(Looking) Back to the Future: using space–time patterns to better predict the location of street crime

Lisa Tompson[‡] and Michael Townsley[†]

‡(Corresponding author) UCL Jill Dando Institute of Crime Science (JDI), University College, Brook House, 2–16 Torrington Place, London, WC1E 7HN. Tel: +44 (0)203 1083126; email: l.tompson@ucl.ac.uk
†School of Criminology and Criminal Justice, Griffith University, Australia. Email: m.townsley@griffith.edu.au

Lisa Tompson joined the JDI in June 2006 as a Research Fellow. She has worked as a UK Police Intelligence Analyst and holds a BA(Hons) in Sociology and Criminology (Keele University) and an MSc in Crime Science (University College London). Her research interests focus on spatial approaches to crime analysis and their relevance to crime reduction.
Dr Michael Townsley is a lecturer at the School of Criminology and Criminal Justice, Griffith University, Brisbane, Australia. His research interests include environmental criminology, crime analysis, crime prevention and policing.

ABSTRACT

Crime analysts attempt to identify regularities in police recorded crime data with a central view of disrupting the patterns found. One common method for doing so is hotspot mapping, focusing attention on spatial clustering as a route to crime reduction (Chainey & Ratcliffe, 2005; Clarke & Eck, 2003). Despite the widespread use of this analytical technique, evaluation tools to assess its ability to accurately predict spatial patterns have only recently become available to practitioners (Chainey, Tompson, & Uhlig, 2008). Crucially, none has examined this issue from a spatio-temporal standpoint. Given that the organisational nature of policing agencies is shift based, it is common-sensical to understand crime problems at this temporal sensitivity, so there is an opportunity for resources to be deployed swiftly in a manner that optimises prevention and detection.

This paper tests whether hotspot forecasts can be enhanced when time-of-day information is incorporated into the analysis. Using street crime data, and employing an evaluative tool called the Predictive Accuracy Index (PAI), we found that the predictive accuracy can be enhanced for particular temporal shifts, and this is primarily influenced by the degree of spatial clustering present. Interestingly, when hotspots shrank (in comparison with the all-day hotspots), they became more concentrated, and subsequently more predictable. This is meaningful in practice; for if crime is more predictable during specific time-frames, then response resources can be used intelligently to reduce victimisation.

INTRODUCTION

Contrary to conventional wisdom, the occurrence of crime is rarely random. The assertion that crime is unevenly distributed

between victims, offenders and places is underpinned by a substantial empirical evidence base (Pease, 1998; Pease & Laycock, 1996; Sherman, Gartin, & Buerger, 1989). In particular, crime concentrates relative to spatial characteristics, reflecting the existence and attractiveness of crime opportunities in locations where offenders and victims interact (Clarke, 1992; Cohen & Felson, 1979; Cornish & Clarke, 1986).

Environmental criminology (Brantingham & Brantingham, 1984) has accumulated substantive insight into the distribution of crime opportunities, both in space and time (daily, weekly or yearly patterns). A number of authors (Felson & Poulson, 2003; Ratcliffe, 2006; Rengert, 1997; Townsley, 2008) have noted the disproportionate attention given to spatial concentrations of crime compared with time patterns. This is surprising given that the prevention of crime almost always begins with a prediction of crime patterns. Furthermore, theory and empirical evidence would suggest that spatial crime patterns should be sensitive to temporal features, possibly at the day, week and month level. Championing this point, Felson and Poulson (2003, p. 1) declare that 'crime varies greatly by hour of day — more than by any other variable'.

Research has repeatedly demonstrated that offenders prefer to return to a location associated with a high chance of success instead of choosing random targets (Pease, 1998). There is empirical evidence that residential burglars favour committing offences at similar times of the day (Coupe & Blake, 2006; Rengert & Wasilchick, 2000; Sagovsky & Johnson, 2007), implying that if they are relatively consistent in where and when they optimally forage for crime opportunities (Johnson & Bowers, 2004) there exist prospects for forecasting and, as a corollary, crime prevention. Rengert (1997) provides empirical evidence that non-criminal routine activities generate the

opportunity structure of vehicle theft in Philadelphia. He found that different periods of the day were associated with different spatial clusters of crime and these were generated by the transient population flux through the day. At a more abstract level, Ratcliffe (2006) argues that offenders operate under temporal constraints (the need to be somewhere at some time) and that these determine the distribution of opportunistic property crimes. He goes on to demonstrate that some properties will be exposed to offender scrutiny for relatively small amounts of time but at roughly the same time each day.

The purpose of this paper is to investigate the degree to which spatial patterns of crime are predictable and the relative improvements in forecasting that can be achieved when temporal information is incorporated into the analysis. The remainder of this paper is structured as follows: the first section offers a brief discussion of the prognostic expectations of crime analysis; the second section outlines the data and methodology used to test the research questions; the analytical results are presented in the third part of this paper; and finally we draw on these results to interpret in what conditions improvements in forecasting are possible.

PROGNOSTIC EXPECTATIONS OF CRIME ANALYSIS

Crime analysis and hotspot identification

A core element of crime analysis is the identification and examination of crime hotspots. Crime hotspots are geographic locations 'of high concentration, relative to the distribution of crime across the whole region of interest' (Chainey & Ratcliffe, 2005, p. 147). Hotspot analysis is possible because of the 'inherent geographical quality' present within crime data (Chainey &

Ratcliffe, p. 1) and moreover an inherent temporal quality too; for all crimes have to happen *somewhere* and *sometime.* Currently the UK police use hotspot analysis as a rudimentary way to forecast crime; using retrospective data to conjecture where crime may concentrate in the future. Hotspot maps are a central part of crime analysis (see Boba, 2005) and hotspot patrolling or policing has thus become a key tactic to inhibit offending in particular areal subunits (Braga, 2001; 2005). Given that most, if not all, tactics within the offender disruption/ targeting doctrine are based on analysts' perception of 'who' is operating 'where', a process described by Ratcliffe (2008) as 'interpreting the criminal environment', it is essential that such perceptions are accurate and have prospective utility.

The ease with which hotspot maps are able to be generated has been greatly aided by the proliferation of Geographical Information System (GIS) computer software. A variety of techniques are available for identifying hotspots, each with its own merits and limitations (Eck, Chainey, Cameron, & Wilson, 2005). Reviews of different mapping methods concur that Kernel Density Estimation (KDE) is the most suitable technique for visualising crime data (Chainey & Ratcliffe, 2005; Chainey, Reid, & Stuart, 2002). Contentions with KDE generally centre on the setting of parameters (such as cell size, bandwidth and thematic class — see McLafferty, Williamson, & Maguire, 2000). The flexibility of these parameter settings means that vastly different maps can be produced with the same data, presenting a credibility problem if map makers are unaware of these influences. No doctrine that suits all purposes has been devised, although suggestions have been offered (Chainey & Ratcliffe, 2005; Eck et al., 2005; Ratcliffe, 1999). Notwithstanding the limitations raised above, KDE remains a popular technique (Brimicombe, Brimicombe, & Li, 2007), both with academics and the crime analysis profession.

'Hot times' of crime

Efforts to integrate the analysis of space and time have historically been hampered by poor temporal functionality in GIS (Lodha & Verma, 2000; Ratcliffe, 2002). Hitherto crime patterns were largely analysed independently in the spatial and temporal dimensions, such that insight in one did not routinely impact analysis in the other. McCullagh (2006) recognises that at best '. . . emphasis is usually placed on the spatial hotspot with only simplistic attempts to tie in temporal changes because of the complexities involved'. However, the space and time of offences are clearly interrelated and are likely to directly affect each other. By synthesising the two domains, analytical insight is increased and resources can be allocated not only at the right places, but also at the right times. An example of this is provided by Ratcliffe (2004) through the use of a hotspot matrix, an intuitive framework used to communicate the optimal allocation of resources based on spatio-temporal patterns in the data. Succinctly put, this methodology starts with analysis of the spatial and temporal patterns within a given data set, examining where each sits on a continuum of concentration. The identified patterns of concentration can then, based on their closest likeness, be classified into one of three typologies. At one extreme are relatively even or uniform patterns (diffused and dispersed for time and space respectively) and at the other extreme highly concentrated patterns (acute and hot point for time and space respectively). Combining these typologies in the hotspot matrix is a clear and effective way for decision-makers to understand the temporal and spatial clustering of a given crime problem, and engenders a more intelligent approach to devising response tactics.

Crime forecasting

There are a growing number of researchers dedicated to predicting where future crime events might occur based on an antecedent event. To date this work has examined residential burglaries (Bowers, Johnson, & Pease, 2004; Johnson, Bernasco, et al., 2007), gun crime (Ratcliffe & Rengert, 2008), theft from motor vehicles (Summers, Johnson, & Pease, 2007) and insurgent activity in Iraq (Townsley, Johnson, & Ratcliffe, 2008). However, these techniques are presently aspirations in the minds of most practitioners as the technological capability is confined to academia. Chainey, Tompson, and Uhlig (2008) have overcome some of these complexities by proposing an evaluative tool, the Prediction Accuracy Index (PAI), which provides practitioners with a means to test the accuracy of hotspot maps at predicting future locations of high crime. A similar, but distinct and independent, approach is given by Bowers, Johnson, and Pease where they test the performance of different hotspot prediction surfaces.

Crime forecasting is only useful if it is performed at an appropriate level of geographical resolution to the resourcing requirement (Johnson, Bernasco, et al., 2007). Police managers would benefit greatly from localised (at least to a neighbourhood level) forecasts, so that resource allocation is apposite and realistic. This leads to the ability to direct patrols to a focused area, remove and redeploy specialist units, schedule training and annual leave in the lower crime periods and generally understand the mechanisms driving crime in their area.

While the importance of spatial resolution is well understood, less well appreciated is the temporal resolution of crime forecasts. Routine activities vary by day, as do crime rates and crime risks (see Felson & Poulson, 2003; Ratcliffe, 2002) and these in turn affect the shift patterns that police personnel typically work. At present, hotspot maps generated by crime analysts routinely fail to capture the fluctuations of crime across the day (although this is slowly changing). If these same maps are used to forecast where spatial patterns of crime might be seen in the future, then making the forecasts more sensitive to temporal patterns could enhance the prediction accuracy. Hence, the main objective of this research is *to explore temporal dynamics within hotspots areas and how these might be used to inform predictions about levels and patterns of crimes in those areas.* An ancillary objective is that the results provide operationally relevant information to the police that may assist with where and when to deploy resources. To reinforce this, it was decided that the methodology should be consistent with Chainey et al. (2008), which is sympathetic to common applications used within the policing environment, thus affording replication by crime analysts or interested parties.

Research questions

This research set out to answer three main questions:

1. Can the predictive accuracy of hotspot forecasts (hot areas based on where crime has clustered in the past) be enhanced by incorporating temporal information into the analysis?
2. If so, are certain temporal periods better at predicting the same periods in the future than others?
3. If so, is the predictive accuracy directly influenced by the presence of spatial clustering, or temporal clustering or both?

Data

Previous research (Chainey et al., 2008) demonstrated that hotspots of 'street crime'

(defined as incidents of either personal robbery or theft from the person) were more accurate predictors of future crime events than other acquisitive crime types. In light of this, it was decided to focus on street crime for the present study. Recorded street crime data were acquired from the Metropolitan Police Service for 1 January 2002 to 31 December 2003 for the boroughs of Camden and Islington, North London: 10,160 events were recorded within the time frame with roughly equivalent numbers in each year, comprising 6,511 robberies and 3,649 theft-from-the-person crimes. We acknowledge the familiar drawbacks of using such data (see Maguire, 1997) but feel that these do not inhibit the quality of the research. These data are the same as used by Chainey et al. (2008).

For the purpose of addressing the research questions the fields extracted for each incident were: (i) geographic coordinates, (ii) time of day, and (iii) date. To comply with UK data protection legislation, the coordinate information was rounded to the nearest 10 metres (so an incident recorded at 530281, 183105 would be transformed to 530280, 183100). This introduced some level of imprecision; however, as this is of relatively small magnitude it was thought not to compromise the representativeness of the spatial distribution of these crime data.

METHODOLOGY

The PAI

The central premise of this research was to determine if the predictive accuracy of hotspot forecasts — 'hot' areas where previous crimes had clustered — could be improved when the time of day was incorporated into the analysis. Prediction accuracy can be measured in various ways; here we used the Prediction Accuracy Index (PAI). The index was devised by Chainey et al. (2008; and

since adopted by Van Patten, McKeldin-Coner, & Cox, 2009) to generate a measurement of prediction accuracy proportionate to different sized study areas. The PAI is the ratio of two statistics; the first is the *Hotspot Hit Rate,* which is the proportion of future (prospective) crime events falling within the hotspot forecasts, generated by the past (retrospective) data. The second ratio is the *Area Proportion,* which is the percentage of the entire study area which the hotspots occupy. In its simplest terms, the PAI is the proportion of the 'prospective' data in the hotspot forecast divided by the proportion of the study area covered by the hotspot forecast. The PAI is calculated using the following formula:

$$\frac{(n/N)\star 100}{(a/A)\star 100} = \frac{Hotspot\ hit\ rate\ (HHR)}{Area\ percentage\ (AP)}$$
$$= PREDICTION\ ACCURACY\ INDEX\ VALUE$$

where *n* is the number of 'future' crimes which fall into the hotspots created from retrospective data; *N* is the total number of 'future' crimes; *a* is the total area of the hotspots created with retrospective data; and *A* is the total area of the study region.

Temporal considerations

Although the PAI is the overarching methodology with which to test the predictive accuracy of hotspot forecasts, a number of decisions were necessary prior to calculating this value. Core to this research was the deliberation of how to analyse the spatio-temporal dynamics within the data so that the prediction accuracy might be improved.

Incorporating temporal information into the hotspots needed to balance two key priorities; first, that the methodology was accordant with extant empirical findings, and second, that the results were operationally relevant. It was known from Johnson, Birks, McLaughlin, Bowers, and Pease

(2007) (who report on the implementation of the ProMap model) that a consistency of time-of-day patterns in crime data justified production of forecasts aligned to police shifts. It was further known, from the authors' prior experience, that police response resources (usually teams of police officers) were often allocated to tasks by shift availability. This then seemed a useful avenue to explore, as the time-of-day information could be used in a way that would complement the allocation of resources typically used to disrupt crime.[1]

After consultation with several police representatives, it was decided to split the data into equally sized temporal 'shifts' that were closely aligned with police personnel working hours. This could not be achieved in full, as police shifts tend to overlap each other so that service delivery is resilient to staff changeover. However, the consultation provided a benchmark from which to create these temporal shifts.

Table 1 shows how the data were temporally partitioned with events allocated to each according to their start times.[2] At a coarse level the temporal clustering of the data is illustrated by the differing percentages seen in this table. The greatest proportion of offences was committed in the 'afternoon' period (nearly one-half of all

## Table 1: Temporal 'shift' periods and descriptive statistics

Name given to temporal 'shift'	Time period included	n events	% of all events
Morning	07:00–14:59	2,149	21.1
Afternoon	15:00–22:59	4,807	47.3
Overnight	23:00–06:59	3,204	31.5

crime events) and the smallest share in the 'morning' period (just over one-fifth). To visualise the temporal clustering within each shift we created Figure 1 — this shows that all shifts had different temporal patterns operating. Temporal clustering peaked at the end point of the morning and afternoon shifts and at the beginning of the overnight shift.

Defining and creating past and future data

To assess the predictive accuracy of the different temporal shifts it was necessary to define what data (time) period we would use to create 'hotspot forecasts' from, and the data period we would subsequently use to evaluate the forecasts. In other words, we needed to define past or 'retrospective' data

Figure 1
The temporal distribution of street crime data (by starting hour of offence). Note: Dashed vertical lines denote the cutoff point between temporal shifts
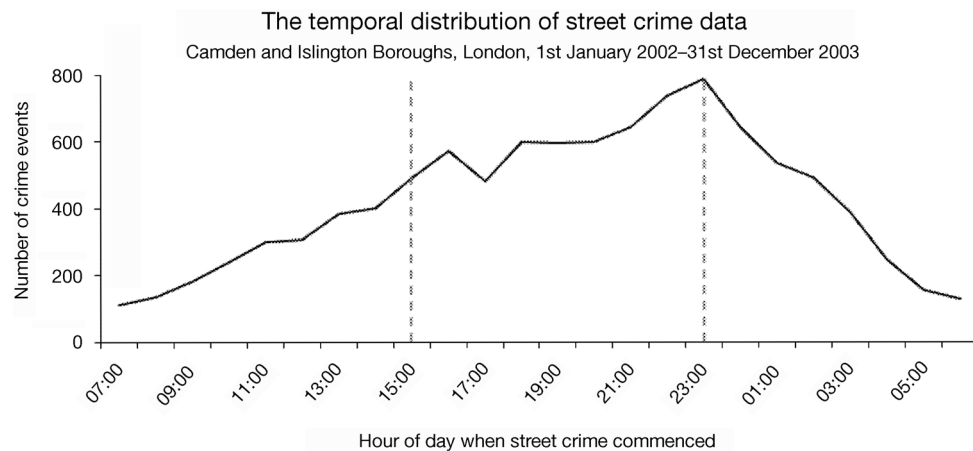
alongside future or 'prospective' data. To do this we created Measurement Dates (MDs) to discriminate the cutoff point between 'retrospective' and 'prospective' data.

In the absence of a means to automate the analysis, the authors had to weigh up the preferred quantity of MDs against the time each one would take to analyse. Previous research had used one, or two MDs (Alvarez, 2004; Chainey et. al., 2008; Keutgen, 2006), and had subsequently suggested that the midpoint of the data — 1 January, 2003 — was a fairly atypical day in the UK calendar. To mitigate for an unrepresentative crime pattern influencing the results of this analysis, we chose four further MDs,[3] which are denoted in Table 2.[4] These had a suitable temporal buffer from the midpoint, so that daily activities could be presumed to be as 'normal' as possible. A further consideration was that they left a sufficient amount of retrospective or prospective data with which to carry out the analysis.

Creating the hotspot forecasts with retrospective data

The question of how much data to use to create a reliable hotspot map is open to debate. When hotspots are transient, the most recent data will be the closest reflection of what will happen next. It logically follows in that scenario that a relatively small retrospective period (relative to the degree of transiency) would be the most appropriate data to use to predict future

hotspots. Conversely, if a hotspot were stable, the recency of the data would not have the same effect, so one might prefer to use longer retrospective data periods to create hotspot forecasts in this case. This would also reduce the influence of another potential spatio-temporal pattern: seasonality in the data. At the time of this research there was no known test for the stability of a hotspot; however, in the meantime Johnson, Lab, and Bowers (2008) have devised a suitable test for just this situation.

In recognition of the potential repercussions of using one retrospective data period, we instead chose to test nine and called them retrospective 'Measurement Periods' or MPs. These increments of data are shown in Table 3 by MD. These were designed to reflect the different horizon lengths in a policing environment (ie, short-, medium- and long-term), but could be adjusted for any relevant operational requirements. These nine retrospective MPs were finally combined with the other temporal breakdowns (the temporal shifts and the five measurement dates).[5] Latterly they would be used to generate *hotspot forecasts*, but the next step was to test their levels of spatial clustering.

Investigating spatial clustering

If a relationship exists between where crime has happened in the past and where crime will happen in the future, therein is an opportunity for crime prediction. Prediction accuracy relies on there being enough clustering in the retrospective data to make hotspot analysis worth pursuing. Mindful of this, prior to generating crime hotspots we conducted tests to determine the degree of spatial clustering in these retrospective data.

The Nearest Neighbour Index (NNI) was computed to identify whether data displayed tendencies of clustering or dispersion, and whether that tendency was statistically significant. All NNI values under one

Table 2: MDs and days of the week

Measurement date	Day of week
12 October 2002	Saturday
11 November 2002	Monday
1 January 2003	Wednesday
11 February 2003	Tuesday
13 March 2003	Thursday

Table 3: How the data were sliced to create retrospective measurement periods (MPs)

Retrospective measurement periods used to create crime hotspots

MD	Max	6 months	3 months	2 months	1 month	2 weeks	1 week	3 days	2 days	1 day
12/10/02	01Jan02–11Oct02	12Apr02–11Oct02	12Jul02–11Oct02	12Aug02–11Oct02	12Sep02–11Oct02	28Sep02–11Oct02	5Oct02–11Oct02	9Oct02–11Oct02	10Oct02–11Oct02	11Oct02
11/11/02	01Jan02–10Nov02	11May02–10Nov02	11Aug02–10Nov02	11Sep02–10Nov02	11Oct02–10Nov02	28Oct02–10Nov02	4Nov02–10Nov02	8Nov02–10Nov02	9Nov02–10Nov02	10Nov02
01/01/03	01Jan02–31Dec02★	01Jul02–31Dec02	01Oct02–31Dec02	01Nov02–31Dec02	01Dec02–31Dec02	18Dec02–31Dec02	25Dec02–31Dec02	29Dec02–31Dec02	30Dec02–31Dec02	31Dec02
11/02/03	11Feb02–10Feb03★	11Aug02–10Feb03	11Nov02–10Feb03	11Dec02–10Feb03	11Jan03–10Feb03	28Jan03–10Feb03	4Feb03–10Feb03	8Feb03–10Feb03	9Feb03–10Feb03	10Feb03
13/03/03	13Mar02–12Mar03★	13Sep02–12Mar03	13Dec02–12Mar03	13Jan03–12Mar03	13Feb03–12Mar03	27Feb03–12Mar03	06Mar03–12Mar03	10Mar03–12Mar03	11Mar03–12Mar03	12Mar03

Note:
★ Retrospective measurement periods containing one year of data.

were taken as evidence of data clustering, and justified the later step of creating the hotspot forecasts. As a rule of thumb the lower the NNI score, the higher the degree of clustering. NNI values for these data were computed using the CrimeStat III application (Levine, 2006) and were complemented by associated z-scores and p-values. The results were used to determine the threshold at which the data became statistically significantly clustered, and hence were appropriate to use to generate hotspot forecasts. Data periods that did not exhibit clustering at a statistically significant level would thus be excluded from the analysis. Nearest neighbour calculations were hence carried out for all combinations of shift (3), MDs (5) and MPs (9), resulting in 135 sets of results.

The mean across the temporal shifts showed the afternoon offences exhibited the least clustering (mean NNI = 0.79) and furthermore, a greater number of offences were required to reach statistical significance (mean n = 129). The overnight offences exhibited the most clustering (mean NNI = 0.73), requiring a mere 46 offences on average to reject the null hypothesis of complete spatial randomness. Table 4 shows the smallest time window (retrospective MP) which produced a statistically significant amount of spatial clustering, i.e. when the NNI value was reliably less than 1.

Determining a threshold for when spatial clustering had reached statistical significance for all possible MD and shift combinations was not straightforward. As Table 4 shows, each temporal shift and MP combination had varying results. It was decided to use one month of data as the threshold as this was the point at which it could be confidently assumed that all shifts had statistical evidence of clustering. All retrospective MPs less than a month were discarded at this point. Thus, hotspot maps were created using one, two, three, six and, where available, twelve months of data for each MD and temporal shift (resulting in 75 retrospective MPs in total). Where twelve months of data were not available — due to the MD occurring before 1 January 2003 — the maximum available was used.[6]

Choice of hotspot methodology

A variety of methods to locate spatial concentrations of crime exist. For this study we selected Kernel Density Estimation (KDE) to identify hotspots as it is commonly employed by UK crime analysts (Weir & Bangs, 2007) and the results would consequently appeal to the police service. KDE has also been found by previous research to be associated with the highest predictive accuracy when compared with alternative hotspot identification techniques (Chainey et al., 2008).

The KDE technique produces a continuous surface, representing the density of crime events over the study area. This is created by overlaying a grid on the study

Table 4: The smallest time window that yields statistically reliable spatial clustering for each combination of MD and temporal shift

	Measurement date				
	12/10/02	11/11/02	01/01/03	11/02/03	13/03/03
Morning	1 month	1 month	1 month	1 month	1 month
Afternoon	2 weeks	1 month	1 month	1 week	2 weeks
Overnight	2 weeks	2 weeks	2 weeks	1 week	3 days

area and computing the number of events in each grid. The calculation is simply a weighted sum of the crimes that occur within a certain range, or bandwidth, for each cell of the grid. Once these summary values are obtained, they can be binned into intervals and colour coded (see Chainey & Ratcliffe, 2005). The hotspot areas are far more precise and compact using this method; Chainey et al. (2008) identify that this is one of the reasons why KDE generates a higher predictive accuracy.

KDE maps were generated using Hotspot Detective (an extension for MapInfo Professional; Ratcliffe, 1999) for each of the 75 retrospective data periods that had been created. Following Chainey et al. (2008), the default values (that Hotspot Detective calculates) for the bandwidth, cell size, kernel type and weighting scheme (ie, none) were used. Five thematic classes (colours) were chosen and default values resulting from choosing the quantile ranges were applied.[7] All hotspot maps adhered to these parameter choices to ensure consistency with each other and the results from Chainey et al.

To create the hotspot forecasts we extracted the top thematic class from the KDE maps using a GIS query. Those cells with the highest density were then saved as a separate file and their total area calculated. This figure was populated into the PAI equation (as *a*) along with the total study region area (*A*) to calculate the Area Proportion (AP) for each of the 75 hotspot forecasts.

Using prospective crime to calculate the PAI

The Hotspot Hit Rate (HHR) — the proportion of future or 'prospective' crimes that fell within the hotspot forecasts (for each Measurement Date (MD)) — was calculated next. In order to determine the temporal utility or 'shelf life' of the hotspot forecasts, different prospective data horizons

were defined in the same manner as the retrospective data. These prospective MPs were one day; three days; one week; two weeks; one month; two months; three months; six months and twelve months (where twelve months were not available due to an MD after 1 January 2003 the maximum data were used). Table 5 shows the resulting prospective MPs for each MD.

Once the prospective MPs were created we entered the total number of crimes in each into the PAI equations as *N*. Next, we took each hotspot forecast in turn (there were fifteen in each MD — five for each temporal shift) and overlaid the prospective MP data from the corresponding MD. We then used a GIS query to count the number of points (the prospective MPs) falling within polygons (the hotspot forecast cells) and entered this value as *n* into the PAI equation. Hence, we were able to test for each MD the predictive power of hotspots, using different amounts of data to generate the forecasts (retrospective MPs), and how these perform over a range of time horizons (prospective MPs). In addition, the data were partitioned by temporal shifts to investigate whether certain times of day yield better predictions. Moreover, we were able to test *all* combinations of the above once the process described above was completed for all MDs. An example of how the PAI equation values were captured is shown in Table 6.

By virtue of the PAI being a ratio, its value can be interpreted as a predictive effect of the retrospective MP, MD and shift combination. For instance, if 20 per cent of all crimes in the prospective MP (ie, the future crimes) were within the hotspot forecast (HHR), and the hotspot covered 10 per cent of the study area (AP), then a PAI of 2 (= 0.2/0.1) implies the hotspot was twice as good as selecting areas at random. Hence, it is not possible to say that a PAI value of *x* translates to a number of

Table 5: How the data were sliced to create prospective measurement periods

Prospective measurement periods used to generate the Hotspot Hit Rate

MD	1 day	2 days	3 days	1 week	2 weeks	1 month	2 months	3 months	6 months	Max
12/10/02	12Oct02	12Oct02–13Oct02	12Oct02–14Oct02	12Oct02–18Oct02	12Oct02–25Oct02	12Oct02–11Nov02	12Oct02–11Dec02	12Oct02–11Jan03	12Oct02–11Apr03	12Oct02–11Oct03
11/11/02	11Nov02	11Nov02–12Nov02	11Nov02–13Nov02	11Nov02–17Nov02	11Nov02–25Nov02	11Nov02–10Dec02	11Nov02–10Jan03	11Nov02–10Feb03	11Nov02–10May03	11Nov02–10Nov03
01/01/03	1Jan03	1Jan03–2Jan03	1Jan03–3Jan03	1Jan03–7Jan03	1Jan03–14Jan03	1Jan03–31Jan03	1Jan03–28Feb03	1Jan03–31Mar03	1Jan03–30Jun03	1Jan03–31Dec03
11/02/03	11Feb03	11Feb03–12Feb03	11Feb03–13Feb03	11Feb03–17Feb03	11Feb03–24Feb03	11Feb03–10Mar03	11Feb03–10Apr03	11Feb03–10May03	11Feb03–10Aug03	11Feb03–31Dec03
13/03/03	13Mar03	13Mar03–14Mar03	13Mar03–14Mar03	13Mar03–19Mar03	13Mar03–26Mar03	13Mar03–12Apr03	13Mar03–12May03	13Mar03–12Jun03	13Mar03–12Sep03	13Mar03–31Dec03

Table 6: Example of how the PAI values were calculated and recorded

Hotspot forecast using an MD of 01/01/03, with 6 months retrospective MP, and the afternoon temporal shift

Morning prospective MP	Total crime events in prospective MP	Count of prospective MP data in hotspot forecast	Hotspot forecast area (km^2)	Total study area (km^2)	Hotspot Hit rate (HHR)	Area percentage (AP)	PAI Score
1 day	0	0	1.25178	36.55	0.00	3.425	0.00
3 days	5	1	1.25178	36.55	20.00	3.425	5.84
1 week	19	4	1.25178	36.55	21.05	3.425	6.15
2 weeks	47	10	1.25178	36.55	21.28	3.425	6.21
1 month	106	24	1.25178	36.55	22.64	3.425	6.61
2 months	185	44	1.25178	36.55	23.78	3.425	6.94
3 months	278	67	1.25178	36.55	24.10	3.425	7.04
6 months	582	134	1.25178	36.55	23.02	3.425	6.72
Maximum	1,048	252	1.25178	36.55	24.05	3.425	7.02

crimes predicted. The value can, though, be used as a relative measure as a comparison against other PAI combinations. The following section outlines the comparisons we made between our generated PAI values.

RESULTS

Have PAI scores improved with the addition of temporal shifts?

The key objective of this research was to determine if incorporating temporal information into crime hotspots enhanced the predictive value they offered in forecasting future areas of high crime concentration. Inferential tests to investigate a difference in PAI values could not be used as the data violated one of the main assumptions required for this type of testing — independence within the data samples. Instead, we chose to employ confidence intervals to visualise the results.

Confidence intervals are a range of values around a statistic within which we believe the true value of the population lies (ie, the parameter). For this research we have used a 95 per cent confidence level, which means that if we conducted the same analysis on 100 different samples, we would expect the mean PAI value to fall within the confidence interval on 95 occasions. To construct the confidence intervals, PAI scores were computed for all MD, MP and shift combinations. The sample mean and standard error for each shift combination were then used to calculate the confidence intervals. Figure 2 illustrates the results for all the PAI scores.

There are three groups of PAI values shown in Figure 2. The baseline PAI is located at the bottom and represents the results obtained without incorporating time of day information. They are, in effect, a replication of Chainey et al. (2008). The next group contains the PAI values for the three same shift predictions (retrospective morning used to predict the prospective morning, etc). These show that the morning–morning combination results in slightly higher PAI values, the afternoon–afternoon PAI are much lower than the baseline and the overnight–overnight combination shows a substantial improvement
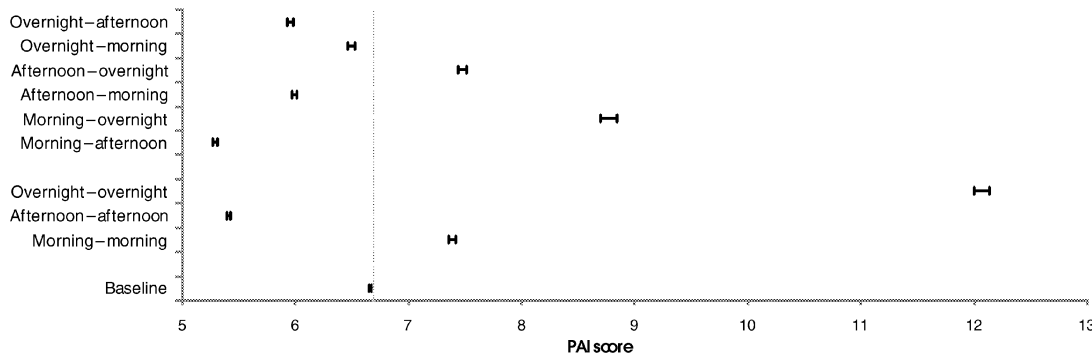
compared with baseline values. The final group in Figure 2 are the PAI values for dissimilar shifts. Combinations that predict the overnight period yield a higher PAI value but otherwise the predictions are worse than the baseline.

It was observed when predicting prospective periods (MPs) that had less than two weeks data that the resultant PAI values tended to fluctuate widely. For this reason, we conducted a replication of the analysis but omitted PAI values with a prospective MP of less than two weeks (we thus removed three prospective MPs), reducing the sample size from 225 to 150 in each shift combination. Figure 3 shows the results.

Figure 3 illustrates the influence of removing the smaller prospective MPs (differing lengths of the immediate fortnight after the MDs). The confidence intervals unanimously shrink in size because the variability in the PAI values is reduced, and in some cases change perceptibly. The overnight-overnight results are consistent in producing superior results. In fact, their predictive accuracy improves from a baseline mean PAI of around 7, to 11.7. Note that the PAI value is a predictive effect of the hotspot area in comparison with an area selected at random. When viewed in this way, these PAI values have improved substantially.

The morning-morning combination also establishes a more defined improvement of PAI results in the second analysis, whereas the afternoon-afternoon counterpart is still poor despite a small increase in value from Figure 2. The remaining shift combinations moved closer to the baseline in this chart, with the exception of the two overnight retrospective periods, both of which im-
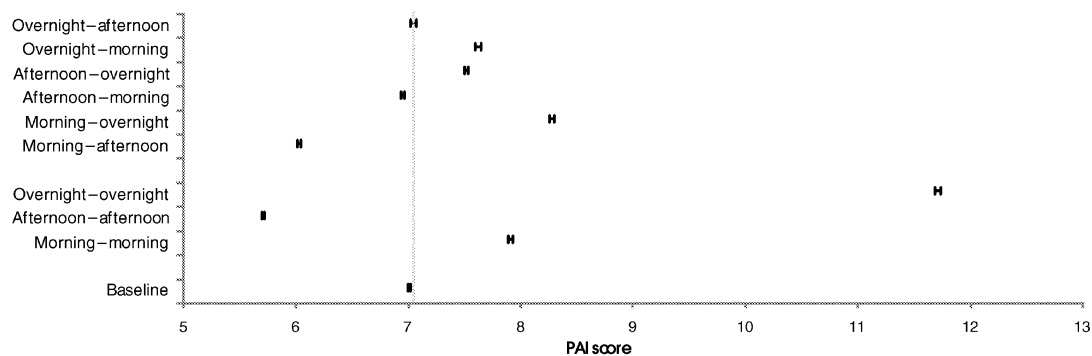
proved noticeably. The overnight-morning combination was now located above the baseline PAI value.

In terms of the research questions, there is conflicting evidence as to whether temporal information increases predictive accuracy. PAI values are reliably higher when incorporating temporal information, but only for specific shifts. The overnight period is by far the most predictable, and the afternoon period the worst at either predicting or being predicted.

The third research question focused on whether any observed predictive gain was a result of spatial or temporal clustering. Table 1 and Figure 1 showed that the overnight period is not the most temporally clustered shift, in so far as it does not contain the highest volume of offences, but instead the afternoon shift contains the bulk of crimes. From this we surmise that temporal clustering is not associated with PAI values. However, we recognise that taking the volume of offences in each shift introduces an aggregation problem, and may disguise the temporal clustering influence within each shift. In contrast, when we reviewed the spatial clustering within each shift (see NNI results and Table 4), the overnight shift illustrated the most clustering and the afternoon shifts the least. We concluded from this that a tentative relationship between the level of spatial clustering in a temporal shift and higher PAI values existed.

Explaining the rise in PAI values

Understanding why some PAI values had improved over others required examining the two components of the PAI, the HHR and the AP. When we examined how each of these contributed to the observed PAI values, the vast bulk of the differences came about due to variability in the AP and not due to the HHR. The baseline analysis generated a mean AP of 0.0294 (SD = 0.0031), equivalent to about 3 per cent of the study area. This was on par with the

mean AP found in this research for the afternoon temporal period (0.0304, SD = 0.049). In contrast, both the morning and overnight shifts saw a reduction in AP of about one-third of the baseline PAI (0.0193 and 0.0175 respectively, with SD = 0.0054 and 0.0032).

Introducing a temporal element into the analysis resulted in making the hotspot forecasts, at least for morning and overnight offences, far more geographically compact than when the temporal element is not included. This 'shrinking' of the hotspot is noteworthy because it suggests that hotspot maps used to forecast can be made more precise when they are restricted to a specific temporal period. The hotspots are then temporally determined and result in smaller areas to target. From an operational perspective, temporally determined hotspot forecasts restrict attention to a more condensed area which is easier to allocate resources to, and police officers can have confidence that those resources are in the appropriate place at the right times.

DISCUSSION AND IMPLICATIONS

This research has demonstrated that the predictive accuracy associated with hotspot maps used to forecast can be enhanced when they are made to be temporally sensitive. Hotspot maps that are used to provide a 'snapshot' of geographical crime concentration lack the temporal variations operating across the course of a day, and consequently can be misleading when the police come to allocate resources on the basis of such maps.

Our results indicated that enhanced PAI values could be achieved, but this did not occur uniformly across the three temporal shifts. The afternoon shift repeatedly showed little improvement in predictive accuracy, and only modest increases were seen when used to predict future overnight offences. In contrast, both the morning and

overnight hotspots exhibited improved PAI values; however these were confined to predicting future morning and overnight offences. The overnight shift outperformed all others in both being predictable and predicting (and when combined generated much higher PAI values).

The most interesting finding to emanate from this research was that there was something unusual — or special — about the overnight shift. One plausible explanation for this may be drawn from routine activity theory (Cohen & Felson, 1979). Given that street crime is wholly dependent on a victim being present, the clustering of offences in time and space is a product of the routine activities of the victims as much as the offenders. Examination of the hotspot locations revealed that they were closely aligned with land used for recreational pursuits (such as restaurants, bars, shops and transport). Commercial areas such as these tend not to be spatially dispersed and remain relatively static over time. Therefore the routine activities of victims, and consequently the opportunities to commit street crime are relatively static. There are fewer commercial premises open (shops shut and bars, restaurants, theatres and transport often have closing times) during hours of darkness, meaning that people are drawn to fewer establishments, often in more compact areas. Additionally, the overnight temporal shift covers hours of darkness in which people may behave in a more spatially constrained way than they would in daylight hours (Ratcliffe, 2006). The crime opportunities thus available within the overnight shift are more geographically consistent, indicating stable hotspot areas that shrink in comparison with the all-day hotspot areas. This is precisely the opposite for the afternoon period, which arguably contains a plethora of routine activities, including the end of school and work for a large segment of society, the beginning of recreational pursuits and the commencement of the evening entertainment industry.

By imposing boundaries on the data, analogies can be drawn with the Modifiable Areal Unit Problem, which applies to spatial boundaries. We concur with other researchers (such as Johnson, Lab, & Bowers, 2008) that there exists a *modifiable temporal unit problem*: when temporal boundaries are used to aggregate data, these boundaries influence the resultant patterns and can change when the boundaries are adjusted. Overcoming this limitation requires techniques to be developed that consider the temporal distribution of crime along a continuum, or continuous spectrum, similar to the way that KDE mapping does for spatial boundaries. In the absence of such a technique (and to the authors' knowledge this is currently the case), one further way of exerting greater temporal precision would be to make the boundaries smaller — say hourly slots of crime data. This may though become impractical for some data sets, as the sample size in each hour shift may not be large enough to conduct robust analysis from.

One revelation from this analysis was that the PAI values increased as a product of the hotspot area proportion (AP) decreasing, rather than the hotspot hit rate (HHR) increasing. This raises the issue of the validity of controlling the size of the predicted area. If reducing the hotspot size serves to increase the predictive accuracy, then future research would do well to concentrate on means of applying this principle with a view to an optimal size.

Considering the current findings, there is clearly a need for replication, both for different crime-types and also for greater time periods. Different temporal dynamics are expected to yield vastly different patterns from those observed here. We postulate that the trends found in our research (eg, hotspots shrinking overnight) may hold for

other crimes where there is a human victim — such as alcohol-related violence. It is unlikely that this is true for crimes relating to property — such as residential burglary — where the victim, or target, is not inclined to behave differently in hours of darkness in relation to space.

In conclusion, the predictive accuracy of hotspots is directly influenced by the level of spatial clustering over the course of a day, although the temporal clustering appeared not to have as much authority over the crime forecast. Thus, using temporally-sensitive hotspots increases the level of predictive accuracy, when the chosen temporal shifts aptly reflect routine activity periods.

NOTES

(1) We encourage other researchers to examine other spatio-temporal dynamics such as seasonal effects or day-of-week effects. Given that the shift patterns that police officers usually work are not on a seven-day cycle, we felt that day-of-week temporal patterns would not be as operationally useful.

(2) The average offence duration for these data was under six minutes, so the impact of selecting the start, end or even mid-point is likely to be trivial.

(3) One of which was the same date used in Chainey et al. (2008) to enable comparison of the results.

(4) A series of analyses, not reproduced here, were conducted to ensure that an appropriate knowledge of the temporal nature of street crime was gained prior to making decisions which could influence the outcome of the results. No patterns were observed, such as highly skewed day-of-week patterns, that we believe would bias the results based on the selection of MDs. Interested readers can obtain these analyses by contacting the corresponding author.

(5) For example, a 'morning' shift, with three months' data before 11 November 2002.

(6) Later on in the paper, 'Maximum' is used to denote where 12 months of data were not available. Instead the greatest volume of data available was used, which was not less than 9.5 months.

(7) The quantile method determines groups by summing all KDE values (to calculate an overall total) and dividing this by the number of thematic classes to give a 'group total'. All cells are then ranked by their KDE values and grouped so that summing the KDE values for each group is less than or equal to the derived group total.

REFERENCES

Álvarez, E. G. (2004). *The influence of kernel density estimation parameters on the predictive accuracy of crime hotspot maps.* Unpublished MSc dissertation, University College London.

Boba, R. (2005). *Crime analysis and crime mapping: An introduction.* Thousand Oaks, CA: Sage Publications.

Bowers, K. J., Johnson, S. D., & Pease, K. (2004). Prospective hot-spotting: The future of crime mapping? *British Journal of Criminology, 44*(5), 641–658.

Braga, A. (2001). The effects of hot spot policing on crime. *ANNALS of American Academy of Political and Social Science, 578,* 104–125.

Braga, A. (2005). Hot spots policing and crime prevention: A systematic review of randomized controlled trials. *Journal of Experimental Criminology, 1,* 317–342.

Brantingham, P. L., & Brantingham, P. J. (1984). *Patterns in Crime.* New York: Macmillan.

Brimicombe, A. J., Brimicombe, L. C., & Li, Y. (2007). Improving geocoding rates in preparation for crime data analysis. *International Journal of Police Science and Management, 9,* 80–92.

Chainey, S., & Ratcliffe, J. (2005). *GIS and crime mapping*. Chichester: Wiley.

Chainey, S. P., Reid, S., & Stuart, N. (2002). When is a hotspot a hotspot? A procedure for creating statistically robust hotspot maps of crime. In D. Kidner, G. Higgs, and S. White (Eds.) *Innovations in GIS 9: Socio-economic applications of geographic information science* (21–36). London: Taylor & Francis.

Chainey, S. P., Tompson, L., & Uhlig, S. (2008). The utility of hotspot mapping for predicting spatial patterns of crime. *Security Journal, 21*, 4–28.

Clarke, R. V. (1992). *Situational crime prevention: Successful case studies*. Albany, NY: Harrow and Heston.

Clarke, R., & Eck, K. (2003). *Become a problem-solving crime analyst in 55 small steps*. London: Jill Dando Institute, University College London.

Cohen, L. E., & Felson, M. (1979). Social change and crime rate trends: A routine activity approach. *American Sociological Review, 44*, 588–608.

Cornish, D., & Clarke, R. V. (1986). *The reasoning criminal: Rational choice perspectives on offending*. New York: Springer.

Coupe, T., & Blake, L. (2006). Daylight and darkness targeting strategies and the risks of being seen at residential burglaries. *Criminology, 44,* 431–464.

Eck, J. E., Chainey, S., Cameron, J., & Wilson, R. (2005). *Mapping crime; understanding hotspots*. Washington, DC: National Institute of Justice.

Felson, M., & Poulson, E. (2003). Simple indicators of crime by time of day. *International Journal of Forecasting, 19*(4), 595–601.

Johnson, S. D., & Bowers, K. J. (2004). The stability of space–time clusters of burglary. *British Journal of Criminology, 44*, 55–65.

Johnson, S. D., Bernasco, W., Bowers, K. J., Elffers, H., Ratcliffe, J., Rengert, G. et al. (2007). Near repeats: A cross national assessment of residential burglary. *Journal of Quantitative Criminology, 23*(3), 201–219

Johnson, S. D., Birks, D. J., McLaughlin, L., Bowers, K. J., & Pease, K. (2007). *Prospective crime mapping in operational context* (Final Report, Online Report 19/07). London, UK: Home Office.

Johnson, S. D., Lab, S. P., & Bowers, K. J. (2008). Stable and fluid hotspots of crime: differentiation and identification. *Built Environment, 34*(1), 32–45.

Keutgen, P. (2006). *Examining the prediction accuracy of the GI⋆ LISA statistic on hotspots of crime*. Unpublished MSc dissertation, University College London.

Levine, N. (2006). Crime mapping and the Crimestat program. *Geographical Analysis, 38*, 41–56.

Lodha, S. K., & Verma, A. K. (2000). *Spatio-temporal visualization of urban crimes on a GIS grid*. Proceedings of the 8th ACM International Symposium on Advances in Geographic Information Systems. Washington, DC, 174–179. Retrieved January 13, 2009 from http://portal.acm.org/citation.cfm?id=355300&jmp=cit&coll=GUIDE&dl=GUIDE&CFID=8275138&CFTOKEN=78266212.

Maguire, M. (1997). Crime statistics, patterns and trends: changing perceptions and their implications. In M. Maguire, R. Morgan, & R. Reiner (Eds.), *The Oxford handbook of criminology* (2nd ed., 135–188). Oxford: Clarendon Press.

McCullagh, M. (2006). *Detecting hotspots in time and space*. Retrieved March 29, 2007 from http://www.isotc211riyadh.org.sa/present/MJM_ISG06_%20Detecting_Hotspots_in_Time_and_Space_High_Quality_Print.pdf.

McLafferty, S., Williamson, D., & Maguire, P. G. (2000). Identifying crime hot spots using kernel smoothing. In V. Goldsmith, P. G. McGuire, J. H. Mollenkopf, & T. A. Ross (Eds.), *Analyzing crime patterns: Frontiers of practice*. Thousand Oaks, CA: Sage. 77-86.

Pease, K. (1998). *Repeat victimization: Taking stock* (Crime Prevention and Detection Series Paper 89). London, UK: Home Office.

Pease, K. & Laycock, G. (1996). *Reducing the heat on hot victims*. Washington: Bureau of Justice Statistics.

Ratcliffe, J. H. (1999). *Hotspot Detective for MapInfo helpfile version 1.0.* (Software Manual. Please contact authors for more information)

Ratcliffe, J. H. (2002). Aoristic signatures and the spatio-temporal analysis of high volume crime patterns. *Journal of Quantitative Criminology, 18*(1), 23–43.

Ratcliffe, J. H. (2004). The hotspot matrix: a framework for the spatio-temporal targeting of crime reduction. *Police Practice and Research, 5*(1), 5–23.

Ratcliffe, J. (2006). A temporal constraint theory to explain opportunity-based spatial offending patterns. *Journal of Research in Crime and Delinquency, 43*, 261–291.

Ratcliffe, J. (2008). *Intelligence-led policing.* Cullompton: Willan Publishing.

Ratcliffe, J. H., & Rengert, G. F. (2008). Near repeat patterns in Philadelphia shootings. *Security Journal, 21*(1–2), 58–76.

Rengert, G. F. (1997). Auto Theft in Central Philadelphia. In R. Homel (Ed.), *Policing for prevention: Reducing crime, public intoxication and injury: Vol. 7* (pp. 199–220) (Crime Prevention Studies). Monsey, NY: Criminal Justice Press.

Rengert, G., & Wasilchick, J. (2000). *Suburban burglary: A tale of two suburbs.* Springfield, IL: Charles C. Thomas.

Sagovsky, A., & Johnson, S. D. (2007). When does repeat burglary victimisation occur? *Australian and New Zealand Journal of Criminology, 40*(1), 1–26.

Sherman, L. W., Gartin, P., & Buerger, M. (1989). Hot spots of predatory crime: routine activities and the criminology of place. *Criminology, 27*, 27–55.

Summers, L., Johnson, S. D., & Pease, K. (2007). El robo de (objetos en) vehiculos y su contagio a traves del espacio y el tiempo: aplicaciones de técnicas epidemiológicas. *Revista Electronica de Investigacion Criminologica, 5*(1), 1–22.

Townsley, M. (2008). Visualising space time patterns in crime: the hotspot plot. *Crime Patterns and Analysis, 1*, 61–74.

Townsley, M., Johnson, S. D., & Ratcliffe, J. H. (2008) Space–time dynamics of insurgent activity in Iraq. *Security Journal, 21*(3), 139–146. doi: http://10.1057/palgrave.sj.8350090.

Van Patten, I.T., McKeldin-Coner, J., & Cox, D. (2009). A microspatial analysis of robbery: prospective hot spotting in a small city. *Crime Mapping: A Journal of Research and Practice, 1*(1), 7–32.

Weir, R., & Bangs, M. (2007). *The use of geographic information systems by crime analysts in England and Wales* (Home Office Online Report Series). London: Home Office.