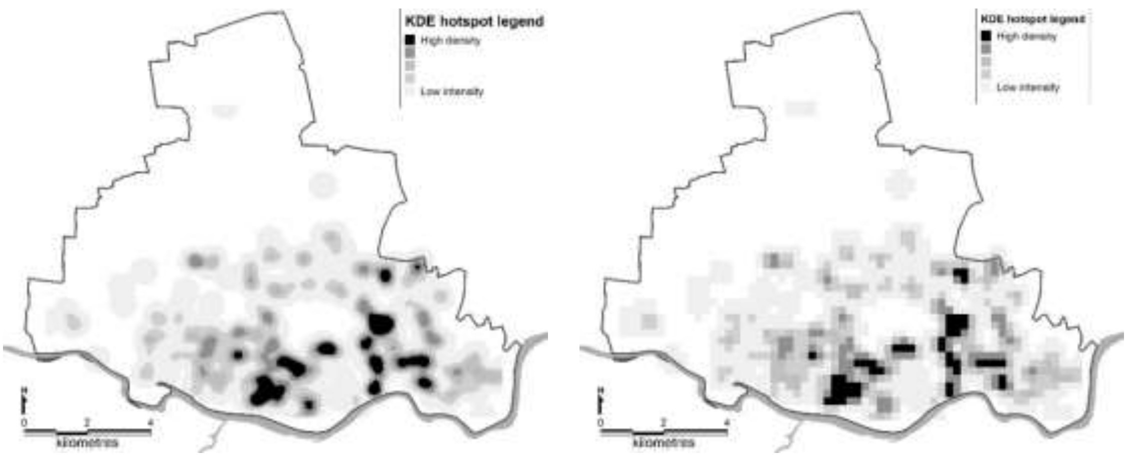

Figure 12.4. A comparison of KDE burglary dwelling hotspot maps for Newcastle, generated using the same bandwidth but with different cell sizes (a) 30 m (PAI of 7.0) and (b) 240 m (PAI of 7.0)

Bandwidth size does though affect the performance of KDE hotspot maps to predict spatial patterns of crime. For example, the maps shown in Figure 12.5 were generated using the same input data, the same cell sizes, but different bandwidth sizes. The PAI values for the KDE map produced using the smaller bandwidth was 119 compared to a PAI of 40 for the KDE map produced using a larger bandwidth. That is, the smaller the bandwidth, the better the KDE map is at predicting spatial patterns of crime.

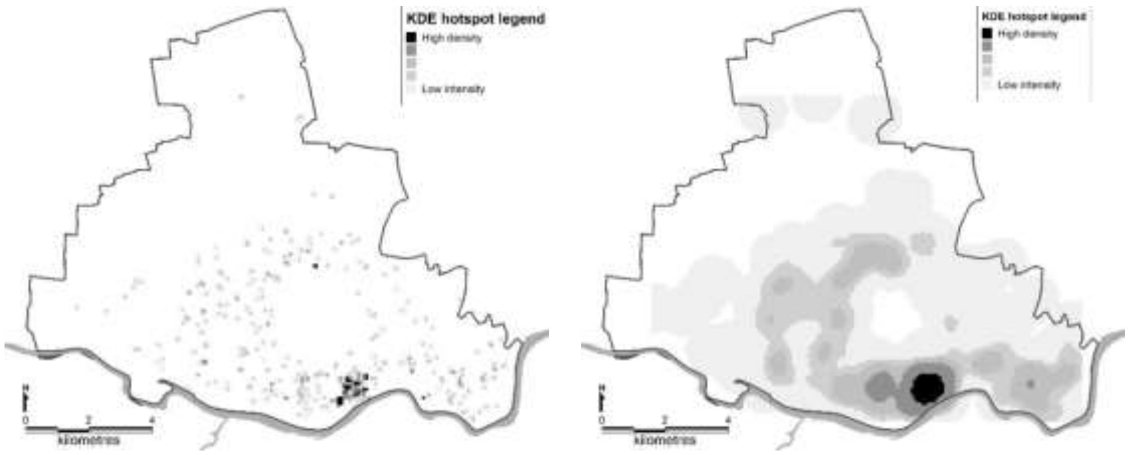

Figure 12.5. A comparison of KDE assault with injury hotspot maps generated using the same cell size but with different bandwidth sizes: (a) 100 m (PAI of 119.3), and (b) 800 m (PAI of 40.4)

The results from this PhD research of different cell sizes and bandwidth sizes now offer practitioners the means to better qualify the default parameter values that are determined by GIS products such as ESRI's ArcGIS Spatial Analyst and Crime Analyst extensions, and Crime Profiler and Hotspot Detective for MapInfo. The results indicate that defaults for cell size such as those generated using Hotspot Detective (which involves dividing the shorter side of the MBR by 150) offer a useful starting point, but reducing this value further will generate maps of greater visual appeal without affecting the map's ability to predict where crime is likely to occur in the future. However, bandwidth default values need scrutiny by practitioners to ensure they are not too large and impair the purpose of the KDE hotspot mapping output. For example, the default Hotspot Detective KDE bandwidth size for three months of violent assaults data for Newcastle-upon-Tyne was 450 m – a bandwidth size that generated a hotspot map with PAI value of 60, compared to a PAI of 143 if a bandwidth of 100 m was used.

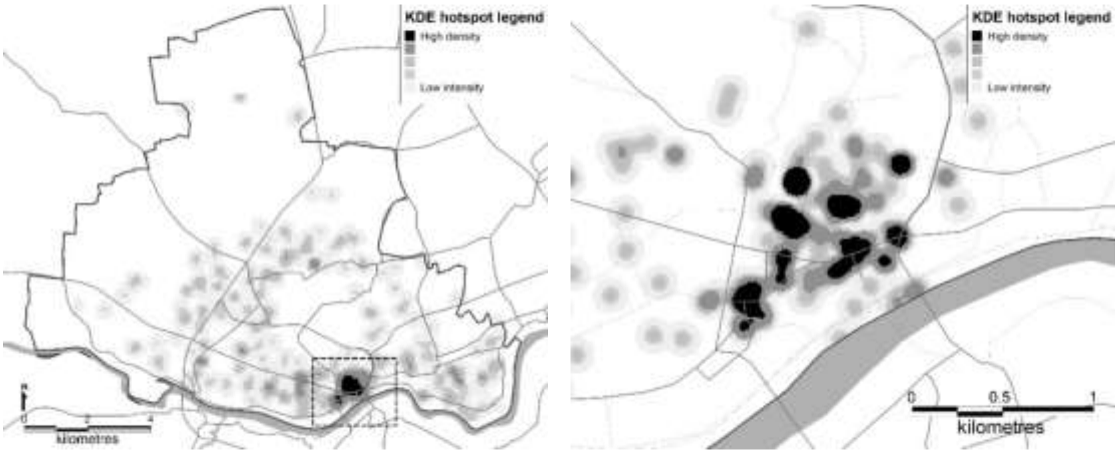

(a) (b)

Figure 12.6. A procedure for creating precise and practical KDE hotspot maps for accurately assisting in the targeting of policing and crime prevention resources: (a) is a KDE hotspot map generated for a large area for identifying the key areas for focus (bandwidth 300 m; cell size 90 m). Once a focus area is identified, data for this area is selected and a KDE hotspot map is generated using a smaller bandwidth (100 m) and smaller cell size (10m).

The arguments so far suggest that all KDE hotspot maps should be produced using small cell sizes and small bandwidths. However, small bandwidth values produce KDE hotspot maps that appear *spikey*, with many small areas identified as hotspots. In practice, this type of hotspot map is often considered unsuitable because it does not sufficiently narrow down the number of areas that require operational attention. Therefore, it is argued that a balance is required between KDE hotspot prediction accuracy, and output that is useful in practice. A way in which this can be overcome is to use a bandwidth size that is large enough initially to identify key hotspot areas, with these areas then being focused upon and a second hotspot map generated based on the distribution of crime in focus area. Figure 12.6 illustrates an example of this. Figure 12.6a uses a bandwidth size of 300 m and cell size of 90 m to identify the assault with injury hotspots in Newcastle-upon-Tyne. The main hotspot then becomes the area of attention, with a second KDE hotspot map generated for this area to more precisely identify the areas that are required for police attention. Figure 12.6b was generated using a bandwidth of 100 m and a cell size of 10 m. This two-step process to KDE hotspot map production offers the benefit of producing a visually accurate and appealing hotspot map for the study area, showing those areas where crime most concentrates, and then precisely showing the areas where crime hotspots exist and retaining visual appeal in the KDE mapping output.

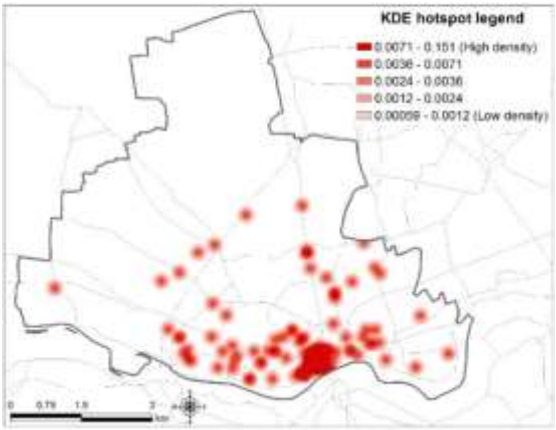

(a) KDE quantile thematic classification

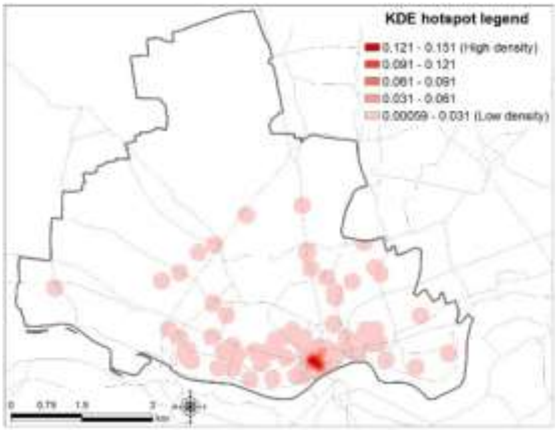

(b) KDE equal intervals thematic classification

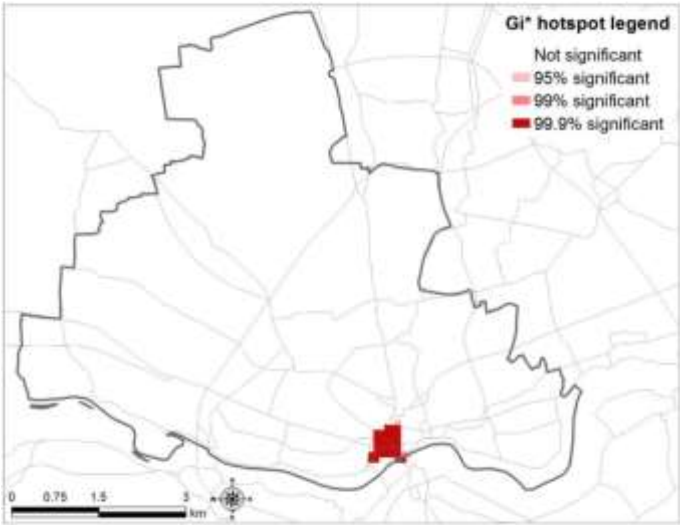

(c) Gi* statistical significance thematic classification

Figure 12.7. Newcastle theft from the person hotspot maps, produced using KDE and (a) the quantile thematic classification method, and (b) the equal intervals thematic classification method, and (c) using the Gi* statistic.

KDE is not without its weaknesses. The procedure described above could fail to identify areas where there is a high and spatially compact concentration of crime when the entire study area is examined because of the smoothing characteristic of KDE. An additional weakness is that the KDE requires the researcher to determine what is *hot* by deciding the value for the top thematic class. In most GIS software, several options are offered for the user to determine a thematic classification method preference, leading to subjectivity in KDE hotspot mapping output. The current research has illustrated how the Gi* statistic can overcome much of the subjectivity in defining the areas that are *hot* in hotspot analysis, and improve on (or at least retain) high levels of spatial crime prediction that good KDE hotspot analysis can offer. That is, in a statistical sense, the Gi* statistic can define those areas that are hotspots. The application of the Gi* statistic in the current research also included using the Bonferroni correction procedure to help address the issue of multiple testing, and set statistical significance thresholds that provide a spatially focused approach to identifying hotspots of crime. For example, Figure 12.7 shows hotspot maps of theft from the person generated using KDE and Gi*. While the hotspot areas identified are similar, the Gi* approach helps remove much of the ambiguity in defining the areas that are *hot* compared to the two KDE examples that define different hotspot areas, simply due to the thematic classification method that was used. The Gi* statistic is increasingly available to analysts and researchers, with it being packaged in ESRI ArcGIS since version 9.3, to MapInfo users in the HS Gridder add-on (Mashford, 2008), and available as a free Excel add-on (importing the results into any GIS) in the Rooks Case tool developed by the University of Ottawa (Sawada, 1999).

Gi*, like KDE, requires the researcher to determine certain input parameters. However, the current research has shown that, compared to KDE bandwidth sizes, the lag distance has less of an effect on the spatial prediction accuracy of the hotspot mapping output, albeit with small rather than large lag distances being preferred. The lag distance is calculated from the cell size that is used, and by including only those immediate neighbours (the eight cells surrounding the cell of interest). This approach to calculating the lag distance helped to retain a focus on generating local detail in Gi* hotspot analysis, minimising the problems of edge effects. However, one advantage that KDE retains over Gi* mapping output is the visual appeal of the hotspot maps that KDE generates. Gi* hotspot maps cannot simply be improved by using smaller cell sizes because this results in many cells containing zero or low counts. The use of small cell sizes has the knock-on effect of the sum of the count of crime in the cells representing the local neighbourhood of analysis to also be low in number, with little then to compare between these local neighbourhood averages and the global average. A procedure to improve this would be to use small cell sizes and to increase the lag distance beyond the eight immediate neighbours, but in doing so this may cause further problems with multiple testing. For example, rather than each cell being used nine times in the calculation of Gi* values, widening the lag distance will increase the number of times each cell is used to calculate Gi* values. This calls for more research to explore how the visual output of Gi* hotspot maps can be improved, while addressing these technical spatial statistical challenges.

The current research has also revealed differences in the spatial prediction performance of hotspot maps between crime types. For example, Figures 12.8a and 12.8b show Gi* hotspot maps of burglary dwelling and theft from the person in Newcastle. Each hotspot map identifies areas that were statistically significant. However in Figure 12.87b, the hotspot areas (determined using a 95% significance threshold) were where 49% of thefts from the person were committed in the six months that followed, compared to the hotspot areas in Figure 12.8a where 18% of burglaries were committed in the same following period. This indicates that the prediction of where certain types of crime occur is easier than it is for others.

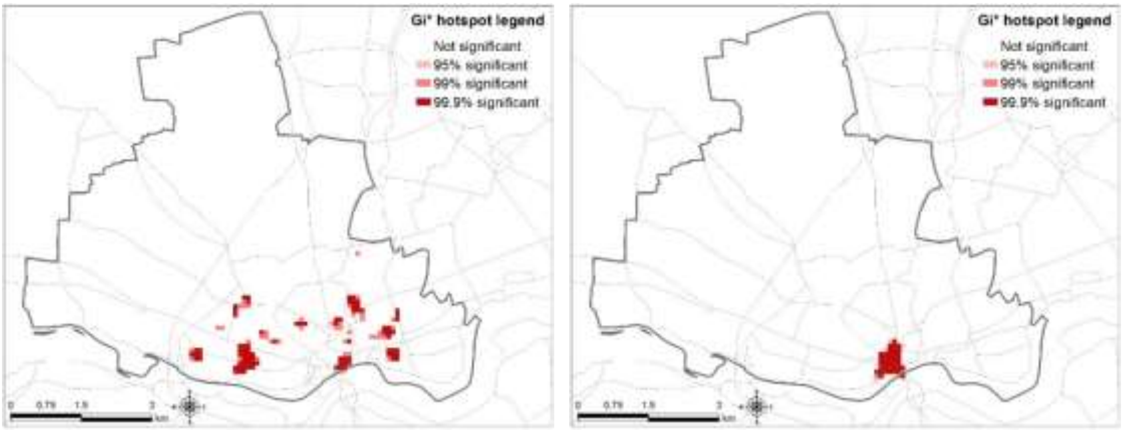

(a) Burglary dwelling (b) Theft from the person

Figure 12.8. Newcastle Gi* hotspot maps of (a) burglary dwelling and (b) theft from the person, generated from six months of input data

The results from the current research have also shown that hotspot maps of street crime offences against the person (such as assault with injury and theft from the person) were consistently better in their spatial prediction performance than any of the other crime types (see Table 12.1). Closer examination of hotspot maps of assaults and thefts from the person suggest that the reason for better levels of prediction performance is most likely due to the manner in which opportunities for these types of crime tend to concentrate. The areas where assaults and thefts from the person offences predominantly occurred were in areas where shops, bars, restaurants, markets and other forms of retail and entertainment concentrate. Areas with these types of land use offer many opportunities to commit street crime offences. This type of land use also tends to be clustered at particular localities, which in turn results in the highly concentrated spatial distribution of people to these localities, suggesting that the opportunity for street crime against the person would similarly be highly concentrated. These types of land use also tend to be static, in that they do not shift around the urban landscape but instead become a stationary part of an area's environmental fabric. Crime patterns tend to follow opportunities to commit crime (Cohen and Felson, 1979; Cornish and Clarke, 1986). Hence, as the opportunities to commit street crime remain highly concentrated and fairly static in geographic space, it is likely that retrospective data on where street crimes have occurred previously would be a good indicator of where street crime may occur in the future. This is the most likely reason that hotspot maps of assaults and thefts from the person generated high spatial prediction values.

Table 12.1. Proportion of crime predicted in Newcastle Gi* hotspots (defined using a 95% statistical significance threshold, 150 m cell size and six months of input data) in the six months that followed

Crime type	**Proportion of crime predicted in Gi* hotspots**
Burglary dwelling	18%
Theft from motor vehicle	26%
Theft of motor vehicle	21%
Theft from the person	49%
Assault with injury	37%

In comparison to the prediction performance of hotspot maps generated for assaults and theft from the person, opportunities to commit thefts from vehicles and thefts of vehicles, while they may be concentrated to certain places, tend to be more diffuse. Parked vehicles can be found in many locations – in garages, on driveways, on the street, and in car parks. This means the opportunity to commit crimes against vehicles is more widely spread than the opportunities that exist for committing assaults and thefts from the person. While hotspots of vehicle crime may occur and be influenced by land use, and the socio-economic and physical characteristics of areas, the wider geographic spread of where vehicles can be found (and targeted by offenders) would result in crime patterns for vehicle crime to be similarly more dispersed than where assaults and theft from the person offences are committed. In turn, the wider spatial distribution of opportunities to commit vehicle crime suggests that retrospective data on where vehicle crime occurs would most likely be less effective than retrospective data on assaults and theft from the person for predicting spatial patterns of future offences. Similarly, residential properties are typically more geographically spread than retail and entertainment facilities. While certain types of property may be more prone to burglary than others, the opportunity to commit burglary is not as heavily concentrated as opportunities to commit assaults and thefts against the person. This difference in the geographic distribution of opportunities is again the most likely reason that the spatial prediction performance of burglary dwelling hotspot maps were lower than those for assaults and thefts from the person.

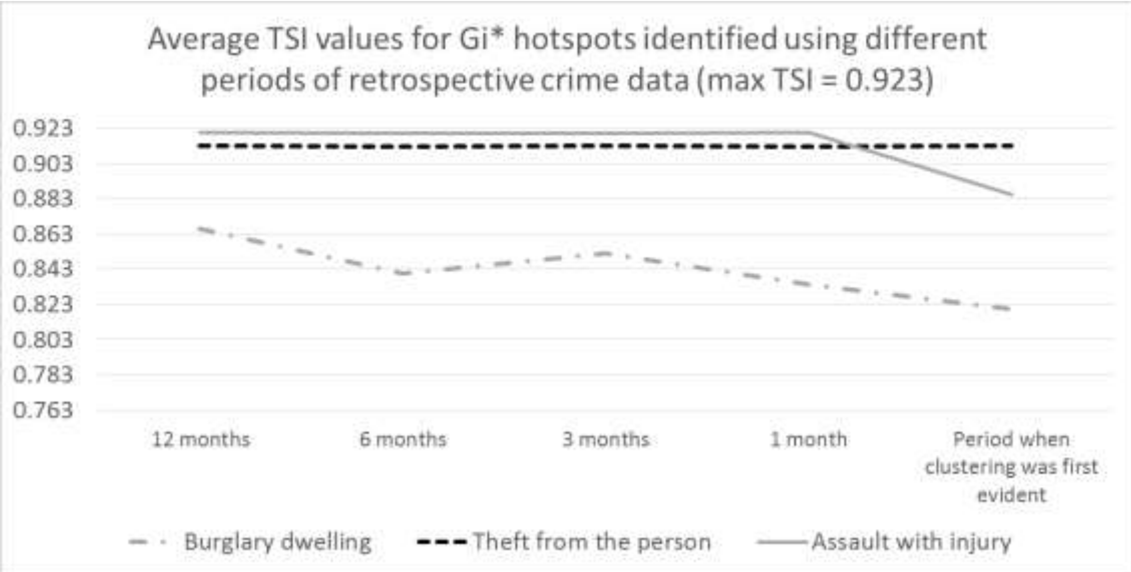

Figure 12.9. Temporal stability index values for Gi* hotspots identified using different retrospective periods of input data, for burglary dwelling, theft from the person and assault with injury. The average TSI value was taken from the TSI values for each hotspot identified from each retrospective input data period.

The differences in the ability to predict spatial patterns of crime for different crime types can be further illustrated through the analysis of the temporal stability of hotspots. Previous research has suggested that spatial patterns of burglary tend to shift frequently (Johnson et al., 2008a; Johnson and Bowers, 2004b). The results from this current PhD research have countered this finding by illustrating a consistent stability in hotspot patterns for each of the crime types analysed. For burglary dwelling, theft from the person and assault hotspots, Temporal Stability Index values were consistently above 0.85 (where perfect stability was indicated with a TSI of 0.923 and perfect instability indicated with a TSI of zero). That is, the results from the current research indicate that where crime has previously concentrated is where crime is likely to persist. However, through the analysis of different input periods for hotspot map generation, different results were found. For theft from the person and assaults, TSI values were very similar for all periods of retrospective crime data that were used. For burglary dwelling, TSI values showed a different trend, reducing in line with shorter input periods (see Figure 12.9). The results into the differences in the temporal stability of hotspots based on the period of retrospective data that were used suggest that for thefts from the person and for assaults, the same hotspot areas were identified when both short and long retrospective periods of crime input data were used, and the spatio-temporal patterns of these types of crime were highly stable. However, for burglary dwelling, the temporal stability of hotspots was lowest at the point that crime concentration was first evident, and increased as more retrospective data were used. These results suggest that spatial patterns of burglary dwelling tend to be less stable when only a short retrospective period of input data are used.

So far, the research results show that where crime has previously concentrated is where crime is highly likely to concentrate again. This is a research finding that has been consistently reported in many previous studies of hotspot analysis. However, the current research offers a comprehensive metric examination across techniques, crime types and for different retrospective periods of crime concentration. Research study 6 (chapter 10) examined whether the prospective mapping prediction method that uses recent individual incidents, and hence before the point that spatial concentration is likely to be evident, produces better spatial predictions of crime than Gi* hotspot mapping output. The results from research study 6 showed that spatial predictions of crime using prospective mapping were highest for the immediate future (i.e., within the next 7 days), but reduced as the temporal period of prediction increased (see Figure 12.10). The results also showed that spatial predictions generated using prospective mapping were only better for burglary dwelling and that Gi* hotspot analysis was just as good at predicting the immediate future for theft from the person and assaults. Beyond the immediate future, the results showed that Gi* hotspot analysis was better than prospective mapping for predicting where crime was likely to occur. These findings that compare the prediction performance of Gi* hotspot analysis and prospective mapping using different periods of retrospective crime data for different crime types suggest that prior to the point of spatial clustering appearing in retrospective crime data, prospective mapping provides an effective means of predicting where crime is likely to occur (and most so for burglary dwelling), but these predictions are most accurate only for the immediate future. The predictions for this immediate future are for single offences that follow repeat and near repeat patterns. Beyond this immediate future is when many single offences begin to cluster, with the results suggesting that hotspot maps generated using Gi* provide a more accurate means of predicting where crime is likely to occur than the prospective mapping approach. That is, once the point of spatial clustering in *prospective* crime data is reached, it is likely that a hotspot map generated from the clustering from *retrospective* crime data will be more accurate than using single very recent incidents in determining where crime is likely to occur. Therefore, from the point of spatial clustering in retrospective crime data, a Gi* hotspot map will provide a more accurate spatial prediction of where crime is likely to cluster than the predictions produced using the prospective mapping approach.

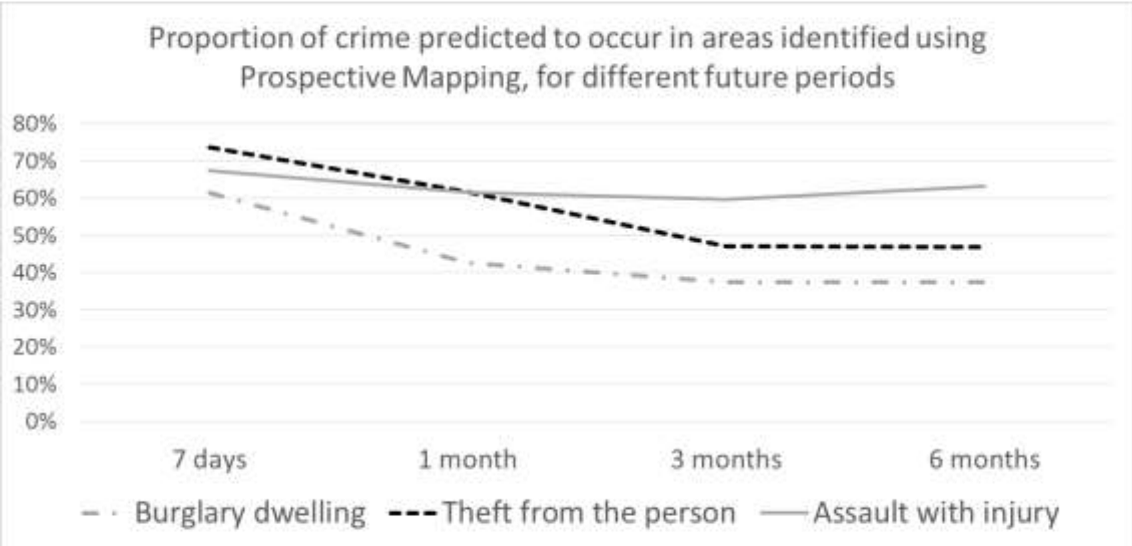

Figure 12.10. A comparison of the proportion of crime predicted to occur in areas identified using prospective mapping, for different future periods.

The findings from the current PhD research suggest the need to consider a two-step process to inform spatial crime prediction. The first step involves using prospective mapping and very recent incidents to predict the immediate future (i.e., where crime may occur over the next few days). The second step uses Gi* hotspot analysis and retrospective incidents that show evidence of clustering to predict where crime is likely to occur beyond the immediate future (i.e., the *near* future). As the vogue of predictive policing has developed, police agencies have increasingly erred to the use of single techniques to inform their spatial crime predictions and operational service delivery. The results from this research suggest caution is required in expecting a single technique to provide accurate predictions for both the immediate and near future. Instead, more accurate predictions can be generated using a two-step approach.

12.4. Explaining why crime hotspots exist: using the *causes* of crime to produce crime predictions

The original motivation for researching spatial regression was to help quantify why hotspots exist, identify causation differences between hotspots, and then use these results to explore whether data, other than retrospective crime data, improved spatial predictions of crime. The research has shown that at present, data for explanatory variables are typically not spatially precise enough to permit causation distinctions to be made between hotspots. In addition, GWR analysis is not spatially precise enough for determining inferential differences between crime hotspots at the spatial scale that hotspots are identified. The spatial regression analysis of crime data also presents several other challenges, particularly with regards to the treatment of highly clustered spatial patterning and the handling of multiple units containing zero counts. While Poisson models provide a viable option for the regression modelling of crime that typically display these spatial qualities, issues such as overdispersion and the spatial clustering of residuals may remain. To identify and address these issues requires the researcher to be knowledgeable and strict in the application of statistical diagnostic procedures that test for model suitability. The researcher may also have to treat their data, such as perform logarithmic transformations and select data sub-sets for processing (e.g., removing geographic units containing zero crime counts), in order to ensure the model is not bias and so that confidence can be placed in the results. The interpretation of model results also relies on good knowledge of theoretical principles for explaining spatial patterns of crime, regardless of whether an exploratory or hypothesis testing approach is used in calibrating explanatory variable inputs. The current research has also illustrated the iterative manner in which spatial regression modelling needs to be applied – involving a process of model variable selection, treatment of variables, interpretation, and model refinement – in order to generate models that not only perform well, but are also reliable and can be explained in theoretical and empirically evidenced terms.

While GWR modelling may not offer the ability to distinguish spatially varying relationships at the scale of hotspot identification, the research literature suggests potential in Bayesian spatially varying coefficient (SVC) models for analysis at more precise spatial scales. A key drawback of GWR modelling is the heavy smoothing that is a feature of the technique. Bayesian SVC models explore spatial relationships based on adjacency rather than a kernel size (Waller et al., 2007). The research literature also shows how Bayesian SVC models can improve model fit by avoiding the use of a single bandwidth applied to all explanatory variables (as applied in GWR) and using instead separate adjacency measures for each explanatory variable (Waller et al., 2007). In addition, it is suggested that Bayesian SVC models allow for more robust and accurate inferences to be determined, including reducing the levels of collinearity that can be a problem with GWR models (Wheeler and Calder, 2007, Wheeler and Tiefelsdorf, 2005) and how statistical significance prediction intervals can be generated (Wheeler and Waller, 2009). This calls for further research that examines whether Bayesian SVC models applied to crime data improve on the findings of the use of GWR from this PhD research.

GWR, though, has shown potential to offer an additional temporal dimension to spatial crime prediction. Through the spatial modelling of crime patterns and explanatory variables, those variables that help explain the spatial distribution of crime can be identified and whether the relationship between these variables spatially varies. While some caution is required in using GWR for inferential purposes, the results from this type of modelling can at least indicate which variables appear to significantly correlate with the distribution of crime and where these variables have their biggest impact. This can then help inform the direction of strategic crime reduction policy through the design, development and spatial targeting of initiatives that aim to reduce the impact these variables have on crime. In statistical terms, this would involve aiming to reduce the significant and spatially varying influence these variables have on crime. For example, where the distribution of the student population is significantly (positively) correlated with spatial patterns of burglary dwelling, a strategic initiative would be to reduce this high level of student vulnerability. Similarly, where the distribution of licensed premises is positively correlated with violent assaults, this finding could inform policy on the issuing of licenses to these premises and any additional premises that wish to open in areas where this relationship is most significant. These findings on the use of GWR for identifying spatially varying relationships between crime and explanatory variables suggests the need to consider a three-step process to inform spatial crime prediction, with spatial regression modelling informing long-term predictions.

12.5. Crime prediction: currency, concentrations and causes

Spatial crime prediction is reliant on the input of data variables for accurately informing where crime is likely to occur in the future. To date, little consideration has been given in practice to the actual prediction period of the future – is the aim to predict where crime is likely to occur in just the next day, the next week, the next month or the next year, or is the prediction assumed to be accurate for all periods of the future? In practice, agencies plan their service provision across all these temporal periods, albeit with some oriented more towards one than the other. Police forces very much focus on the here and now, oriented towards tactical responses for the next operational police shift that support reductions in crime; Community Safety Partnerships have a focus towards designing responses that may require contributions from a number of agencies, oriented towards implementing initiatives that may take several weeks to organise and that deliver sustainable reductions in crime; whereas a City Mayor or Police and Crime Commissioner may have a focus towards policy changes, introducing new programmes that reduce victimisation against those who are most vulnerable, and achieving reductions in crime. These variations in focus require consideration of the input data variables that are most suitable for informing these crime predictions for different temporal periods.

This PhD research has shown the value of hotspot mapping for accurately predicting spatial patterns of crime. It has shown that perhaps the most powerful variable for predicting where crime is likely to occur is where crime has previously occurred. Other variables such as land use, demography, and socio-economic conditions that influence the opportunity for crime to occur may explain why crime patterns spatially vary, but typically these data are not spatially precise enough to be used alongside or in replacement of crime data to improve spatial predictions of crime. Indeed, even if these variables were available at the detailed resolution of recorded crime data they may actually offer little to improve spatial predictions of crime that use only retrospective recorded crime data.

To date, KDE has become the hotspot analysis technique of choice amongst police and crime reduction practitioners and researchers. This PhD research has illustrated the accurate spatial crime predictions that the KDE technique can generate, and how it can be improved with attention to bandwidth choice. The current research has also shown how hotspot analysis can be improved further by using the Gi* statistic to help remove much of the ambiguity in defining hotspots and produce predictions that at least equal or are better than those produced using KDE. For some crime types, the spatial crime predictions generated using Gi* hotspot analysis were exceptionally high, illustrated by CPI values that were almost perfect. This finding helps set a benchmark for which other techniques can be compared.

The current research has also shown that hotspot analysis is better at predicting where hotspots are likely to exist in the future in comparison to prospective mapping. That is, patterns of retrospective spatial clustering are of great value in predicting patterns of future spatial clustering. When predicting the immediate future, differences were observed in the performance of prospective mapping and hotspot analysis. Prospective mapping, rather than hotspot analysis, appeared to be more accurate in determining where crimes against property, such as burglary dwelling incidents, were likely to occur. However, prospective mapping did not appear to offer any additional value in predicting where crimes against the person were likely to occur.

These findings on the differences in the prediction performance of prospective mapping for different crime types appears to be related to the spatio-temporal stability of crime patterns, with those crime types that are less stable and susceptible to spates committed by the same *foraging* offender being more suited to prospective mapping treatment. Where there is a very high degree of spatio-temporal stability in crime patterns and the commission of these offences is less to do with spates of offending being committed by a returning offender, prospective mapping appears to offer no improvement in the spatial crime predictions that are generated by good hotspot analysis. Therefore, while currency in retrospective crime data appears to be relevant for informing the choice of technique that is used to generate spatial predictions of crime, this currency needs to be considered alongside the type of crime that is being predicted and a clear theoretical understanding of why crime may tend to take place in the locations where it is predicted to occur. In practical operational terms, while accurately predicting where crime may occur in the next day, week, month or year will support the targeting of policing and crime reduction resources, understanding *why* crime is likely to occur in these locations will inform the actual activity conducted in these locations. Optimal foraging theory and the boost account help explain why an offender may return to an area where they recently committed a crime, with crime pattern theory helping to explain the preference in choice of area based on the individual's awareness of opportunities. The combination of these theories provides the foundation for explaining where crime is likely to be committed in the immediate future. As a result, several police agencies have already adopted tactics to directly counter this predictable foraging and boost behaviour (Chainey, 2012b). A gap in the theory appears, however, in drawing together crime pattern theory and other environmental criminology concepts for explaining why clusters begin to form in the near future, and why they may go on to persist for some time.

12.6. Theoretical developments: explaining why crime hotspots emerge using the watering hole principle

Optimal foraging theory, boost account theory and crime pattern theory help explain the preference of the individual offender to return to the same target to commit a crime (i.e., commit a repeat offence) or to a similar target in close proximity (i.e., commit a near repeat). In particular, optimal foraging and boost account theories help explain why these repeat and near repeat offences occur swiftly after an initial incident, with crime pattern theory helping explain why the offence is likely to occur in a place that is familiar to the offender, based on their routine activities and the distribution of opportunities. Each of these theories draws from the theoretical concepts of least effort: the knowledge and selection of opportunities to commit crime tend to result in a short journey to crime (and short journey to the place of foraging); and repeated offending trips take place because the offender is boosted in the knowledge learnt from the initial offence, rather than applying more effort to seek new opportunities. These theoretical principles of foraging, boost account, crime pattern theory and least effort help to explain individual offending behaviour and predictions on where crime may occur within the immediate period following an initial incident (i.e., within one week).

A gap appears, however, in how we can fully explain the individual behaviour of offenders that, when aggregated, helps to explain why the crimes they commit will begin to form hotspots. We can, of course, draw on crime pattern theory, the principles of least effort, the routine activity approach and the rational choice perspective to help explain why crime patterns tend not to be spatially random, but concentrate at certain locations. However, a specific theory that draws these theoretical principles together to explain where and why hotspots may begin to form in the *near* and more distant future (similar to how optimal foraging theory helps us to explain where and why crime may occur at certain locations in the *immediate* future) has yet to be articulated in the environmental criminology literature. Brantingham and Brantingham (1995) introduced the concepts of crime generators, crime attractors, crime enablers and crime-neutral sites to help make descriptive distinctions between different types of spatial crime distribution. In this section I introduce the watering hole principle to help advance the understanding of why crime tends to form hotspots, and make clear the distinctions between this new theoretical concept and the Brantingham's typology of spatial crime distribution.

There are also two distinctions to consider for explaining why hotspots of crime will tend to form. The first is that the crime committed in a hotspot may appear as an aggregation of many individual offenders committing multiple crimes in the same area and who are boosted by their commission of recent offences (i.e., several offenders committing spates of crime in the same area). The second is where the crime committed in a hotspot may appear as an aggregation of many individual offenders committing single offences that are not boosted by their commission of recent offences (i.e., multiple offenders committing single offences in an area). The similarity in both cases is that the areas where hotspots begin to form are likely to be where conditions are favourable for the commission of crime.

The watering hole principle helps advance the understanding of why crime tends to form hotspots. The principle also offers additional theoretical validity to why hotspot analysis can be effective in predicting where crime is likely to occur in the near future (at the point when individual offences begin to form clusters). The main distinction between the watering hole principle and the concepts of attractors, generators, enablers and neutral-sites is that the latter offer a set of *a posteriori* descriptive typologies for the range of spatial crime distribution, while the former provides an *a priori* theoretical explanation of spatial crime concentration. The distinction between the watering hole principle and crime attractors, generators, enablers and neutral-sites is elaborated on below after introducing the watering hole principle in full.

Purposely, the watering hole principle draws from ecology because it provides the best analogy for explaining this principle in a crime sense, and because it naturally fits alongside the criminological explanations introduced by Johnson et al. (2009) in optimal foraging theory. In addition, in the same way that the ecological concept of foraging links to the boost account, the ecological concept of the watering hole, as explained further below, links to the flag account. The idea of using a watering hole analogy is not new to crime. For example, Felson (1987: 914) stated, ‘Just as lions look for deer near their watering holes, criminal offenders disproportionately find victims in certain settings or high-risk occupations’. However, the idea of using this ecological analogy has not been expanded upon by Felson or others. As elaborated here, the analogy provides the primary theoretical explanation for locations where conditions are favourable for the commission of crime.

A watering hole in ecological terms is a pool where animals come to drink. For the watering hole to exist, favourable conditions have to be present: it can only form in a natural geographic depression, and not on flat land and where the water table reaches the surface; the soil and underground rock conditions must not cause the water to drain away; and rainfall needs to be frequent enough to replenish the watering hole. In an environmental criminology sense, a cluster of crime will form in areas where favourable conditions are present: it can only form in a place where suitable targets are concentrated; the targets are not removed; or if removed are frequently replenished. These core favourable conditions establish the theoretical basis for explaining prior to crime commission the locations that are likely to experience spatial concentrations of crime.

The watering hole principle provides an *a priori* explanation of spatial crime concentration. That is, the principle explains pre-event, the locations that are likely to experience spatial concentrations of crime. This is distinct to the concepts of crime generators, attractors, enablers and neutral-sites which provide a set of *a posteriori* descriptive typologies for the range of spatial crime distribution. That is, generators, attractors, enablers and neutral-sites describe post-event the range in levels of the geographic distribution of crime in relation to the characteristics of areas. Additionally, the typologies of generators, attractors, enablers and neutral-sites were constructed to describe the spatial distribution of crime from an offending viewpoint, with the focus being towards the offender's choice of targets, their target areas, and motivational desire (Brantingham and Brantingham, 1995). The watering hole principle considers the spatial distribution of crime from both an offending motivation viewpoint and from a vulnerability of victimisation perspective. That is, the principle considers favourable conditions in locations, both from the perspective that provide favourable conditions for the commission of crime and the conditions that are *favourable* for increasing the risk of crime. Furthermore, the application of generators, attractors, enablers and neutral typologies were devised for 'types of urban sites that need to be considered' (Brantingham and Brantingham, 1999: 7), whereas the theoretical concept of the watering hole principle is applicable for all landscape settings – from urban to rural, and all categories between. Generators, attractors, enablers and neutral-site typologies are also more suited to the micro and meso geographic scales, whereas the application of the watering hole principle can be applied to micro, meso and macro scales. For example, the notion of favourable conditions for offending or risk of victimisation can refer to particular targets (e.g., a unit on an industrial estate that has no perimeter fencing and easy egress to the estate's main entrance and exit exhibits more favourable conditions for burglary than those units located deep into the industrial estate and that have secure perimeter fencing), particular areas (e.g., city centre night-time economy areas where there are more favourable conditions for assault), and regions (e.g., student residential areas within a city where there are more favourable conditions for domestic burglary). Table 12.2 provides a summary of the conceptual differences between the watering hole principle and generators, attractors, enablers and neutral-sites.

Table 12.2. Conceptual differences between the watering hole principle and generators, attractors, enablers and neutral-sites

Watering hole principle	**Generators, attractors, enablers, and neutral-sites**
• Definition: locations where there are favourable conditions for crime to occur • An *a priori* explanation of spatial crime concentration – explaining pre-event the locations that are likely to experience spatial concentrations of crime • Considers crime from an offending viewpoint and from a vulnerability of victimisation perspective • Can be applied to the full range of urban to rural settings – focus is towards all types of favourable conditions for crime to occur: from the built environment (e.g., land use, housing type, street morphology) to demographic and socio-economic conditions that influence the spatial distribution of crime • Can be applied at micro, meso and macro geographic scales. For example, favourable conditions can refer to particular targets, particular areas, and regions • Can be used to explain the flag account theory – flags are enduring conditions that make certain targets more favourable than others	• Definition: description of different types of spatial crime distribution • A set of *a posteriori* descriptive typologies for the range of spatial crime distribution – describing post-event the range in levels of the distribution of crime in relation to an area’s characteristics • Describe the spatial distribution of crime mainly from an offending viewpoint rather than from a vulnerability of victimisation perspective • Have an urban and situational focus – focus is towards the built environment rather than also considering the influences of demographic and socio-economic conditions on the spatial distribution of crime • More suitably applied to the micro and meso geographic scales, and not the macro scale. For example, describing the range in levels of the distribution of crime in relation to characteristics of particular targets (e.g., differences between the types of shops) and particular areas (e.g., city centres and shopping malls), rather than differences between regions • Can not be used to explain the flag account theory of offender target selection as it provides a description for a range of different types of spatial crime distribution

Flag account theory is used to help explain why some enduring characteristic about a target makes it a higher risk for being vulnerable to crime. Watering holes are places where many *flags* may exist, offering a plentiful supply of suitable targets. In this sense, similar to how foraging theory links to the boost account to explain spates of crime committed by the same offender, the watering hole principle links to the flag account to explain the stable nature of crime concentration in places where many favourable conditions to commit crime exist.

The watering hole principle also has the flexibility in being able to be applied to the two offending distinctions described previously on why hotspots of crime form. Firstly, the watering hole is where many individuals (who are not prolific offenders) commit single offences because the conditions are favourable for the commission of crime. Secondly, the watering hole is where individual offenders who are more prolific in their offending activity commit many crimes and where their spates of offending overlap. The use of the watering hole principle is illustrated for burglary dwelling. Burglaries of residential properties can be predicted in the immediate future by using the spatial and temporal attributes of recent incidents – offenders will be attracted to forage for further opportunities to commit burglary by returning to the same property or nearby properties. Hotspots of burglary dwellings can be predicted for the near future by identifying where hotspots previously occurred. These hotspots can be predicted because offenders will be attracted to the *watering holes* where a plentiful pool of suitable targets are available, displayed either through the aggregation of multiple offender's foraging and boost behaviour (i.e., the collective aggregation of their spates of crime), through the aggregation of multiple offenders committing single offences, or a combination of the two. The current research has indicated that for the study area of Newcastle, the favourable geographic conditions that resulted in high levels of burglary dwelling were influenced by the area being in close proximity to where known offenders live, it being an area with a large student population, and/or an area with a large Asian population.

For some types of crime, the theoretical concepts of foraging and the boost account do not apply. For example, violent assaults in a busy town centre do not tend to result from a person who got into a drunken brawl the previous night returning the next night or the next weekend, foraging for another drunken brawl, boosted from the previous event. Therefore, it is not suitable to use the concepts of foraging and boost behaviour to explain why crimes of this type may occur. Spatial patterns of repeats and near repeats may still be observed in the distribution of violent assaults, but an explanation other than foraging and the boost account is required to explain these patterns. This current research has shown that good hotspot mapping is just as accurate, if not better, than prospective mapping for predicting spatial patterns of violent crime. To provide an explanation behind the effective prediction of this type of crime, the watering hole principle, rather than foraging and boost behaviour, is more suitable: the locations where future violent assaults are most likely to occur will be in the watering holes where a plentiful pool of suitable targets are available, displayed through the aggregation of multiple offenders committing single offences. This research has indicated that for the study area of Newcastle, the favourable geographic conditions that result in high levels of violent assaults are influenced by where pubs, bars and nightclubs are located, and/or by areas that experience high levels of income deprivation.

By using the watering hole principle to help explain where and why hotspots are likely to occur (and in clearer theoretical terms than offered by the *a posteriori* descriptive typologies of crime generators, attractors, enablers and neutral-sites), the explanations associated with the watering hole principle can also help determine how these predictable hotspots can be tackled. As the watering hole is where favourable geographic conditions are present, policing and crime reduction resourcing should be focused towards addressing these favourable conditions. Strategies could involve protecting the most vulnerable individuals, raising awareness that helps reduce the vulnerability of certain groups, making it more risky for offenders to commit crime, increasing the effort it takes for offenders to commit crime, and reducing the rewards from the commission of crime. With reference to the theoretical link between the watering hole principle and the flag account theory, the targeted focus of crime prevention activity in the watering holes would be akin to removing the enduring characteristics that make the targets vulnerable to crime.

12.7. Practice and policy implications: the Crime Prediction Framework - a temporal framework for spatial crime prediction

Previous research (e.g., Johnson et al., 2009) and the current research have shown how the spatial attributes of recent offences can be very effective in predicting where individual crimes may occur in the immediate future. The ability to be able to make these predictions of crime in the immediate future is based on the well-researched and frequently empirically observed spatial and temporal patterns of repeat and near repeat victimisation. However, the effectiveness of this prospective approach for predicting where crime is likely to occur decays as the temporal horizon extends – it is effective at predicting the immediate future, but beyond this timeframe the accuracy in the predictions begin to reduce. Once crime patterns begin to form into hotspots, the places where crimes previously formed hotspots appear to provide more accurate spatial crime predictions than using just the patterning principles of repeats and near repeats. A third temporal frame for predicting spatial patterns of crime is the more distant, long-term future. Through the spatial modelling of crime patterns against variables that are hypothesised to explain the spatial distribution of crime, the relationship with these explanatory variables can be quantified and used to inform the direction of strategic policy and predict how crime levels may change as a result.

To date, the attention to spatial crime prediction (so called predictive policing) has been towards using single, all-encompassing techniques to produce predictions. Often, little thought is given to how the currency of data may influence these predictions and to whether these predictions are more suitable for the immediate future (i.e., the next day), the near future (i.e., the next week or month) or are better at providing a long-term forecast (i.e., for several months and beyond). Little thought has also been given to whether the technique of choice is equally suitable for providing accurate predictions for all types of crime. The findings from the current research suggest it is not sufficient to consider that a single spatial analysis technique will be accurate for predicting where crime is likely to occur for all crime types and for all periods of the future.

To help illustrate this, I use a weather forecasting analogy. Data on current and very recent weather conditions are perhaps the best predictors of what the weather is likely to be like in the immediate future. To forecast what the weather may be like next month, data in addition to recent conditions would be used. To forecast what the weather may be like next year, data other than that on recent conditions and from just the last month would be used. Similarly, the analytical technique or model that is used to forecast what the weather may be like tomorrow is different to the technique or model used to forecast the weather outlook for next month, with another different technique or model being used to forecast what the weather may be like next year. Also, different models are used in different areas to reflect the predominant type of weather each area experiences. For example, different models are used in temporal climate zones to those used in tropical areas. Using this analogy for crime, it would appear unsuitable to use a single technique, with little thought given to the input data, to determine accurate spatial predictions of crime for different periods of the future. Therefore, a temporal framework is suggested for spatial crime prediction – the crime prediction framework.

The crime prediction framework consists of three temporal prediction periods – predictions for the immediate future, predictions for the near future and predictions for the distant future. However, due to important distinctions that explain spatial patterns of different types of crime, two frameworks are suggested – one where spatial patterns of crime in the immediate future can be explained using foraging and boost account behaviours, and a second framework where foraging and boost account behaviours cannot be used to explain spatial patterns of crime.

Table 12.3. The crime prediction framework - a temporal framework for spatial crime prediction, where the commission of crime can in part be explained using foraging and boost account theoretical principles

Time frame	Spatial analysis technique	Input data	Theoretical reasoning	Responses
Immediate future	Prospective mapping	• Crime records showing recent incidents	• Optimal foraging • Boost account	• Police tactics • Targeted offender supervision
Near future	Gi* hotspot analysis	• Crime records showing recent hotspots	• Watering hole principle • Flag account	• Crime prevention initiatives
Distant future	Geographically weighted regression	• Crime records showing persistent hotspots • Explanatory variables	• Background norms	• Strategic interventions and changes in policy

Table 12.3 shows the crime prediction framework and illustrates the key characteristics of the first temporal prediction period for spatial crime prediction. The framework suggests that for the purposes of predicting the immediate future, the prospective mapping technique should be used. These predictable patterns of crime can be explained using

foraging and boost account behavioural principles. These types of predictions for the immediate future are most likely suited to targeting police patrols in those areas where incidents are predicted, using the patrols' high visibility to deter any further offending, utilising stop and search on known offenders who are suspected to have recently committed incidents, and speaking to people who live or frequent this area, encouraging them to carry out practical crime prevention activity that will minimise their risk of victimisation. The immediate activity should also involve minimising the heightened risk of victimisation to the person or other target that has recently been victimised, and utilising offender supervision resources to help disrupt and deter the activity of those suspected to be involved in the commission of crime in this area.

For the purpose of predicting the near future, the crime prediction framework (as illustrated in Table 12.3) shows that hotspot mapping using the Gi* statistic should be used. These predictable hotspots of crime can be explained using the watering hole principle and the flag account theory. However, further analysis would need to be conducted on the hotspot to determine the favourable geographic conditions that cause this area to be a watering hole. Police and other agency activity should therefore be focused on addressing these favourable and enduring conditions that make crime particularly conducive in this area.

For the purpose of predicting the distant future and long-term change, the crime prediction framework illustrated in Table 12.3 shows that a GWR analysis of crime, against a hypothesised set of explanatory variables, is required. The variables that are significantly correlated in this type of modelling and that can be explained in clear theoretical terms would inform the direction of strategic interventions. Activity that is focused on addressing the influence these variables have on crime (where the relationship is significantly positive) and improving the influence these variables have on crime (where the relationship is significantly negative) would help bring long-term reductions in crime that focus on addressing the underlying norms that influence crime levels.

The second version of the crime prediction framework is applied to crime types where the theoretical concepts of foraging and boost account behaviour are not valid for explaining the commission of crime. This second framework would, therefore, only utilise the spatial analysis techniques of Gi* hotspot mapping and GWR analysis for predicting spatial patterns of crime. The second version of the crime prediction framework is illustrated in Table 12.4 and shows that the same principles apply to the framework illustrated in Table 12.3, with the exception of any prospective mapping approach. The immediate and near future time frames are combined, with response activity on where crime is likely to occur within these two time frames following the analysis of where recent hotspots occurred. The police tactics and focused offender supervision described in relation to the temporal period of the immediate future from Table 12.3 would instead be targeted towards where hotspots are predicted to occur.

Table 12.4. The crime prediction framework - a temporal framework for spatial crime prediction, where the commission of crime cannot be explained using foraging and boost account theoretical principles

Time frame	Spatial analysis technique	Input data	Theoretical reasoning	Responses
Immediate and near future	Gi* hotspot analysis	• Crime records showing recent hotspots	• Watering hole principle • Flag account	• Police tactics • Targeted offender supervision • Crime prevention initiatives
Distant future	Geographically weighted regression	• Crime records showing persistent hotspots • Explanatory variables	• Background norms	• Strategic interventions and changes in policy

These two crime prediction frameworks help direct a realistic response structure for reducing crime in the areas where it is predicted to occur – in the immediate, near and more distant temporal terms. In order to determine the types of response most suitable, a clear theoretical explanation for the patterns needs to be provided. These theoretical explanations also need to be sensitive and aligned to the different temporal response arrangements of different agencies – where the focus on police services is to respond quickly with tactics, while for other agencies some further planning may be required to organise the response activity. The crime prediction framework points towards the theoretical principles that explain these spatial patterns and the types of service response that would be most suitable. To identify and predict spatial patterns of crime, different techniques are required, using different types of input data. The frameworks also point towards the input data that are required and the spatial analysis techniques that are most suitable for each time frame. The frameworks are also sensitive to what the current research has distinguished as broadly two groupings of crime – where the commission of crime can, in part, be explained by foraging and boost behaviour principles, and where these principles do not apply. It is hoped, therefore, that with wide promotion, the adoption of the crime prediction framework will further improve how police and public safety agencies respond to and reduce crime.

12.8. Contributions to the field

This research offers many technical and practical contributions to the study of geographical patterns of crime and spatial crime prediction. In addition, the research suggests a new contribution to environmental criminology theory and offers a comprehensive benchmark analysis of hotspot mapping techniques against which other spatial crime prediction techniques can be compared. Specific contributions that the research offers include:

- The introduction of new techniques for measuring spatial crime prediction. These are the Prediction Accuracy Index, the illustration of the area under the curve measure applied to spatial crime analysis and the introduction of the Crime Prediction Index
- A template for capturing the key metrics for documenting the spatial prediction performance of a crime mapping output
- Illustration of the use of the Nearest Neighbour Index for determining if clustering is present in the crime data being examined and whether these data can be considered for hotspot analysis
- The first comprehensive metric examination of the commonly used hotspot mapping techniques for their ability to predict spatial patterns of crime – kernel density estimation, grid thematic mapping, thematic mapping of administrative areas, and standard deviation spatial ellipses
- The first comprehensive metric examination of the influence that cell size and bandwidth size has on KDE hotspot mapping output and this output's spatial crime prediction performance
- The first comprehensive metric examination of the Gi* statistic for producing spatial predictions of crime. This included comparing the standard approach for determining statistical significance thresholds to the Bonferroni corrected approach, examining the

influence of cell size and lag distance, and illustrating how the Gi* statistic helps remove much of the ambiguity in defining areas that are hotspots

- The first rigorous assessment of how the temporal stability of hotspot mapping output compares in relation to the temporal period of input data
- The first critical assessment of prospective mapping for generating spatial crime predictions for different periods of the future and how this compares to Gi* hotspot mapping
- The first comprehensive assessment of how GWR can be applied to crime data. While previous research has used GWR on crime data, little to date has been documented using crime data to illustrate the rigorous diagnostic statistical processes that are applied, the treatment that is often required to make crime data suitable for GWR modelling, and the interpretation of the results. The research was the first that has examined the differences in the results between using a Poisson GWR approach and a Gaussian GWR approach to crime data. In addition, this research was novel in examining both the hypothesis and exploratory approaches to GWR modelling and how GWR results could inform long-term spatial crime prediction
- The watering hole principle as an advancement to existing environmental criminology theory to help explain why crimes tend to form clusters. This has included describing how optimal foraging theory and the boost account theory are often not sufficient on their own for explaining where crimes such as street crimes against the person are likely to be committed in the immediate future
- The concept of a temporal framework for spatial crime prediction (the crime prediction framework), recognising that a single technique is unlikely to be suitable for accurately predicting all spatial patterns of crime and the importance of using different input data for different prediction time frames. The crime prediction framework also helps indicate why certain patterns of crime are likely to occur and is designed to help practitioners recognise the different response roles for countering the spatial crime patterns that can be accurately predicted.

12.9. Possible new areas of research

The results from this research, while comprehensive, are still limited by the testing of hypotheses in only one or two study areas and using a sub-set of crime types. However, analysis of different crime types has helped show where there are consistencies and differences in the results, with these then being explained by drawing on key theoretical principles to assist their interpretation. Analysis for two study areas, with differences in geography and the spatial distribution of crime has also allowed for a greater level of robustness in the findings and the practical and policy impacts of these results. That said, further research that repeats the methods used in this research using data for other areas, for the same and other crime types and for different time periods would of course be of value.

The research into KDE focused on using a fixed bandwidth approach. An alternative approach using an adaptive approach was not applied. The adaptive approach involves applying a bandwidth size that varies across the study area, determined by an upper limit in the number of neighbours that are used in any KDE calculation. A focus on only a fixed bandwidth approach in the current research calls for experiments similar to those used in this research to test the spatial crime prediction performance of an adaptive bandwidth approach to KDE, and how the limit on neighbours influences the results. A weighted approach to KDE can also be applied. The weighted approach involves weighting the input data by some attribute. For example, this attribute could include weighting recent events more heavily in the calculation of KDE values. Experiments that examine and test different weights for different attributes and how they influence spatial crime predictions may also be a worthwhile area for future research.

The Gi* statistic compares local averages to global averages, and in the experiments conducted in this PhD research, all areas were considered in the study area for the Gi* calculations. For certain areas in the study area, the commission of certain types of crime is not possible. For example, burglaries cannot be committed in lakes or on areas of open land. Cells representing areas where the commission of crime is not possible could be removed from the Gi* calculation process to produce a more accurate measure that compares local averages for where crime is possible, to the global average where crime is possible. This refinement may improve the prediction performance of Gi* hotspot analysis.

Throughout, a Bonferroni correction procedure was applied to the Gi* statistic results. While concerns had been expressed about the highly conservative nature of the Bonferroni correction approach in an earlier section (section 8.3.2), the approach produced results where it was easy to discern hotspots of crime. Indeed, the conservative determination of statistical significance thresholds from applying the Bonferroni correction appeared to offer an advantage for spatial crime analysis because it produced results that led to identifying small, specific areas that were extremely accurate in predicting where crime was likely to occur. However, further research may involve experimenting with other types of correction procedures to examine whether these alternative correction approaches produce spatial predictions of crime that are better than the Bonferroni corrected Gi* results.

The examination, results and discussion of GWR applied to crime data also introduced Bayesian spatially varying coefficient modelling as an alternative approach to spatial regression. The research literature suggests that Bayesian SVC modelling has several advantages over GWR. However, Bayesian SVC modelling involves significantly more computer processing in comparison to GWR and at present the availability of the software is limited to certain specialist products. To date, no application of Bayesian SVC modelling of crime data has been documented, but again, opportunities for exploring the application of Bayesian SVC modelling to crime data would be of value.

The focus of the current research has been on examining crime hotspots – the locations where crime concentrates. However, the examination of crime *coldspots* – locations where there is a significant absence of crime – is also of interest. The identification and analysis of coldspots could, for example, help to better understand the characteristics that are present that lead to the absence of crime. Research into crime coldspots may then help identify the conditions that practitioners may then want to promote in helping to prevent crime in areas where crime hotspots exist.

Finally, the introduction of the watering hole principle, although supported by the empirical observations from this research, requires further testing. Introducing the watering hole principle to a wide range of academics for their review and the further testing of the principles against crime data would be of benefit. For example, this could include the specific examination of the watering hole principle to crime data, similar to Bowers and Johnson's (2004) test of the boost account theory.

12.10. Research summary and conclusions

Hotspot analysis involves identifying areas where there is a high concentration of crime, relative to the distribution of crime across the study area of interest. Hotspot analysis has become a common feature of police and community safety analysis to assist in the targeting of resources. In this sense, hotspot mapping is perhaps the most basic form of spatial crime prediction: using crime data from the past to predict where crimes are likely to occur in the future. In recent years, new mapping techniques have emerged under the developments of predictive policing, with the claim of offering more accurate spatial predictions of crime. This PhD research has involved a comprehensive metric examination of hotspot analysis to investigate the extent to which hotspot analysis techniques accurately predict spatial patterns of crime and set a benchmark measure for comparison to these new predictive mapping techniques. The current research has also developed new technical and theoretical knowledge to support improvements in the geographical analysis of crime and spatial crime prediction.

The findings from this PhD research have shown the high degree of stability that exists in spatial patterns of crime. While previous research has also illustrated similar results, the findings from the current research have extended these previous results in two ways: examining the influence of different temporal periods of retrospective crime data on the stability of spatial crime patterns; and examining when crimes that occur in the future reach a stage of spatial stability.

The findings from this PhD research show that concentrations of crime that are likely to exist in the future can be determined from relatively short retrospective periods of recorded crime data. This clustering in retrospective crime data can be identified using the Nearest Neighbour Index, after which a hotspot mapping technique can be used to determine where these clusters are evident. A number of techniques are available for producing hotspot maps of crime. The techniques that have been commonly used in practice include spatial ellipses, thematic mapping of geographic administrative units, thematic mapping of grids and kernel density estimation. The current research has shown that of these techniques, KDE is the one that performs best for determining where crime is likely to occur in the future.

To compare the performance of different mapping techniques for predicting spatial patterns of crime, this PhD research has developed a new set of measures for spatial crime prediction. The first of these was the Prediction Accuracy Index, a simple global measure that provides an easy to calculate means of comparing the spatial prediction performance of mapping techniques. During the time this PhD research has developed, the use of the accuracy concentration curve was illustrated by Johnson et al. (2008b), and was shown to measure spatial prediction accuracy in more detail by comparing the performance of a mapping technique across the full analytical range of spatial scales. This PhD research has added to the measure provided by the accuracy concentration curve by borrowing from other fields of scientific study, in particular clinical trials, and introduced the area under the curve and the Crime Prediction Index to quantify spatial crime prediction performance more clearly.

While KDE may be the best of the common hotspot mapping techniques for predicting spatial patterns of crime, practice shows that the hotspot maps generated using this technique can significantly vary visually when different settings are used for the two main technical input parameters – the cell size and bandwidth size. To date, analysts and researchers have been offered little guidance on suitable cell size and bandwidth values for crime analysis, and no advice on how these parameters influence the predictive accuracy of the KDE mapping output that is generated. This research has shown that while cell size has little influence on the prediction performance of KDE hotspot maps, small cell sizes produce maps with greater visual appeal. Bandwidth size does though have an influence, with small sizes improving the prediction performance of the KDE hotspot map.

KDE hotspot mapping, though, is not without its weaknesses, the most problematic being the difficulty in determining the areas that are hot. This PhD research has examined the use of the Gi* statistic for identifying areas where the spatial concentration of crime is statistically significant. In-so-doing, this has helped remove much of the ambiguity in determining hotspot areas and has shown that this technique consistently performs better than KDE for predicting where crime is likely to occur. The Bonferroni correction measure has also been used to further improve the determination of these hotspot areas by helping to address the issues of multiple testing.

As the interest in predictive crime mapping has developed over the last few years, little attention has been given to the retrospective period of crime data to use to inform these predictions, and whether these predictions are consistently accurate for all periods of the future. In addition, little attention has also been given to whether these spatial crime predictions are stronger for some crime types than others, whether certain techniques are more suitable for certain crime types, and why. If the reason why crime is likely to occur in a certain place at certain time in the future cannot be clearly explained, this then limits its practical use and the determination of the types of response that would be most suitable for addressing the crime that is predicted to occur. This is seen to be a weakness in many commercially driven predictive mapping techniques which use data for multiple retrospective periods, on multiple crime types, expecting a single algorithm to produce accurate predictions for periods of the future.

Prospective mapping stands aside from these new predictive policing techniques because it is built directly on the frequently observed patterns of repeat and near repeat victimisation, and the theoretical principles of optimal foraging theory, boost account theory and crime pattern theory. The prospective mapping technique is sensitive to how offenders tend to operate in spates, and has become a popular technique of choice for predicting spatial patterns of burglary. When comparing, in the current research, the output from prospective mapping of burglary to Gi* hotspot analysis, prospective mapping performed better for predicting where burglaries were likely to occur in the immediate future (within the next week). However, for street crimes against the person such as theft and assaults, Gi* hotspot analysis more accurately predicted where these types of crimes were most likely to occur. The reason for Gi* hotspot analysis performing better in predicting spatial patterns of crime was considered to be due to the offending behavioural concepts of foraging and the boost account being less relevant for theft from the person and assault offences. Instead, the spatial clustering of theft from the person and assaults was considered to be more related to occurring in areas where the geographic conditions for this type of crime are highly favourable. In addition, the findings showed that the predictions generated using prospective mapping were less accurate than those generated using Gi* when the time frame was for the near future (beyond the next week). Collectively, these findings that compared the results for Gi* hotspot analysis to prospective mapping, for different crime types, for different retrospective periods of crime data and for predicting crime for different time periods of the future have resulted in the introduction of the watering hole principle to help theoretically explain why crime is attracted to areas where favourable geographic conditions exist. While crime pattern theory, the routine activity approach and the rational choice perspective help to explain why spatial patterns of crime concentrate in places, the watering hole principle helps further explain where these concentrations are most likely to emerge and why. As a result, in many cases a combination of prospective mapping and hotspot mapping should be used for predicting where crime is likely to happen in the near future. This combined approach that uses these two spatial analysis techniques is captured in the crime prediction framework. The crime prediction framework shows when and why hotspot analysis and/or prospective mapping should be used, the input data required, the theoretical reasoning that explains the spatial patterns that are likely to occur, and the types of response activity most suitable.

The original ambition for examining the use of geographically weighted regression was to determine at the local level the reasons why hotspots exist, and whether these reasons differed between hotspots. At present, the spatial precision of most of the data of interest for exploring spatial relationships with crime is not precise enough to enable these explanatory distinctions between hotspots. Additionally, the heavy smoothing nature of GWR does not permit analysis at the spatial scale that hotspots are identified. While techniques such as Bayesian spatially varying coefficient models appear to offer some potential in exploring these spatially precise relationships, GWR modelling still offers the ability to identify more macro spatially varying relationships between crime patterns and other variables. The findings from the current research have illustrated the use of GWR for identifying the variables that spatially vary in their relationship with the distribution of crime has introduced a third temporal prediction period to the crime prediction framework. This third temporal prediction period is used for informing the direction of strategic policy and making long-term predictions on how crime may change.

This PhD research has shown that extremely accurate predictions of where crime is likely to concentrate in the future can be best determined by using good hotspot analysis of where crimes have concentrated in the past. While prospective mapping continues to offer better predictions of where crime against property is likely to occur in the immediate future, good hotspot analysis using the Gi* statistic is better at identifying where street crimes against the person are likely to occur in the immediate future and for identifying where crimes are likely to form hotspots in the near future. The reasons for these hotspots can now be better explained with the improvement in the theory (the watering hole principle) on why crime is likely to form clusters, which in turn can better assist practitioners in determining how to address the favourable conditions that lead to these concentrations of crime. Spatial crime prediction is expected to continue to gain interest in the coming years. This PhD research offers a significant step forward in helping to inform practitioners about how spatial crime patterns can be predicted, contributes to developments in environmental criminology on why crime hotspots are likely to form, and offers a comprehensive metric examination into the technical application of hotspot mapping for improving the geographical analysis of crime.

13.References

Anselin, L. (1988a), *Spatial econometrics: methods and models*. Dordrecht: Kluwer.

Anselin, L. (1988b), "Lagrange Multiplier test diagnostics for spatial dependence and spatial heterogeneity", *Geographical Analysis* 20:1–17.

Anselin, L. (1995), "Local indicators of spatial association – LISA", *Geographical Analysis* 27:93-115.

Anselin, L., and Griffith, D.A. (1988), "Do Spatial Effects Really Matter in Regression Analysis?", *Papers, Regional Science Association* 65:11-34.

Ashton, J., Brown, B. Snr, and Pease, K. (1998), "Repeat victimisation: offender accounts", *International Journal of Risk, Security and Crime Prevention* 3:4:269-279.

Astun Technology (2013), *Vigilance Modeller*. https://astuntechnology.com/labs/vigilance-modeller [Accessed 22nd January 2014].

Atkinson, P., and Molesworth, A. (2000), "Geographical analysis of communicable disease data", in P. Elliot, J. Wakefield, N. Best and D. Briggs (Eds) *Spatial epidemiology: methods and applications*, Oxford: Oxford University Press.

Austin, M.P.; Meyers, J.A. (1996), "Current approaches to modelling the environmental niche of eucalypts: implication for management of forest biodiversity", *Forest Ecology and Management* 85:95-106.

Babor, T., Caetano, R., Casswell, S., Edwards, G., Giesbrecht, N., Graham, K., Grube, J., Gruenewald, P., Hill, L., Holder, H., Homel, R., Osterberg, E., Rehm, J., Room, R., and Rossow, I. (2003), *Alcohol : No ordinary commodity. Research and public policy*, Oxford: Oxford University Press.

Bailey, T.C., and Gatrell, A.C. (1995), *Interactive Spatial Data Analysis*, Reading, MA: Addison-Wesley.

Beato, C. (2008), *Comprehendo e Avaliando: Projetos de Segurança Publica*, Belo Horizonte: Editora UFMG.

Bini, L.M., Diniz-Filho, J.A.F., Rangel, T.F.L.V.B., Akre, T.S.B., Albadalejo, R.G., Alburquerque, F.S., Aparicio, A., Araújo, M.B., Baselga, A., Beck, J., Belloq, M.I., Böhning-Gaese, K., Borges, P.A.V., Castro-Parga, I., Chey, V.K., Chown, S.L., de Marco Jr., P., Dobkin, D.S., Ferrer-Castán, D., Field, R., Filloy, J., Fleishman, E., Gómez, J.F., Hortal, J., Kerr, J.T., Kitching, I.J., León-Cortés, J.L., Lobo, J.M., Montoya, D., Morales-Castillo, I., Moreno, J.C., Olalla-Tárraga, M.A., Pausas, J.G., Qian, H., Rahbeck, C., Rodriguez, M.A., Sanders, N.J., and Williams, P. (2009), "Coefficient shifts in geographical ecology: an empirical evaluation of spatial and non-spatial regression", *Ecography*, 32:1-12.

Bithell, J.F. (1990), "An Application of Density Estimation to Geographical Epidemiology", *Statistics in Medicine* 9:691–701.

Blau, P. (1977), *Inequality and Heterogeneity*, New York, NT: Free Press.

Block, R., and Block, R.B. (2000), "The Bronx and Chicago – Street Robbery and the Environs of Rapid Transit Stations", in V. Goldsmith, P.G. McGuire, J.H. Mollenkopf, and T.A. Ross (Eds.) *Analysing Crime Patterns: Frontiers and Practice*, Thousand Oaks, CA: Sage.

Block, R., and Perry, S. (1993), *STAC News Vol. 1, No. 1*, Chicago: Illinois Criminal Justice Information Authority - Statistical Analysis Center.

Bottoms A. E. and Wiles P. (2002), "Environmental criminology", in M. Maguire, R. Morgan, R. Reiner (eds), *The Oxford Handbook of Criminology 3rd edition*, Oxford: Oxford University Press.

Bowers, K.J., and Hirschfield, A. (1999), "Exploring Links Between Crime and Disadavantage in North-West England – An Analysis Using Geographical Information Systems", *International Journal of Geographic Information Science* 13: 2:159-184.

Bowers K.J., and Johnson S.D. (2004), “Who commits near repeats? A test of the boost explanation”, *Western Criminology Review* 5:3.

Bowers, K.J., Johnson, S., and Pease, K. (2004), “Prospective Hotspotting: The Future of Crime Mapping?”, *British Journal of Criminology* 44:5:641-658.

Bowers, K.J., Newton, M., and Nutter, R. (2001), “A GIS-linked Database for Monitoring Repeat Domestic Burglary”, in A. Hirschfield and K. Bowers (Eds.) *Mapping and Analysing Crime Data – Lessons from Research and Practice*, London: Taylor and Francis.

Bowman, A.W., and Azzalini, A. (1997), *Applied Smoothing Techniques for Data Analysis: the Kernel Approach with S-Plus Illustrations*, Oxford: Oxford University Press.

Braga, A. (2007), *Effects of Hot Spots Policing on Crime: A Campbell Collaboration Systematic Review*. http://www.aic.gov.au/campbellcj/reviews/titles.html [Accessed 22nd January 2014].

Brantingham, P. J. and Brantingham, P. L. (1981), *Environmental Criminology*, London: Sage.

Brantingham, P.J., and Brantingham, P.L. (1984), *Patterns in crime*, New York: Macmillan.

Brantingham, P., and Brantingham, P. (1995), “Criminality of place: Crime generators and crime attractors”, *European Journal of Criminal Policy and Research* 3:5-26.

Brunsdon, C. (1995), “Estimating probability surfaces for geographical point data: An adaptive kernel algorithm”, *Computers and Geosciences* 21:877–894.

Brunsdon C, and Charlton M, (2011), "An assessment of the effectiveness of multiple hypothesis testing for geographical anomaly detection", *Environment and Planning B: Planning and Design* 38(2) 216 – 230.

Burgess, E. W. (1916), “Juvenile delinquency in a small city”, *Journal of the American Institute of Criminal Law and Criminology* 6:724-728.

Burgess, E.W. (1925), “The growth of the city: An introduction to a research project”, in R.E. Park, E.W. Burgess and R.D. McKenzie (Eds.), *The City*, Chicago: University of Chicago Press pp. 47-62.

Caldas de Castro, M., and Singer, B.H. (2006), “Controlling the False Discovery Rate: A New Application to Account for Multiple and Dependent Test in Local Statistics of Spatial Association”, *Geographical Analysis* 38:180-208.

Cahill, M., and Mulligan, G. (2007), “Using geographically weighted regression to explore local crime patterns”, *Social Science Computer Review* 25:2:174-193.

Can, A., (1990), “The measurement of neighborhood dynamics in urban house prices”, *Economic Geography* 66:254–272.

Chainey, S.P. (2001), “Combating Crime Through Partnership: Examples of Crime and Disorder Mapping Solutions in London, UK”, in A. Hirschfield. and K. Bowers, K. (Eds.) *Mapping and Analysing Crime Data – Lessons from Research and Practice*, London: Taylor and Francis.

Chainey, S.P. (2011), *Identifying hotspots: an assessment of common techniques*, Presentation at the International Crime and Intelligence Analysis Conference November 2011, Manchester, England.

Chainey, S.P. (2012a), “Improving the explanatory content of analysis products using hypothesis testing”, *Policing: A Journal of Policy and Practice* 6:2: 108-121.

Chainey, S.P. (2012b). *JDI Brief: Predictive mapping (predictive policing)*, London: UCL.

Chainey, S.P. (2013), “Examining the influence of cell size and bandwidth size on kernel density estimation crime hotspot maps for predicting spatial patterns of crime”, *Bulletin of the Geographical Society of Liege* 60:7-19.

Chainey, S.P., and Chapman, J. (2013), "A problem-oriented approach to the production of Strategic Intelligence Assessments", *Policing: An International Journal of Strategies and Management* 36:3:474-490.

Chainey, S.P., and Desyllas, J. (2008), "Modelling pedestrian movement to measure on-street crime risk", in L. Liu, and J. Eck (Eds.) *Artificial Crime Analysis Systems: Using Computer Simulations and Geographic Information Systems*, Hershey, PA: Information Science Reference.

Chainey, S.P. and Ratcliffe, J.H. (2005), *GIS and Crime Mapping*, London: Wiley.

Chainey, S.P., Reid, S., and Stuart, N. (2002), "When is a Hotspot a Hotspot? A Procedure for Creating Statistically Robust Hotspot Maps of Crime", in G. Higgs (Ed.), *Innovations in GIS 9 Socio-economic Applications of Geographic Information Science*, London: Taylor and Francis.

Chainey, S.P., Tompson, L., Uhlig,S. (2008a), "The utility of hotspot mapping for predicting spatial patterns of crime", *Security Journal* 21:1-2.

Chainey,S.P., Tompson,L., and Uhlig,S. (2008b), "*Response to Pezzuchi"*, *Security Journal* 21:4.

Chainey,S.P., Tompson,L., and Uhlig,S. (2008c), "Response to Levine", *Security Journal* 21:4.

Clark, P.J., and Evans, F.C (1954), "Distance to nearest neighbour as a measure of spatial relationships in populations", *Ecology* 35:445-453.

Clarke, R.V. (1999), *Hot Products: understanding, anticipating and reducing demand for stolen goods*, Police Research Group: Police Research Series, Paper 112, 48.

Clarke, R.V., and Eck, J. (2003) *Become a Problem-Solving Crime Analyst in 55 Small Steps*, London: University College London.

Clarke, R.V., and Felson, M. (1993), "Introduction: Criminology, routine activity, and rational choice" in R.V. Clarke and M. Felson (Eds.), *Routine Activity and Rational Choice*, New Brunswick, NJ: Transaction Publishers, pp. 259-294.

Clarke, R.V.G., and Goldstein, H. (2003), "Thefts from Cars in Centre-City Parking Facilities: A Case Study in Implementing Problem-Oriented Policing" in Knutsson, J. (Ed.), *Problem-Oriented Policing: From Innovation to Mainstream*, Crime Prevention Studies, Monsey, NY: Criminal Justice Press, pp. 257–298.

Cohen, L.E., and Felson, M. (1979), "Social change and crime rate trends: A Routine Activity Approach", *American Sociological Review* 44:588-608.

Cornish, D., and Clarke, R. (1986), *The Reasoning Criminal: Rational Choice Perspectives on Offending*, New York, NY: Springer-Verlag.

Costello, A., and Wiles, P. (2001), "GIS and the journey to crime", in A. Hirschfield and K. Bowers (Eds.), *Mapping and Analysing Crime Data*, London: Taylor and Francis, pp. 27-60.

Cromwell, P., Olson, J.N., and Avary, D. A. W. (1999), "Decision strategies of residential burglars", in P. Cromwell (Ed.), *In their own words: Criminals on crime*, Los Angeles, CA: Roxbury, pp. 50-56.

Crowe, T.D. (2000), *Crime prevention through environmental design – application of architectural design and space management concepts* (2nd edition), Woburn, MA: Butterworth-Heinemann.

de Smith, M.J., Goodchild, M.F., and Longley, P.A. (2007), *Geospatial Analysis: A Comprehensive Guide to Principles, Techniques and Software Tools*, Leicester: Winchelsea Press.

Dent, B. (1999), *Cartography – Thematic Map Design*, London: McGraw-Hill.

Dolan, M. (2000), "Violence risk prediction: Clinical and actuarial measures and the role of the Psychopathy Checklist", *The British Journal of Psychiatry* 177:303-311.

Dubin, R.A. (1992), “Spatial autocorrelation and neighbourhood quality”, *Regional Science and Urban Economics* 22:433–452.

Dubin, R.A. (1998), “Spatial autocorrelation: a primer”, *Journal of Housing Economics* 7:304–327.

Eck, J. (1995), “A general model of the geography of illicit retail marketplaces”, in D. Weisburd and J.E. Eck (Eds.), *Crime and place*, Monsey, NY: Criminal Justice Press, pp. 67-93.

Eck, J.E., Chainey, S.P., Cameron, J.G., Leitner, M. and Wilson, R.E. (2005), *Mapping Crime: Understanding Hot Spots*, Washington DC: National Institute of Justice.

Elliot, P., Wakefield, J., Best, N., and Briggs, D. (2000), *Spatial epidemiology: methods and applications,* Oxford: Oxford University Press.

Ericsson, U. (1995), “Straight from the Horse's Mouth”, *Forensic Update* 43:23-25.

ESRI(UK) (2013) *Crime Analyst* http://www.esriuk.com/software/arcgis/crimeanalyst [Accessed 22nd January 2014].

Farrell, G., and Pease, K. (1993), *Once bitten, twice bitten: repeat victimisation and its implications for crime prevention*, Crime Prevention Unit Paper 46. London: Home Office.

Fischer, M.M. (1997), “Computational Neural Networks: A New Paradigm for Spatial Analysis”, *Environment and Planning A* 29:1873-1891.

Fielding M., and Jones V. (2012), “Disrupting the optimal forager: predictive risk mapping and domestic burglary reduction in Trafford, Greater Manchester”, *International Journal of Police Science and Management* 14:1:30-41.

Felson, M. (1987), “Routine activities and crime prevention in the developing metropolis”, *Criminology* 25:911-932.

Felson, M. (1995), "Those who discourage crime", in D. Weisburd & J.E. Eck (Eds.), *Crime and place*, Monsey, NY: Criminal Justice Press, pp. 53-66.

Felson, M. (1998), *Crime and everyday life: Impact and implications for society*, Thousand Oaks, CA: Pine Forge Press.

Foody, G.M. (2003), "Geographical weighting as a further refinement to regression modelling: an example focussed on the NDVI-rainfall relationship", *Remote Sensing of Environment* 88:3:283-293.

Fotheringham, A.S., Brunsdon, C. and Charlton, M. (2002), *Geographically weighted regression: the analysis of spatially varying relationships*, Chichester: Wiley.

Fox, J.C., Ades, P.K., Bi, H. (2001), "Stochastic structure and individual-tree growth models", *Forestry Ecology Management* 154:261-276.

Gao, X., Asami, Y., and Chung, C-J.F. (2006), "An Empirical Evaluation of Spatial Regression Models", *Computers and Geosciences* 32:1040-1051.

Getis, A., and Aldstadt, J. (2004), "Constructing the spatial weights matrix using a local statistic", *Geographical Analysis* 35:2:90-104.

Getis A., and Ord J.K. (1996), "Local spatial statistics: an overview", in P. Longley and M. Batty (Eds.) *Spatial analysis: modelling in a GIS environment*, Cambridge: GeoInformation International.

Gibbs, J.P., and Martin, W.T. (1962), "Urbanization, technology and the division of labor", *American Sociological Review* 27:667–677.

Goldsmith, V., McGuire, P.G., Mollenkopf, J.H. and Ross, T.A. (2000), *Analysing Crime Patterns: Frontiers of Practice*, New York: Altamira Press.

Goldstein, H. (1979), "Improving policing: a problem-oriented approach", *Crime and Delinquency* 25:236-258.

Goldstein, H. (1990), *Problem-oriented policing*, New York, NY: McGraw-Hill.

Graham, K., and Homel, R. (2008), *Raising the bar: preventing aggression in an around bars, clubs and pubs,* Cullompton: Willan Publishing.

Gray, P. (2000), "Repeat Victimisation in the Asian Community: A Study of Domestic Burglary", *Crime Prevention and Community Safety* 2:53–65.

Griffith, D.A. (2008), "Spatial-filtering-based contributions to a critique of geographically weighted regression (GWR)", *Environment and Planning A* 40:2751-2769.

Groff, E.R., Weisburd, D. and Morris, N. (2008), "Where the Action is at Places: Examining Spatio-Temporal Patterns of Juvenile Crime at Places Using Trajectory Analysis and GIS", in D. Weisburd, W. Bernasco and G. Bruinsma (Eds.) *Putting Crime in it's Place: Units of Analysis in Spatial Crime Research*, New York, NY: Springer.

Guerry, A.-M. (1833), *Essai sur la statistique morale de la France: precede d'un rapport a l'Academie de sciences*, Paris: Chez Crochard.

Guisan A., Edwards, T.C., and Hastie, T. (2002), "Generalized linear and generalized additive models in studies of species distribution: setting the scene", *Ecological Modelling* 157:89-100.

Haberman, C.P., and Ratcliffe, J.H. (2012), "The predictive policing challenges of near repeat armed street robberies", *Policing: A Journal of Policy and Practice* 6:2:151-166.

Haiganoush, P.K., Rappaport, N.G., and Wood, D.L. (1997), "Regression methods for spatially correlated data: an example using beetle attacks in a seed orchard", *Forest Science* 43:1:71-77.

Haining, R. (1990), *Spatial data analysis in the social and environmental sciences*, Cambridge: Cambridge University Press.

Hanley, J.A., and McNeil, B.J. (1982), "The Meaning and Use of the Area under a Receiver Operating Characteristic (ROC) Curve", *Radiology* 143:1: 29–36.

Harries, K. (1999), *Mapping Crime: Principle and Practice*, Washington DC: United States National Institute of Justice.

Harris, P., Fotheringham, S., Crespo, R., and Charlton, M (2010), "The use of geographically weighted regression for spatial prediction: an evaluation of models using simulated data sets", *Mathematical Geosciences* 42:6:657-680.

Hart, T.C., and P.A. Zandbergen, (2012), *Effects of data quality on predictive hotspot mapping*, Washington DC: United States National Institute of Justice.

Hastie, T.J., and Tibshirani, R.J., (1990), *Generalized Additive Models*, London: Chapman and Hall.

Ching-Chi H. and Pugh, M.D. (1993), "Poverty, Income Inequality, and Violent Crime: A Meta-Analysis of Recent Aggregate Data Studies", *Criminal Justice Review* 18:2:182-202.

Hirschfield, A., Bowers K.J., and P.J.B. Brown (1995), "Exploring relations between crime and disadvantage on Merseyside", *European Journal on Criminal Policy and Research* 3:3:93-112.

Home Office (2001), *Crime Reduction Toolkits Focus Area and Hotspots*, London: Home Office Crime Reduction Unit.

Home Office (2005) *Crime Mapping: Improving Performance. A Good Practice Guide for Front Line Officers*, London: Home Office.

Hough, M. and Tilley, N. (1998), *Getting the Grease to the Squeak: Research Lessons for Crime Prevention*, Crime Prevention and Detection Paper 85. London: Home Office.

IBM (2010a), *Public Safety: From 'Sense and Respond' to 'Predict and Act,* Somers, NY: IBM Corporation.

IBM (2010b), *City of Lancaster takes a predictive approach to policing: enhancing operational planning with deep insight into crime data*, Somers, NY: IBM Corporation.

Illinois Criminal Justice Information Authority (1996), *Spatial and Temporal Analysis of Crime*, Chicago: State of Illinois.

Institute of Alcohol Studies (2013) *Crime and social impacts of alcohol factsheet*. http://www.ias.org.uk/uploads/pdf/Factsheets/Crime%20FS%20HM%20May%202013.pdf [Accessed 22nd January 2014].

International Association of Crime Analysts (IACA, 2013). *Predictive Policing Forum*, International Association of Crime Analysts Training Conference 9-13 September 2013, Fort Lauderdale, Miami.

Jefferis, E. (1999), *A Multi-Method Exploration of Crime Hot-spots: A Summary of Findings*. Crime Mapping Research Centre Intramural project, Washington, DC: National Institute of Justice.

Jeffery, C.R. (1971), *Crime Prevention Through Environmental Design*, Beverly Hills, CA: Sage Publications.

Johnson, S.D., and Bowers, K.J. (2004a), "The burglary as clue to the future: the beginnings of prospective hotspotting", *European Journal of Criminology* 1:2:237-255.

Johnson S.D., and Bowers K.J. (2004b), "The stability of space-time clusters of burglary", *British Journal of Criminology* 44:1:55-65.

Johnson, S.D., Lab, S.P., and Bowers, K.J. (2008a), "Stable and Fluid Hotspots of Crime: Differentiation and Identification", *Built Environment* 34:1:32-45.

Johnson S., Bowers K.J., Birks D., and Pease K. (2008b), "Predictive Mapping: Accuracy for Different Units of Analysis and the role of the Environmental Backcloth", in D. Weisburd, W. Bernasco and G. Bruinsma (Eds.) *Putting Crime in it's Place: Units of Analysis in Spatial Crime Research,* New York, NY: Springer.

Johnson, S.D., Summers, L., and Pease, K. (2009). “Offender as Forager? A Direct Test of the Boost Account of Victimization”, *Journal of Quantitative Criminology* 25:181-200.

Johnson, S.D., Bowers, K.J., and Pease, K. (2012), “Towards the Modest Predictability of Daily Burglary Counts”, *Policing: A Journal of Policy and Practice* 6:2:167–176.

Kennedy, D.M. (2009), *Deterrence and Crime Prevention: Reconsidering the prospect of sanction*, New York, NY: Routledge.

Kennedy, L.W., Caplan, J.M., and Piza, E. (2011), “Risk clusters, hotspots, and spatial intelligence: Risk Terrain Modeling as an Algorithm for Police Resource Allocation Strategies”, *Journal of Quantitative Criminology* 27:3:339-362.

Knox, P.L. (1994), *Urbanization: an introduction to urban geography*, Englewood Cliffs, NJ: Prentice-Hall.

Langworthy, R.H., and Jefferis, E. (2000), “The Utility of Standard Deviation Ellipses for Evaluating Hotspots”, in V. Goldsmith, P.G. McGuire, J.H. Mollenkopf, and T.A. Ross (Eds.) *Analysing Crime Patterns: Frontiers and Practice*, Thousand Oaks, CA: Sage.

LaVigne, N., and Wartell, J. (Eds.) (1998), *Crime Mapping Case Studies: Successes in the Field, Vol. 1,* Washington, DC: Police Executive Research Forum.

LaVigne, N., and Wartell, J. (Eds.) (1999), *Crime Mapping Case Studies: Successes in the Field, Vol. 2,* Washington, DC: Police Executive Research Forum.

LeBeau, J.L. (1987), “The methods and measures of centrography and the spatial dynamics of rape”, *Journal of Quantitative Criminology* 3:125-141.

LeBeau, J.L. (1992), “Four case studies illustrating the spatial-temporal analysis of serial rapists”, *Police Studies* 15:124-145.

LeBeau, J.L. (2001), "Mapping Out Hazardous Space for Police Work", in A. Hirschfield and K. Bowers, *Mapping and Analysing Crime Data – Lessons from Research and Practice*, London: Taylor and Francis.

Legendre, P. (1993), "Spatial autocorrelation: trouble or new paradigm?", *Ecology* 74:6:1659-1673.

LeSage, J.P. (1997), "Bayesian Estimation of Spatial Autoregressive Models", *International Regional Science Review* 20:1 and 2:113-129.

Levine, N. (2008), "The Hottest Part of a Hotspot: comments on The Utility of Hotspot Mapping for Predicting Spatial Patterns of Crime", *Security Journal* 21:4:295-302.

Levine, N. (2010), *CrimeStat III: A Spatial Statistics Program for the Analysis of Crime Incident Locations*, Houston, TX: Ned Levine and Associates, Washington, DC: National Institute of Justice.

Littell, R.C., Milliken G.A., Stroup W.W., and Wolfinger, R.D. (1996), *SAS system for mixed models*, Cary, N.C: SAS Institute Inc.

Loftin C., and Ward S.K. (1983), "A spatial autocorrelation model of the effects of population density on fertility", *American Journal of Sociology* 48:1:121-28.

Maguire, M., and Hopkins, M. (2003), "Data analysis for problem-solving: alcohol and city centre violence", in K. Bullock and N. Tilley (Eds.) *Crime Reduction and Problem-Oriented Policing*, Cullompton: Willan Publishing.

Martin, D., Barnes, E., and Britt, D. (1998), "The Multiple Impacts of Mapping it Out; Police, Geographic Information Systems (GIS) and Community Mobilization During Devil's Night in Detroit, Michigan", in N. La Vigne and J. Wartell (Eds.) *Crime Mapping Case Studies: Successes in the Field*, Washington, DC: Police Executive Research Forum.

Mashford, T. (2008), "Methods for implementing crime mapping within a large law enforcement agency: experiences from Victoria, Australia", in S.P. Chainey and L.

Tompson (Eds.) *Crime Mapping Case Studies: Practice and Research*, Chichester: Wiley.

Mayhew, H. (1862), *London Labour and the London Poor*, London: Griffin Bohn.

McCreith, S. and Parkinson, S. (2004), *Crimes Against Students: Emerging Lessons For Reducing Student Victimisation*, London: Home Office.

McDonald, P.P. (2002), *Managing Police Operations: Implementing the New York Crime Model – CompStat*, Belmont, CA: Wadsworth.

McGuire, P.G., and Williamson, D. (1999), *Mapping Tools for Management and Accountability*, Paper presented to the Third International Crime Mapping Research Center Conference, Orlando, Florida, December 11-14, 1999.

Messner, S.F. and Anselin, L. (2004), Spatial analyses of homicide with areal data, in M.F.Goodchild and D.G. Janelle (eds) *Spatially integrated social science*, New York, NY: Oxford University Press.

Mohebbi M., Wolfe, R., Jolley, D., Forbes, A.B., Mahmoodi, M., and Burton, R.C. (2011), "The spatial distribution of esophageal and gastric cancer in Caspian region of Iran: an ecological analysis of diet and socio-economic influences", *International Journal of Health Geographics* 10:13.

Monmomier, M. (1996), *How to Lie with Maps*, Chicago: University of Chicago Press.

Mossman, D. (1994), "Assessing predictions of violence: Being accurate about accuracy", *Journal of Consulting and Clinical Psychology* 62:783 -792.

Nakaya T., Fotheringham, A.S., Charlton, M., and Brunsdon, C., (2005), "Geographically weighted Poisson regression for disease associative mapping", *Statistics in Medicine* 24:2695-2717.

Nakaya, T., Charlton, M., Lewis, P, Brunsdon, C., Yao, J., and Fotheringham, A.S. (2012), *GWR4 User Manual: Windows Application for Geographically Weighted Regression Modelling*.

Newman, G. (1997), "Introduction: Towards a theory of situational crime prevention", in G. Newman, R. Clarke and S.G. Shoham (Eds.), *Rational Choice and Situational Crime Prevention: Theoretical Foundations*, Dartmouth, UK: Ashgate, pp. 1-23.

Newman, O. (1972), *Defensible Space: Crime Prevention Through Urban Design*, New York: Macmillan.

Openshaw, S. (1984), *The Modifiable Areal Unit Problem*, Concepts and Techniques in Modern Geography 38. Norwich: Geobooks.

Ord, J.K., and Getis, A. (1995), "Local Spatial Autocorrelation Statistics: Distributional Issues and an Application", *Geographical Analysis* 27:4.286-306.

Osborne, D.A. and Wernicke, S.C. (2003), *Introduction to Crime Analysis: Basic Resources for Criminal Justice Practice*, New York: Haworth Press.

O'Sullivan, D., and D. J. Unwin (2003), *Geographic Information Analysis*, Hoboken, NJ: Wiley.

Pease, K. (1998), *Repeat victimization: taking stock*, Home Office Police Research Group, crime detection and prevention series, paper 90, London: Home Office.

Pezzuchi, G. (2008), "A brief commentary on The Utility of Hotspot Mapping for Predicting Spatial Patterns of Crime", *Security Journal* 21:4:291-292.

Polivi, N., Looman, T., Humphries, Ch., and Pease, K., (1991), "The Time-Course of Repeat Burglary Victimization", *British Journal of Criminology* 31:411-414.

PredPol (2013a), *Predict crime in real time*. Available at http://www.predpol.com/ [Accessed 22nd January 2014].

PredPol (2013b), *Policing meets big data*. Available at http://www.predpol.com/about/ [Accessed 22nd January 2014].

PredPol (2013c), *Scientifically proven field results.* Available at http://www.predpol.com/results/ [Accessed 22nd January 2014].

Quetelet, A. (1842), *A Treatise in Man,* Edinburgh: Chambers.

Ratcliffe, J.H. (2002), "Aoristic signatures and the temporal analysis of high volume crime patterns", *Quantitative Criminology* 18:23-43.

Ratcliffe, J. (2004) *HotSpot Detective for MapInfo Helpfile Version 2.0*, Philadelphia: JHR Systems.

Ratcliffe, J.H. (2008), *Intelligence-led Policing*, Cullompton: Willan Publishing.

Ratcliffe, J.H. (2010), "The spatial dependency of crime increase dispersion", *Security Journal* 23:18-36.

Ratcliffe, J.H., and McCullagh, M.J. (1999), "Hotbeds of Crime and the Search for Spatial Accuracy", *Journal of Geographical Systems* 1:4:385-398.

Ratcliffe, J.H., and McCullagh, M.J. (2001), "Crime, Repeat Victimisation and GIS", in K. Bowers and A. Hirschfield, *Mapping and Analysing Crime Data – Lessons from Research and Practice*, London: Taylor and Francis.

Ratcliffe, J.H., Taniguchi, T., Groff, E., and Wood, J. (2011), "The Philadelphia Foot Patrol Experiment: A randomized controlled trial of police patrol effectiveness in violent crime hotspots", *Criminology* 49:3:795-831.

Rengert, G. (1989). "Behavioural geography and criminal behaviour", in D.J. Evans and D.T. Herbert (Eds.), *The Geography of Crime*, London: Routledge, pp. 161-175.

Rengert, G. F. (1996), *The Geography of Illegal Drugs*, Boulder, CO: Westview Press.

Rengert, G.F., and Wasilchick, J. (1985). *Suburban burglary: A time and place for everything*, Springfield, IL: C.C. Thomas Publishing.

Rogerson, P.A. (2004), "The statistical significance of the maximum local statistic", in A. Getis, J. Mur and H.G. Zoller (Eds.), *Spatial Econometrics and Spatial Statistics*, Basingstoke: Palgrave McMillan.

Roncek, D.W. (1981), "Dangerous Places: Crime and Residential Environment", *Social Forces* 60:1:74-96.

Rosenshein, L., Scott, L., and Pratt, M. (2011), *Finding a meaningful model*, Redlands, CA: ESRI.

Rossmo, K. (2000), *Geographic profiling*, Boca Raton, FL: CRC Press.

Sampson, R.J., (2012), *Great American City: Chicago and the enduring neighbourhood effect*, Chicago: University of Chicago Press.

Sampson, R.J., Raudenbush, S.W., and Earls, F. (1997), "Neighborhoods and violent crime: a multilevel study of collective efficacy", *Science* 277:918-924.

Sawada, M. (1999), "ROOKCASE: An Excel 97/2000 Visual Basic (VB) Add-in for Exploring Global and Local Spatial Autocorrelation", *Bulletin of the Ecological Society of America* 80:4:231-234.

Schabenberger, O., and Pierce, F.J. (2002), *Contemporary Statistical Models for the Plant and Soil Sciences*, New York, NY: CRC Press.

Schick, W. (2004), "CompStat in the Los Angeles Police Department", *Police Chief* 71:1:17-23.

Shaw, C.R. and McKay, H.D. (1942), *Juvenile delinquency and urban areas,* Chicago: Chicago University Press.

Sherman, L. (2009), “Hot spots”, in Wakefield, A. and Fleming, J. *The SAGE Dictionary of Policing*, London: Sage Publications.

Sherman, L., Gartin, P. and Buerger, M.E. (1989), “Hot spots of predatory crime: Routine activities and the criminology of place”, *Criminology* 27:27–55.

Shin, M. and A.L. Goel (2000), “Empirical data modeling in software engineering using radial basis functions”, *IEEE transactions on software engineering* 26:6:1-10.

Silverman, B.W. (1986), *Density Estimation for Statistics and Data Analysis*, London: Chapman and Hall/CRC.

Tobler W., (1970), “A computer movie simulating urban growth in the Detroit region”, *Economic Geography* 46:2:234-240.

Townsley, M., Mann, M., and Garrett, K. (2011), “The Missing Link of Crime Analysis: A Systematic Approach to Testing Competing Hypotheses”, *Policing: A Journal of Policy and Practice* 5:2:158-171.

Voss, P.R., Long, D.D., Hammer, R.B., and Friedman, S (2006), “County child poverty rates in the US: a spatial regression approach”, *Population Research and Policy Review* 25:369-91.

Waller, L.A., Zhu, L., Gotway, C.A., Gorman, D.M., and Gruenewald, P.J. (2007), “Quantifying geographic variations in associations between alcohol distribution and violence: a comparison of geographically weighted regression and spatially varying coefficient model”, *Stochastic Environmental Research and Risk Assessment* 21:5:573–588.

Walsh, W. (2001), “Compstat: An Analysis of an Emerging Police Paradigm”, *Policing: An International Journal of Police Strategies and Management* 24:3:347-363.

Wang, Q., Ni, J. and Tenhunen, J. (2005), “Application of a geographically-weighted regression analysis to estimate net primary production of Chinese forest ecosystems”, *Global Ecology and Biogeography* 14:4:379–393.

Weir, R., and Bangs, M. (2007), *The Use of Geographic Information Systems by Crime Analysts in England and Wales*, Home Office Online Report Series, London: Home Office.

Weisburd, D.L., and Braga, A.A. (2006), "Hot spots policing as a model for police innovation" in D.L. Weisburd and A.A. Braga (Eds.) *Police Innovation: Contrasting Perspectives*, New York: Cambridge University Press.

Weisburd, D.L., Bushway, S., Lum, C., and Yang, S.M. (2004), "Trajectories of crime at places: a longitudinal study of street segments in the city of Seattle", *Criminology* 42:2:283-321.

Wheeler D., and Calder, C. (2007), "An assessment of coefficient accuracy in linear regression models with spatially varying coefficients", *Journal of Geographical Systems* 9:145-166.

Wheeler D., and Tiefelsdorf, M. (2005), "Multicollinearity and Correlation among Local Regression Coefficients in Geographically Weighted Regression", *Journal of Geographical Systems* 7:161-18.

Wheeler, D., and Waller, L. (2009), "Comparing spatially varying coefficient models: A case study examining violent crime rates and their relationships to alcohol outlets and illegal drug arrests", *Journal of Geographical Systems* 11:1: 1-22.

Wigton, R.S., Connor, J.L., and Centor, R.M. (1986), "Transportability of a decision rule for the diagnosis of streptococcal pharyngitis", *Archives of Internal Medicine* 146:81-83.

Wiles, P., and Costello, A. (2000), *The 'road to nowhere': the evidence for travelling criminals*, London: Home Office.

Williamson, D., McLafferty, S., McGuire, P., Goldsmith, V., and Mollenkopf, J. (1999), *A Better Method to Smooth Crime Incidence Data*, ArcUser Magazine**,** January – March 1999, 1-5.

Williamson, D., McLafferty, S., McGuire, P., Ross, T., Mollenkopf, J., Goldsmith, V., and Quinn, S. (2001), "Tools in the Spatial Analysis of Crime", in A. Hirschfield and K. Bowers (Eds.), *Mapping and Analysing Crime Data: Lessons from Research and Practice*, London: Taylor and Francis.

Wright, R., and Decker, S. (1994), *Burglars on the Job*, Boston, MA: Northeastern University Press.

Zhang, L., Bi, H., Cheng, P., and Davis, C.J.(2004), "Modeling spatial variations in tree diameter-height relationships", *Forest Ecology and Management* 189:317-329.

Zhang, L. and H. Shi. (2004), "Local modeling of tree growth by geographically weighted regression", *Forest Science* 50:2:225-244.

Zhang, L., Gove, J.H., and Heath, L. (2005), "Spatial residual analysis of six modeling techniques", *Ecological Modelling* 186:154-177.

Zhang, L., Ma, Z., and Guo, L. (2009), "Spatial autocorrelation and heterogeneity in the relationships between tree variables", *Forest Science* 55:6:533-548.

Zipf, G.K. (1949), *Human behaviour and the principle of least effort*, Oxford: Addison-Wesley Press

Zou, K.H., O'Malley A.J., and L. Mauri (2007), "Receiver-Operating Characteristic Analysis for Evaluating Diagnostic Tests and Predictive Models", *Circulation* 115:654-657.