Algorithmic prediction in policing: assumptions, evaluation, and accountability

Lyria Bennett Moses & Janet Chan

Published online: 08 Nov 2016.

Submit your article to this journal ↗

View related articles ↗

View Crossmark data ↗

Routledge
Taylor & Francis Group

Algorithmic prediction in policing: assumptions, evaluation, and accountability

Lyria Bennett Moses ⬤ and Janet Chan

Faculty of Law, UNSW Australia and Data to Decisions Cooperative Research Centre, Sydney, Australia

ABSTRACT

The goal of predictive policing is to forecast where and when crimes will take place in the future. The idea has captured the imagination of law enforcement agencies around the world. Many agencies are purchasing software tools with the goal of reducing crime by mapping the likely locations of future crime to guide the deployment of police resources. Yet the claims and promises of predictive policing have not been subject to critical examination. This paper provides a review of the theories, techniques, and assumptions embedded in various predictive tools and highlights three key issues about the use of algorithmic prediction. *Assumptions*: The algorithms used to gain predictive insights build on assumptions about accuracy, continuity, the irrelevance of omitted variables, and the primary importance of particular information (such as location) over others. In making decisions based on these algorithms, police are also directed towards particular kinds of decisions and responses to the exclusion of others. *Evaluation*: Media coverage of these technologies implies that they are successful in reducing crime. However, these claims are not necessarily based on independent, peer reviewed evaluations. While some evaluations have been conducted, additional rigorous and independent evaluations are needed to understand more fully the effect of predictive policing programmes. *Accountability*: The use of predictive software can undermine the ability for individual officers or law enforcement agencies to give an account of their decisions in important ways. The paper explores how this accountability gap might be reduced.

1. Introduction

'Predictive policing' is a term applied to a range of analytic tools and law enforcement practices. What links these together is the claimed ability to 'forecast where and when the next crime or series of crimes will take place' (Uchida 2014, p. 3871), combined with changes in law enforcement decision-making, particularly deployment of officers, based on those forecasts. As practised, mostly within the US but also elsewhere, the analytic element typically involves an off-the-shelf or adapted software tool that analyses historic crime data (and sometimes other data such as social media, weather, and mortgage defaults) to predict most commonly where, but sometimes by whom or to whom, crime will take place. In addition to forecasting, predictive policing involves taking a proactive response to crime where the goal is to 'change outcomes' (Beck and McCue 2009) through the identification of 'crime prevention tactics/strategies' and the evaluation of police programmes (Wilson *et al.* 2009) (see also section 2).

CONTACT Lyria Bennett Moses ✉ lyria@unsw.edu.au

www.

As a phenomenon, predictive policing is more than a set of tools. Predictive policing is also premised on the assumptions that it is possible to use technology to predict crime before it happens (van Brakel and De Hert 2011), that forecasting tools can predict accurately, and that police will use this knowledge effectively to reduce crime. But such positive beliefs around predictive policing are often based on a mythological and unrealistic view of actual capabilities and practices. As an example, one media article suggested that '[t]his complex equation can in theory predict, with pinpoint accuracy, where criminal offences are most likely to happen on any given day' (Adams 2012). It is clear that this statement is flawed, since the complexity of software goes beyond solving a single equation, no tool offers 'pinpoint accuracy' but rather larger blocks (such as 500 square foot boxes or street sections) and not all 'criminal offences' are equally suitable for forecasting (Sherman et al. 1989, p. 47, Hart and Zandbergen 2012, p. 58).

While the media report quoted above may represent an example of hype around predictive policing, an increasing number of police organisations are reported as using predictive policing tools. Most are in the US (Police Executive Research Forum 2014) and the UK, although predictive policing is reportedly being taken up to varying extents elsewhere. This paper will provide an important counter to existing rhetoric; it will examine the assumptions underlying the models on which predictive policing is based (section 3), shortcomings in existing evaluations (section 4) and challenges for accountability (section 5). While it is written as a critique, the goal is not to argue that predictive policing is ineffective or more problematic than alternatives, but rather to enhance understanding of what it entails.

2. Underlying model of predictive policing

The underlying model of predictive policing is represented in Figure 1. As Perry et al. (2013, p. 128) point out, predictive policing is 'not fundamentally about making crime-related predictions' but about implementing a prediction-led policing business *process*, which consists of a cycle of activities and decision points: data collection, analysis, police operations, criminal response, and back to data collection. At each stage of the cycle, choices are made regarding, for example, the types of data to collect, the duration and frequency of data collection and update, the types of analytical tools to

The Prediction-Led Policing Business Process

Figure 1. The prediction-led policing business process (from Perry et al. (2013, p. 128, Figure 5.4)). Copyright RAND included with permission.

employ, the dependent variables to focus on, the types of police operations to employ, how and when to evaluate the success of interventions, and what changes in interventions should be implemented following the evaluation. These choices can be taken and implemented in a variety of ways under a range of organisational conditions, including staffing, resourcing, institutional support, and technical expertise (Perry *et al*. 2013).

Conceptually, predictive policing is closely connected to, but distinguished from, a range of other approaches to law enforcement, including intelligence-led policing (ILP), data-driven policing, risk-based policing, 'hot spots' policing, evidence-based policing, and pre-emptive (Sherman *et al*. 1989, Lum *et al*. 2011, van Brakel and De Hert 2011, Ratcliffe 2012). In some cases, such as data-driven policing and hot spots policing, the primary difference is the explicit projection into the future. Rather than assuming continuity of current crime patterns and hot spots, predictive policing explicitly models change over time, often relying on evidence of statistically broader geographical impact of a single crime event (e.g. Bowers *et al*. 2004). Pre-emptive policing refers to action taken prior to an event occurring where it is believed an event would otherwise occur (as in the case of preventative detention). It can also suggest a broader shift towards a pre-crime society that seeks security against risk through prevention of crime, which is said to be characteristic of community and problem-oriented approaches to policing (van Brakel and De Hert 2011). Unlike predictive policing, however, pre-emptive policing need not be based on data-driven forecasting.

Despite differences, all of these types of policing are part of a broader evolution of policing strategies and technologies that increase the role of intelligence within law enforcement agencies. Predictive policing uses new analytic tools within a pre-existing ILP approach, which involves:

> A strategic, future-oriented and targeted approach to crime control, focussing upon the identification, analysis and 'management' of persisting and developing 'problems' or 'risks' … rather than on the reactive investigation and detection of individual crimes. (Maguire 2000, p. 315)

Another relevant initiative is the support offered to 'building evidence-based, data-driven law enforcement tactics and strategies that are effective, efficient and economical' by the Bureau of Justice Assistance in the US under the Smart Policing Initiative. This is not limited to predictive policing; it includes 'place-based policing' and 'offender-based policing' in a variety of jurisdictions, as well as 'predictive' approaches in Miami, Cambridge and Chicago.

Most analytic tools used for predictive policing seek to make probabilistic forecasts as to the location and timing of future crime, but there are tools more focussed on social media monitoring, CCTV monitoring and interpretation or the creation of individual profiles. Perry *et al*. (2013, p. 18) have sorted predictive techniques used in policing into four categories: (i) classical statistical techniques, using processes such as 'regression, data mining, time-series analysis, and seasonality adjustments'; (ii) simple methods such as using checklists and indexes; (iii) complex applications that require sophisticated computer programmes and large volumes of data; and (iv) tailored methods that use existing techniques to provide visualisation of data to support predictive policing. Table 1 (from Perry *et al*. 2013, p. 19) provides an overview of predictive techniques.

Software used for predictive policing can range from simple spreadsheet programs to complex algorithms. It can be open-sourced or inexpensive (Olesker 2012), purposed off-the-shelf, or specially tailored. Information on the tools themselves is often limited and source code is often a trade secret.

One popular tool on which some published information is available is PredPol®. PredPol® uses an earthquake prediction model to predict crime, as set out in the work of one of its founders (Mohler *et al*. 2011). This model assumes that crimes such as residential burglary or gang violence operate as a self-excited point process, similar to earthquake aftershocks. An earthquake or crime may be a 'background event' or it may be the aftershock or near repeat of another earthquake or crime nearby in space and time. In the case of both burglary and earthquake, the background rate will vary with time and location, but there will also be temporarily high rates in particular areas (due to aftershocks or near-repeat crimes) that may not be indicative of a longer-term structural difference. In order to determine the variables, there is an iterative process of estimation from a sample of events.

Table 1. Classes of predictive techniques (from Perry et al. (2013, p. 19)). Copyright RAND included with permission/

Classes of predictive techniques

Analytic category and primary application	Predictive technique	Class			
		Classical	Simple	Complex	Tailored
Hot spot analysis (where, using crime data only)	Grid mapping	X			X
	Covering ellipses	X			
	Kernel density	X			
	Heuristics		X		X
Regression methods (where, using a range of data)	Linear	X	X		
	Stepwise	X		X	
	Splines			X	X
	Leading indicators	X			X
Data mining (where, using a range of data)	Clustering	X		X	
	Classification	X		X	
Near-repeat (where, over next few days, using crime data only)	Self-exciting point process			X	
	ProMap			X	
	Heuristic		X		
Spatiotemporal analysis (when, using crime and temporal data)	Heat maps	X	X		X
	Additive model			X	
	Seasonality	X			
Risk terrain analysis (where, using geography associated with risk)	Geo-spatial predictive analysis			X	X
	Risk terrain modelling		X		X

PredPol® has generated a largely positive media response, some of which may be linked to a contractual requirement allegedly put on participating agencies to engage in promotional activities (Bond-Graham and Winston 2013, Cushing 2013). We are not arguing that PredPol® is better or worse than other software products deployed in predictive policing, and indeed could not do so without a proper evaluation of the different products provided by organisations such as IBM® (e.g. Blue CRUSH in Memphis and Blue PALMS in Miami), Information Builders® (Law Enforcement Analytics), Azavea® (HunchLab), SPADAC® (Signature Analyst®), Accenture®, and Hitachi®. However, its greater transparency makes it a useful example for illustrating assumptions about patterns of crime made by predictive policing software. Challenges for accountability presented by commercial confidentiality restricting transparency are discussed further in section 5.

3. Assumptions

All of the predictive tools mentioned in section 2 are based on assumptions about crime risk and its social, spatial, and temporal correlates. Even in a world of perfect surveillance, complexity implies that it would not be possible to predict future crime with certainty. The underlying model of predictive policing (Figure 1) relies on a host of technical, organisational, social, and ethical assumptions implicit in the successful operation of predictive policing. In the analysis below, we list the assumptions underlying the four stages of the predictive policing cycle: data collection, analysis, police operations, and criminal response.

3.1. Stage one: data collection

Assumption 1: Data used accurately reflect reality.

The crime data inputted into predictive policing software are not necessarily an accurate representation of criminal activity in a given area. It is necessarily limited by what individuals choose to report and what law enforcement officers directly observe. Further, the gap between crimes committed and crimes reported is not random but systemic (Harcourt 2007, Gassaway et al. 2011), linked to the type of crime (e.g. domestic violence vs murder), characteristics of victims, and differential availability and accessibility of data (e.g. social media compared to all communications).

Because the data are being used as the basis for law enforcement decision-making, victims excluded from data collection will often continue to marginalised and ignored by law enforcement (Lerman 2013).

Not only is it impossible to capture every 'crime' that is committed, but such data that are captured will not always be categorised accurately or consistently (Calders and Žliobaitė 2013, p. 48). Questions of whether an event reported to police constitutes a 'crime', how that 'crime' is classified or categorised, what thresholds exist for an event to be recorded, whether and how the seriousness of a crime is recorded are all matters of discretion (Barocas and Selbst 2016). The problem is exacerbated where data sets from multiple sources are combined in the analysis, as similarly described classifications may have different meanings in different organisations. As in the title to Gitelman's (2013) book '*Raw data' is an oxymoron* (see also Kitchin 2014).

The practice of predictive policing itself affects the data collected. The presence of more police officers in a particular neighbourhood may spur the recording of crime in that neighbourhood (Custers 2013, pp. 3–4). Predictions can accordingly become self-affirming. This feedback loop is self-perpetuating, potentially resulting in observed stability of crimes, locations and individuals monitored by police despite potential changes in the *actual* crimes being committed. Thus the predictive accuracy of an algorithm needs to be measured independently from an evaluation of the impact of that algorithm (which requires changes to police practice).

Where historic data embeds discrimination either in the underlying reality or in the collection of data, this 'fact' can become self-perpetuating. As a result, there is a risk that what might seem an objective process becomes a means of perpetuating historic discrimination or bias with the 'imprimatur of impartiality on the resulting decisions' (Barocas and Selbst 2016).

3.2. Stage two: data analysis

Assumption 2: The future is like the past.
Any projection of historic crime patterns into the future assumes a degree of continuity. This is not necessarily unreasonable, but the appropriateness of this assumption depends on the context. It may be more appropriate for some types of crimes (e.g. burglary) than for others (e.g. kidnapping) and the degree of continuity can be affected by changes in policy as well as social and cultural changes in a particular community (Chan and Bennett Moses 2016). Putting more police officers in hot spots can (and is intended to) reduce crime (due to police presence) and change crime rates in adjacent areas (Bowers *et al.* 2011). Again, because predictive policing itself affects what is being measured (crime), predictive accuracy cannot be measured simultaneously with programme effectiveness.

Even where predictive software models change over time, it will focus on particular kinds of change. For example the model used in Mohler *et al.* (2011) assumes that historical crime data are reflected in the future in two distinct ways. There are structural differences in the crime rates in different areas and at different times that are assumed to be constant (although the actual crime events will change). In addition, there are highly localised series or clusters of events related to the 'near repeat' theory of crime. Thus the model assumes that fluctuations in the background rate of crime over time are not spatially localised. It is thus important to understand the models on which particular software products rely. While Mohler *et al.* (2011) publicises a model that is likely the basis for the PredPol® platform, we cannot check that the same principles apply to its current version and have even less knowledge of other predictive policing software. Commercial confidentiality thus presents challenges for accountability of police decision-making (see section 5). However, the impact on predictive accuracy of simplified modelling *may* be small so that these models remain useful.

Assumption 3: Irrelevance of omitted variables.
Any analytic tool will focus on a limited, even if large, set of variables. In the case of predictive policing, variables will depend on the information routinely collected about crime which in turn depends on the available fields, how much is known about a given crime and what data are entered into the

system. Some variables are omitted because it was not realised *ex ante* that they would prove rel-evant for prediction or because they are expensive or difficult legally to procure. Not all variables con-sidered relevant in criminological theory are necessary included, and different software will rely on different variables.

The omission of relevant variables affects the accuracy of a predictive model. A machine learning algorithm may learn to classify based on a coarse generalisations when greater granularity would provide for better performance. For example, an algorithm may learn to generalise based on the basis of a general attribute such as 'young and male' in relation to dangerous driving when another variable, such as 'aggressive' (not captured directly) would be more precise (Calders and Žlio-baitė 2013, p. 52). The selection of datasets and variables is thus crucial, both to maximise predictive accuracy and in evaluating the fairness of discriminating based on particular characteristics. However, as with the driving example, predictions from available data can be useful even if not optimal.

Assumption 4: Algorithms are neutral.
Not only are there potential issues with the data set itself (as to comprehensiveness, accuracy, con-sistency and choice of variables) but the algorithm used to analyse data will have both advantages and disadvantages. Selection of the appropriate algorithm will depend on an existing theory or model (see Assumption 2). While different algorithms will have different limitations, it is important to understand that all will make assumptions, in some cases based on an assumed model of crime and in other cases based on general factors such as simplicity versus flexibility (with the associated potential of overfitting) and predictive power versus other goals such as comprehensibility, preser-vation of provenance or non-discrimination. Such bias is not problematic as such, but it does require awareness and testing to ensure the continued applicability of the assumptions underlying the choice of any particular prediction method. Further, the danger that an algorithm may be chosen specifically to optimise performance against particular evaluation metrics needs to be borne in mind in analysing an evaluation.

Assumption 5: Data analytics does not unjustly discriminate.
To understand this assumption, it is necessary to briefly consider what makes discrimination unjust. If predictive policing identifies a correlation between feature X and probability of offending, in what circumstances is it unjust to treat a person with feature X differently? Or, in the case of location-based tools, in what circumstances is it unjust to increase surveillance of neighbourhoods whose resi-dents share particular characteristics? Historically, concern around unjust discrimination is focussed on the targeting of individuals with particular types of characteristics (race, ethnicity, etc.) or locations populated by such individuals. While this remains important, and will be the focus here, it is also poss-ible that predictive tools will create new groupings of targeted individuals or places that are not associated with any historical category of discrimination (Zarsky 2013).

The first potential problem with discrimination is when the correlation is spurious – in this situ-ation such profiling is not only unjust but also ineffective. Such false correlations are often linked to the assumptions discussed above, as where crimes perpetrated by individuals of a particular race are more likely to be detected and treated as criminal.

More difficult is the question of when it is unjust to profile based on a characteristic that is indeed predictive of propensity to commit a crime. Schauer (2003, pp. 147–151, pp.189–190) argues that society may wish to prohibit even statistically justifiable discrimination (1) where as a general matter, such categories are more commonly the basis for inappropriate generalisation than for appro-priate generalisation such that a general prohibition on discrimination is justified, (2) on a related point, where the use of that generalisation is unfairly selective, as where there are other more reliable generalisations (see Assumption 3), and (3) where use of such generalisations in particular circum-stances would be stigmatising or produce excessive separation. The goal of predictive policing, which is to discriminate based on data-identified differences across individuals or locations, can raise concerns about unjust discrimination, particularly of the third type, even where there is no

intentional human animus towards a particular group and even where there are no other issues in data collection and analysis.

There are thus important issues to consider even where a sensitive characteristic (such as race associated with individuals or neighbourhoods) is truly predictive so that profiling becomes 'rational' (Harcourt 2007, Sklansky 2015). When it comes to historically sensitive information, such as the race of the offender or local population, Schauer's concerns, or those of Manning (2003, p. xiii, pp. 186–191) around democratic policing, may justify omitting a variable (such as race) either from the database generally or else from the fields used in the analytic process to make predictions even where it has predictive power.

However the simple solution of omitting sensitive variables is insufficient since the potential for differential treatment based on a sensitive characteristic (and thus potentially stigmatising effects) remains (Calders and Žliobaitė 2013). This is because there are many other variables that correlate, to varying degrees, with sensitive variables as in the case of 'red-lining' where geography is used as a proxy for race. At the same time, removing all variables that correlate with sensitive variables would reduce predictive accuracy significantly to the point that predictive analytics becomes imposs- ible (Calders and Verwer 2010, p. 279) because many such variables are relevant in any event, as in the case of the number of prior convictions. The appropriate balance between ensuring predictive accu- racy and protecting historically disadvantaged groups from the discriminatory impact of predictive policing is a difficult, and subjective, one.

While there are tools that can reduce the potentially discriminatory impact of algorithmic predic- tion, that can only be done via a 'positive' discrimination that may reduce predictive accuracy. For example, techniques described by Kamiran *et al.* (2013) involve altering the historical dataset, chan- ging a machine learning procedure by restricting the search space to non-discriminatory models and adjusting the models in post-processing. As explained by Verwer and Calders (2013, p. 255), without a causal model explaining why a particular characteristic is statistically relevant, the only choices are either (1) to avoid all variables that correlate with that characteristic, thus reducing accuracy, (2) accept that the algorithm may 'learn' to predict this effect (even if characteristic itself is not a variable in the database), or (3) change the data, algorithm or inference using discrimination-aware tech- niques by positively discriminating in favour of the relevant characteristic. Our point here is not to argue in favour of or against deploying anti-discrimination techniques; it is rather to emphasise that the use (or non-use) of such techniques is a controversial choice. As a result, it would be unwise to assume that predictive policing tools are inherently neutral so that the risk of perpetuating historical discrimination need not be considered in adopting a predictive policing solution.

To our knowledge, there is no study of the differential impact of predictive policing on particular groups. Measurement of such effects, and understanding the different ways in which such effects can occur, would be important in a full evaluation of predictive policing.

We are not suggesting that human decision-makers are less likely to perpetuate unfair discrimi- nation and thus preferable to the use of predictive policing software. Predictive analytics can be used to discover potentially discriminatory practices using association rule mining to identify decision rules in a dataset, filtering on the basis of legally grounded measures (Pedreschi *et al.* 2013), providing an important tool to those advocating for change. Even predictive analytics that perpetuates histori- cal discrimination may be preferable to human decision-making that is motivated by animus towards or unwarranted distrust of a particular group (Schauer 2003, p. 179, Zarsky 2003, p. 27). Thus decision- making based on analytics should be compared not to perfection, but rather to the human alternative (Zarsky 2003, p. 50). But it remains important to understand the impossibility of neutrality with respect to the way in which the existence of historic discrimination is treated in the analysis of his- torical data.

Assumption 6: Primacy of place.
Most prediction software focuses on *place* rather than *people* (Wilson *et al.* 2009, p. 7). As Uchida (2013, p. 1) points out, 'predictive policing is based on the premise that crime can be accurately

forecasted or predicted for small enough geographic areas and on reasonable enough time frames to enable police agencies to adequately deploy officers to potential problem areas'. While there are exceptions, such as in the case of approaches focussing more on identifying potential criminals and victims (as in Chicago, IL), semantic analysis of social media data to predict events (Wang *et al.* 2012) or analysis of physiological and behavioural cues for travellers at screening points, most predictive policing tools focus on the location of crime.

The focus on location is not new, and is similar to older 'hot spot' policing techniques, whereby police were encouraged to focus on locations with high historic crime rates (see generally, Braga and Weisburd 2010; Weisburd 2008, p. 2). The assumption that place is the crucial variable in predicting crime may be reasonable and is supported by a variety of empirical findings including patterns of burglary victimisation (Bowers *et al.* 2004), the impact of geographical features (Caplan and Kennedy 2010) and routine activities theory (Cohen and Felson 1979, Felson 2002). It is also subject to critique (as summarised in Manning 2010, pp. 168–173).

However, a focus on location is not equally appropriate for all crime types. As Sherman *et al.* (1989, p. 47) suggest, '[c]rimes arising out of intimate or market relationships may be much less dependent on place than predatory stranger crimes'; their research found that 'even predatory stranger offenses vary substantially by type of offense', for example, '[o]ne can avoid robbery twice as effectively by staying away from certain places than one can avoid sex crimes or auto thefts'. Hart and Zandbergen (2012, p. 58) found that crime type has an effect on predictive accuracy and for crimes such as homicide may have very low hit rates (from 1% to 4%). Thus an approach that focusses on place, as prevent in predictive policing practice, risks also focusing on only particular categories of crime. These limitations do not restrict the usefulness of predictive policing within the domains for which it is likely to be effective, however, to the extent that it diverts policing resources towards such domains, it may nevertheless have a negative impact overall.

3.3. Stage 3: police operations

Assumption 7: Targeting police deployment should be the primary intervention.
Perry *et al.*'s (2013, p. 14) model of predictive policing ideally includes not only *generic* interventions such as deployment of additional officers, but also *crime-specific* and *problem-specific* interventions. However, many predictive policing programmes envisage a single response to the forecasts made, namely that police deployments will change to focus on the locations at greater risk of future crime. In theory, predictive policing does not prevent other responses such as a problem-oriented response to predictions including 'recreational opportunities, mentor programs, community organising, lever-pull strategies, and youth sports programs' (Casady 2011, p. 2) and neighbourhood meetings are being used in jurisdictions such as Dallas in predicted hot spots. However, some predictive policing programmes are primarily about mobilising police patrols to particular blocks based on predictive analytics, not about attempting to understand the causes of crime in a particular area.

The ability to calculate (whether accurately or inaccurately) the geo-spatial distribution of future crime itself changes the context in which police strategies are formulated. Police responses in a hotspot policing approach tended 'to be pre-packaged, cookie-cutter reactions rather than tailored, researched strategic plans for solving or eliminating the problem over the long haul' (Rosenbaum 2006, p. 249). Strategic responses are even more difficult for predictive policing because predictive policing is explicitly about future crime, which is not necessarily limited to areas with high 'background' or structural crime rates for which a longer-term problem-oriented approach may be designed. Where the focus is on police deployments as the primary response, even a demonstrated reduction in crime rates does not imply that the most effective intervention has been implemented.

Assumption 8: Perfect implementation.
It cannot be assumed that law enforcement officers comply absolutely with intervention strategies informed by the prediction analysis. Even perfect predictions are only useful to the extent they

change police practice. Much will depend on the level of trust that police officers have in what they are being told. Sometimes police officers have knowledge not captured by the data (as where they know that the data they themselves enter into the system are flawed or incomplete) and may thus be less inclined to trust the forecasts or the analysts (see Cope 2004, p. 191). Perry *et al.* (2013, p. 129) have nominated police resistance as a barrier to the successful implementation of predictive policing.

In general, 'strong interest and enthusiasm among the staff involved' is regarded as an important characteristic that can improve the likelihood of successful police operations in response to predictions (Perry *et al.* 2013, p. 135). Assuming willingness to act, there also needs to be an ability to act effectively. Thus hot spots must be small enough to patrol effectively and there must not be too many of them (Perry *et al.* 2013). Similarly 'heat lists' of crime-prone individuals must not be overly long.

Hunt *et al.*'s (2014, p. xiii) evaluation of the Shreveport Predictive Policing Experiment, discussed below, certainly did not find perfect implementation of the prevention model:

> … treatment districts did not follow all aspects of the prevention model. Most important, the monthly planning meetings to set and maintain intervention strategies did not occur. These meetings were to be a key mechanism to ensure the prevention strategies were the same across police commands, and consequently to increase the statistical power needed for the impact analysis. Instead, the experimental districts made intervention-related decisions largely on their own … the strategies and levels of effort employed varied widely by district and over time.

Ethnographic studies on ILP have uncovered various organisational and cultural issues that may have led to unintended processes or consequences. For example, research in six Canadian police services found that the use of 'crime science' and analytic technologies to support ILP is more rhetorical than real (Sanders *et al.* 2015, p. 711).

3.4. Stage four: criminal response

Assumption 9: Changing police deployments prevents crime.
One assumption made in the use of predictive policing software is that changing police deployments prevents crime. There is credible empirical evidence that the use of hot spots policing can lead to lower crime levels. The most recent systematic review of the research evidence (based on 16 experimental or quasi-experimental studies) suggests that 'hot spots policing generates small but noteworthy crime reductions' (Braga *et al.* 2012, p. 633). However, there are variations in the size of the mean effect depending on the type of crime and the nature of police intervention (Braga *et al.* 2012, p. 653). When interventions are separated into 'problem-oriented policing' and 'increased traditional policing',[1] the meta-analysis shows that the former produced an overall mean effect that was twice the size of the latter (Braga *et al.* 2012, p. 653). While it is possible to combine predictive policing with a problem-oriented approach, there are challenges (see Assumption 7). The effects (measured using Cohen's *d*)[2] of increasing traditional policing activities alone are in fact quite small: 0.157 for violent crimes, 0.087 (not statistically significant) for property crimes, 0.139 for drug offences and 0.063 for disorder offences (Braga *et al.* 2012, p. 653). Weisburd *et al.* (2016, p. 47) have argued that even a small decrease in crime rate in a large city can bring 'meaningful social benefits' such as reduction in the costs of crime.

However, some scholars remain sceptical about the impact of changing police deployments on crime rates. Tonry (2014, p. 1), for example, argues that evidence from a historical analysis of the fall in crime rates throughout the Western world suggests that whatever the explanations may be for both long- and short-term declines in the rates of homicide, burglary, auto theft, and other property crimes since the 1990s, 'they do not include direct effects of changes in policing or sanctioning policies'. Eck and Maguire's (2000) review of 27 research studies on the effect of police strength on violent crime found mixed results. In relation to new policing strategies, they found 'little evidence that generic changes in policing are responsible for reducing violent crime', stronger evidence for 'focused' policing strategies to have contributed to the drop in violent crime, but uncertainties remain regarding the effectiveness of these strategies (2000, p. 245). At a broader level, some

researchers have argued that 'the police do stop some crime, although increasingly they will rely on non-police personnel for assistance in doing so' (Dixon 2005, p. 19).

The impact of policing strategy on crime is obviously a larger debate than can be fairly captured here. What is important for our purposes is that predictive policing programmes assume that crime can be reduced through changed police practices, most commonly changes in police deployments.

Assumption 10: The focus on crime is always appropriate.

Another assumption upon which predictive policing rests is that the focus on crime is always appropriate. Rosenbaum (2006, p. 247) asks this question directly in relation to hot spots policing, questioning why it is 'crime' that determines the hot spot rather than, for example, fear of crime or slum landlords which are also geographically focussed. Crime is not a simple category, so prioritising its prevention is not always optimal. Assuming it could be done, few would argue against the benefits of preventing murder. However, looking back in history, we now regard police focus on some 'crimes' (such as homosexuality) as oppressive and inappropriate. This should make us at least somewhat sceptical of the benefits of more perfect 'crime prevention', even if it were achievable. Dixon (2005, pp. 17–18) has cited examples of police role (e.g. in the policing of domestic violence, drug offences and overdoses, stop/search and public order) where an exclusive focus on crime prevention can conflict with considerations of equity, due process, public health, community relations, and order maintenance.

4. Evaluations of predictive policing

Predictive policing has been credited with substantial reductions in crime in some police agencies. For example, the PredPol® website reports drop in particular categories of crime in particular jurisdictions employing its software. Media and other sources also report percentage reductions in crime (e.g. Olesker 2012, Ibrahim 2013, Mitchell 2013, Turner *et al.* 2014), each time without evidence or references to published evaluations. There are also statements about the predictive accuracy of particular tools (e.g. Rich 2011), however, because of the assumptions above, this does not necessarily indicate the usefulness of a predictive policing programme.

Yet, as far as we are able to ascertain, very few formal evaluations of predictive policing have been conducted. In his review of predictive policing techniques Uchida (2014, p. 3878, *our emphasis*) has concluded that:

> The statistical techniques used in predictive analytics *are largely untested and have not been rigorously evaluated.* Risk Terrain Modeling, ProMap, neural networks, feature space, statistical learning models, the self-exciting point process, and predictive analytics developed by private vendors (IBM SPSS, SAS, IBI, and others) are yet to be fully tested in the field by independent evaluators.

Perry *et al.*'s (2013, p. 7, *our emphasis*) report on predictive policing in the US has similarly come to the conclusion that '*with a few exceptions, there have been no formal controlled evaluations* (though some were under way at the time of this writing)'. Perry *et al.*'s (2013, p. 123) interviews with police practitioners found 'very few said that they had evaluated the effectiveness of the predictions they produced or the interventions that followed the prediction'. This lack of emphasis on evaluation is one of five pitfalls,[3] in addition to the seven listed by Ridgeway (2013),[4] that the authors suggest must be avoided in order to make effective use of predictive policing.

What is available by way of published evaluation of predictive policing is very thin. Friend's (2013) report[5] of substantial crime reduction in Los Angeles Police Department (LAPD) lacks details and appears to be based on preliminary results. Nevertheless, we have been able to identify two systematic evaluations of predictive policing that provide some valuable insights into its effectiveness in crime reduction. In this paper we are focusing on evaluating the effectiveness of predictive policing in crime reduction and fair treatment, rather than evaluating the accuracy of the algorithm in predicting crime, even though verifying the claims of predictive policing tools is also important (Perry *et al.* 2013, p. 124). For the reasons explained above (particularly Assumption 1), evaluations of accuracy need to be conducted independently of evaluations of impact. Only evaluations of impact as to

effectiveness and fairness, however, take account of all assumptions (except Assumption 10, which cannot be tested empirically).

4.1. The Shreveport predictive policing experiment

An evaluation of the 2011 predictive policing experiment in Shreveport Police Department in Louisiana was conducted by Hunt *et al.* (2014). Funded by the US National Institute of Justice, it was both a process and an outcome evaluation. The Predictive Intelligence Led Operational Targeting (PILOT) program was an 'analytically driven policing strategy using special operations resources focused on narrow locations more likely to incur property crime' (Hunt *et al.* 2014, p. ix). The experiment was based on a clearly articulated set of parameters: programme theory, prediction model, prevention models, intermediate outcomes, and final outcomes.

The evaluation of outcomes was based on comparing three districts that made use of the predictive strategy and three control group districts that continued with existing policing approach. Predictions of future property crimes were based on multivariate logistic regression analyses for 400-foot grid cells to identify 'hot spots' where the probability of at least one property crime happening was predicted to be 40–60% (orange) or over 60% (red). Prevention of property crimes in the 'experimental' districts was based on strategies determined at monthly strategic planning meetings; they were not allowed to request operations from the citywide Crime Response Unit. In 'control' districts hot spots 'were derived from conventional crime mapping rather than from a predictive model' (Hunt *et al.* 2014, p. xi).

The evaluation study found no statistically significant difference in crime rates between the experimental and the control districts. The authors suggest that this null effect could be explained by three factors: the low statistical power of the tests, 'given the small number of experimental and control districts, as well as the low and widely varying crime counts per month and district in Shreveport' (Hunt *et al.* 2014, p. xiii); a failure of programme implementation as there were variations in the extent to which the prevention model was implemented between districts and over time; or a failure of programme theory, that is, the programme design was 'insufficient to generate crime reduction' (Hunt *et al.* 2014, p. xv).

4.2. Randomised controlled field trials of predictive policing in Los Angeles

Mohler *et al.* (2015)[6] describe two randomised controlled experiments they conducted to test their epidemic-type aftershock sequence (ETAS) crime forecasting tool (see Mohler *et al.* 2011). The first experiment was carried out in three divisions of the Los Angeles Police Department, while the second was conducted in two divisions of the Kent Police Department in the UK. Our discussion will focus only on the LA experiment as the trial in Kent was aimed at evaluating the predictive accuracy of the forecasting tool rather than the impact of predictive policing on crime.

The LA experiment involved comparing the predictions produced by the ETAS algorithm (treatment) with hot spots maps produced by dedicated crime analysts (control) in single-blind field trials in three LAPD Divisions from November 2011 to January 2013. Days were randomly assigned to treatment or control: 'each day all officers within an entire division received either ETAS predictions or analyst predictions that were identical in appearance save for the hot spot locations' (2015, p. 1400). The results were reported as follows:

> Our results show that ETAS models predict 1.4–2.2 times as much crime compared to a dedicated crime analyst using existing criminal intelligence and hotspot mapping practice. In the three divisions in Los Angeles, police patrols using ETAS forecasts led to an average 7.4% reduction in crime volume as a function of patrol time, whereas patrols based upon analyst predictions showed less than half of the treatment effect at a level that was not statistically significant (Mohler *et al.* 2015, p. 1400).

The study has demonstrated ecological validity in the sense that the 'experiment was conducted head-to-head with practitioners under all the pressures that confront crime analysts and police on a day-to-day basis' (2015, p. 1409). The extent to which the results are generalisable depends on

whether the policing practices at the LAPD Divisions in the study were representative of population of police officers, for which the authors rely on survey results in the US and the UK on police use of crime analysis in relation to larger police departments.

The authors acknowledged that predictive policing is a short-term strategy to disrupt criminal opportunities; it is not a replacement for long-term crime prevention strategies aimed at fixing more fundamental causes of crime.

These two evaluations of predictive policing, one using multivariate regression (Hunt et al. 2014) and the other using an ETAS model (Mohler et al. 2015) for predicting 'crime hotspots' are useful for understanding the prospects and limits of predictive policing in practice. Hunt et al.'s (2014) evaluation is particularly useful as it is multi-faceted, 'utilisation focused' (Patton 2008)[7] and 'realistic' (Pawson and Tilley 1997).[8] It is not only focused on final outcomes (reduction in crime and increase in arrests), but also on programme theory and its implementation. Mohler et al.'s evaluation is less concerned with implementation; the authors explain that '[s]ince each day was randomly assigned to control or treatment, variation in both policing tactics and patrol time was independent of experimental condition' (2015, p. 1401). Nevertheless, paying attention to implementation issues is an important aspect of evaluation, as the outcomes of predictive policing may depend on the quality of its implementation (see Assumption 8 in section 3). Given the null effect of the Shreveport Experiment and the positive findings of the LA trial, further rigorous evaluations of predictive policing would be required to build a firm evidence base. The independence of the evaluation is also a very important issue as it affects the credibility of the findings. In general, our understanding of the effectiveness of predictive policing would be advanced by additional independent and rigorous evaluations, not only in relation to hot spots, but also its impact on larger geographical areas.

5. Accountability

It is crucial that law enforcement agencies are accountable for their decisions. Accountability has two competing meanings in the context of policing: it is either concerned with *control* of the police or the requirement that police must *explain* their conduct – '[b]oth models of accountability seek to legitimate the conferring of extraordinary powers upon the police … by reassuring citizens that police are not out of control or their actions free from appropriate scrutiny' (Chan 1999, p. 253). More generally, accountability involves (1) the provision of answers (2) to others with a legitimate claim to demand an account, (3) with consequences (Bovens et al. 2014). Here we are concerned primarily with the challenges that predictive policing presents in relation to the first of these. We leave aside the question of the scope of accountability, particularly as regards *public* accountability, as this raises complex questions about competing objectives (see e.g. Luna 2000, Zarsky 2013, Ditchley Foundation 2015).

Predictive policing raises two difficulties for accountability in the sense of the giving of explanations: (1) the challenges of multi-agent responsibility for deployment decisions and (2) the challenges of explanation itself. As in the case of discrimination, we are not suggesting that existing practices involve full accountability – there are critiques of current accountability practices within law enforcement agencies (e.g. Manning 2010, pp. 69–70). Our goal is to explain the particular challenges raised by the new practice of predictive policing for police accountability, not to argue that it would be worse than current alternatives.

An example of the first problem occurs where a senior police officer is asked to account for deployment decisions made. That officer could refer to a decision to use particular software (approved by different officers) and an output given by that software on a particular day. However, this deflects rather than answers the inquiry. The challenge here is the temptation to outsource aspects of the decision-making process, and thus responsibility and accountability for the decision itself, to technological tools. If the tool is relevantly neutral, such as word-processing software used to record a human decision, this will rarely be a problem. However, as we have seen, predictive policing tools depend on multiple assumptions about which decision-makers themselves may be unaware. In Latour's (2005) terms, an important component on the decision-making process is outsourced to a

non-human 'actant', being the software itself. In such situations, there is a potential accountability gap whenever the software itself becomes an acceptable basis for decision-making so that those to whom account is given do not (or are unable to) dig deeper.

The second concern relates to the ability to dig deeper. An explanation that includes the software itself can either be based on an explanation of the operation of the software (requiring transparency and comprehensibility) or on an explanation of its effectiveness.

Transparency, defined as 'availability of information about an actor allowing other actors to monitor the workings and performance of this actor' (Meijer 2014, p. 511), is closely related to accountability (Pasquale 2015, p. 175). If there is comprehensible information as to the circumstances in which data are collected, the circumstances in which data may be altered (including data storage security protocols), the circumstances in which data are matched or combined (including statistical probabilities of incorrect matches and any systemic bias), the algorithms and processes that are used on the data (the 'source code'), as well as the assumptions and biases inherent in those, the output in its original form, and any limitations in the process by which any visualisations are rendered, then the person to whom the account is given has a comprehensive understanding of the reasons for any given decision. The decision, including reliance on the software, is then open to evaluation, critique, feedback, and formal sanction where relevant.

However, for a range of reasons, such full transparency and comprehensibility is rarely possible in predictive policing. Software may be subject to commercial-in-confidence provisions in licensing contracts. Even where the source code is available, there may be emergent properties of the algorithm (particularly in the case of machine learning) that cannot be anticipated. In addition, the person to whom an account is given may lack the expertise to deduce biases inherent in the choice of algorithm or process of analysis. Other assumptions discussed above, such as those inherent in processes of data collection, may be poorly understood so that there is little evidence to confirm or counter the assumptions inherent in the process. Thus, while full transparency and comprehensibility constitute an ideal form of accountability, one needs to consider alternatives that may be 'good enough' in practice.

A second and more fruitful alternative to transparency is to have proper evaluations of the software, not only for effectiveness in reducing (rather than relocating) crime, but also for differential impact on historically disadvantaged communities. A full evaluation framework was the subject of section 4 above. This would enable accountability by putting those making decisions in a position where they can describe the benefits of using a tool, despite not being able to explain *why* those the tool brings those benefits. It thus allows the decision-maker to justify the use of the software relied on to make a decision as both effective and non-stigmatising, thus providing a basis for accountable decision-making.

6. Conclusion: problems and prospects of predictive policing

This paper has provided an overview of predictive policing – what it is, how it is supposed to work, some of the analytic tools used, and three key issues that must be considered for this model of policing. First of all, we draw attention to the assumptions being made at every stage of the predictive policing cycle. We then highlight the importance of independent evaluations of this model of policing and propose evaluation designs that are not simply focused on outcomes. Finally, we discuss some of problems of accountability in predictive policing, and the need for some combination of transparency and evaluation in order to explain police deployment decisions.

Our goal is not to argue that predictive policing tools are bad or shouldn't be used. Such an argument would require at a minimum a proper comparison with other bases for law enforcement decision-making which may be equally or more flawed. Rather our goal is to understand predictive policing so that those using it, or considering it, are aware of its limitations as well as its potential, can build in opportunities for evaluation, and can construct knowledge-sharing and regulatory frameworks that preserve accountability. Predictive policing provides tools that have the potential to leverage developments in data science in order to improve the effectiveness and efficiency of law

enforcement agencies. However, this ultimately depends on humans understanding the limitations and assumptions embedded in the software and ensuring that ineffective and unjust outcomes are avoided.

As Citron and Pasquale (2014, p. 18) argue in the context of credit scoring, systems that have the capacity to contribute to or create the situation they claim merely to predict are a normative matter. As a society we have an interest in crime prevention and efficient policing, but we also have an interest in ensuring that law enforcement strategies, including deployment and surveillance decisions, are effective, fair, and just. This requires understanding, testing, and governance.

Notes

1. Braga et al.(2012, pp. 655–656) define 'problem-oriented policing' as 'police-led efforts to change the underlying conditions at hot spots that cause them to generate recurring crime problems' and 'traditional policing' as 'policing activities concentrated at specific places to prevent crime through general deterrence and increased risk of apprehension'.
2. As with other meta-analyses, Braga et al. (2012, p. 639) used 'the standardized mean difference effect size (also known as Cohen's d, see Cohen, 1988)' to measure the effect of the intervention on outcomes. Cohen's d is the difference between two means divided by a standard deviation for the data. A value of d below 0.2 is regarded as a small effect size.
3. The other four pitfalls are: focusing on prediction accuracy instead of tactical utility, relying on poor-quality data, misunderstanding the factors behind the prediction, and overlooking civil and privacy rights.
4. See Perry et al. (2013, p. 121, fn 5) quoting Ridgeway (2013).
5. Friend is apparently not an independent commentator but an affiliate of PredPol, see Bond-Graham and Winston (2013). See also Cushing (2013).
6. The article acknowledges that two of the authors (Mohler and Brantingham) are co-founders of PredPol®, and five of the seven authors hold stock in PredPol®.
7. Utilisation-focused evaluation is 'evaluation done for and with specific intended primary users for specific, intended uses'.
8. Realistic evaluation emphasises the importance of programme evaluation not to focus solely on outcomes, but to understand contexts and mechanisms (theory), so that the question is not simply 'what works?' but 'what works for whom in what circumstances?'.

ORCID

Lyria Bennett Moses 🔗 http://orcid.org/0000-0003-2571-1469

References

Barocas, S. and Selbst, A.D., 2016. Big Data's disparate impact. California law review, 104, 671.
Beck, C. and McCue, C., 2009. Predictive policing: what can we learn from WalMart and Amazon about fighting crime in a recession. The police chief, 76 (11), 20–29.

Bovens, M., Schillemans, T. and Goodin, R.E., 2014. Public accountability in MAP. *In*: M. Bovens, R.E. Goodin and T. Schillemans, eds. *The Oxford handbook of public accountability*. Oxford: Oxford University Press, 1–18.

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

Bowers, K.J., *et al*., 2011. Spatial displacement and diffusion of benefits among geographically focused policing initiatives: a meta-analytical review. *Journal of experimental criminology*, 7, 347–374.

Braga, A.A. and Weisburd, D.L., 2010. *Policing problem places: hot spots and effective prevention*. New York: Oxford University Press.

Braga, A.A., Papachristos, A.V., and Hureau, D.M., 2012. The effects of hot spots policing on crime: an updated systematic review and meta-analysis. *Justice quarterly*, 31 (4), 633–663.

van Brakel, R. and De Hert, P., 2011. Policing, surveillance and law in a pre-crime society: understanding the consequences of technology based strategies. *Journal of police studies*, 20 (3), 163–192.

Calders, T. and Verwer, S., 2010. Three naive Bayes approaches for discrimination-free classification. *Data mining and knowledge discovery*, 21 (2), 277–292.

Calders, T. and Žliobaitė, I., 2013. Why unbiased computational processes can lead to discriminative decision procedures. *In*: B. Custers, T. Calders, B. Schermer and T. Zarsky, eds. *Discrimination and privacy in the information society: data mining and profiling in large databases*. Heidelberg: Springer, 43–60.

Caplan, J.M. and Kennedy, L.W., 2010. *Risk terrain modelling manual: theoretical framework and technical steps of spatial risk assessment for crime analysis*. Newark, NJ: Rutgers Center on Public Security.

Casady, T., 2011. Police legitimacy and predictive policing. *Geography and public safety*, 2 (4), 1–2.

Chan, J., 1999. Governing police practice: limits of the new accountability. *British journal of sociology*, 50 (2), 251–270.

Chan, J. and Bennett Moses, L., 2016. Is Big Data challenging criminology? *Theoretical criminology*, 20 (1), 21–39.

Citron, D.K. and Pasquale, F.A., 2014. The scored society: due process for automated predictions. *Washington law review*, 89 (1), 1–33.

Cohen, J., 1988. *Statistical power analysis for the behavioral sciences*. 2nd ed. Hillsdale, NJ: Lawrence Erlbaum.

Cohen, L. and Felson, M., 1979. Social change and crime rate trends: a routine activity approach. *American sociological review*, 44 (4), 588–608.

Cope, N., 2004. 'Intelligence led policing or policing led intelligence?': Integrating volume crime analysis into policing. *British journal of criminology*, 44 (2), 188–203.

Custers, B., 2013. Data dilemmas in the information society: introduction and overview. *In*: B. Custers, T. Calders, B. Schermer and T. Zarsky, eds. *Discrimination and privacy in the information society: data mining and profiling in large databases*. Heidelberg: Springer, 3–26.

Ditchley Foundation. *Conference on intelligence, security and privacy*, 14–16 May 2015, Ditchley Park, UK. Available from: http://www.ditchley.co.uk/conferences/past-programme/2010-2019/2015/intelligence [Accessed 8 February 2016].

Dixon, D., 2005. Why don't the police stop crime? *Australian & New Zealand journal of criminology*, 38 (1), 4–24.

Eck, J.E. and Maguire, E.R., 2000. Have changes in policing reduced violent crime? An assessment of the evidence. *In*: A. Blumstein and J. Wallman, eds. *The crime drop in America*. Cambridge: Cambridge University Press, 207–265.

Felson, M., 2002. *Crime and everyday life: insight and implications for society*. 3rd ed. Thousand Oaks, CA: Sage.

Gassaway, B., Armon, S., and Perez, D., 2011. Engaging the community: operation heat wave. *Geography and public safety*, 3 (1), 8–13.

Gitelman, L., ed., 2013. *'Raw Data' is an oxymoron*. Cambridge, MA: MIT Press.

Harcourt, B.E., 2007. *Against prediction: profiling, policing and punishing in an actuarial age*. Chicago, IL: University of Chicago Press.

Hart, T.C. and Zandbergen, P.A., 2012. *Effects of data quality on predictive hotspot mapping: final technical report*. Washington, DC: National Institute of Justice, 239861.

Hunt, P., Saunders, J., and Hollywood, J.S., 2014. *Evaluation of the Shreveport predictive policing experiment*. Santa Monica, CA: RAND.

Kamiran, F., Calders, T., and Pechenizkiy, M., 2013. Techniques for discrimination-free predictive models. *In*: B. Custers, T. Calders, B. Schermer, and T. Zarsky, eds. *Discrimination and privacy in the information society: data mining and profiling in large databases*. Heidelberg: Springer, 223–240.

Kitchin, R., 2014. *The data revolution: Big Data, open data, data infrastructures and their consequences*. Thousand Oaks, CA: Sage.

Latour, B., 2005. *Reassembling the social: an introduction to actor-network-theory*. New York: Oxford University Press.

Lerman, J., 2013. Big Data and its exclusions. *Stanford law review* [online], 66. Available from: http://www.stanfordlawreview.org/sites/default/files/online/topics/66_stanlrevonline_55_lerman.pdf [Accessed 2 February 2016].

Lum, C., Koper, C.S., and Telep, C.W., 2011. The evidence-based policing matrix. *Journal of experimental criminology*, 7 (3), 3–26.

Luna, E., 2000. Transparent policing. *Iowa Law Review*, 85 (4), 1107–1194.

Maguire, M., 2000. Policing by risks and targets: Some dimensions and implications of intelligence-led crime control. *Policing and Society*, 9 (4), 315–336.

Manning, P.K., 2010. *Democratic policing in a changing world*. Boulder: Paradigm.

Meijer, A., 2014. Transparency. In: M. Bovens, R.E. Goodin, and T. Schillemans, eds. The Oxford handbook of public account-ability. Oxford: Oxford University Press, 507–521.

Mohler, G.O., et al., 2011. Self-exciting point process modeling of crime. Journal of the american statistical association, 106 (493), 100–108.

Mohler, G.O., et al., 2015. Randomized controlled field trials of predictive policing. Journal of the American statistical association, 110 (512), 1399–1411.

Pasquale, F., 2015. The black box society: the secret algorithms that control money and information. Cambridge, MA: Harvard University Press.

Patton, M.Q., 2008. Utilization-focused evaluation. 4th ed. Saint Paul, MN: Sage.

Pawson, R. and Tilley, N., 1997. Realistic evaluation. London: Sage.

Pedreschi, D., Ruggieri, S., Turini, F., 2013. The discovery of discrimination. In: B. Custers, T. Calders, B. Schermer, and T. Zarsky, eds. Discrimination and privacy in the information society: data mining and profiling in large databases. Heidelberg: Springer, 91–108.

Perry, W.L., et al., 2013. Predictive policing: the role of crime forecasting in law enforcement operations. Santa Monica, CA: RAND.

Police Executive Research Forum, 2014. Future trends in policing. Washington, DC: Office of Community Oriented Policing Services.

Ratcliffe, J.H., 2012. Intelligence-led policing. Cullompton: Willan.

Ridgeway, G., 2013. The pitfalls of prediction. National institute of justice journal, 271, 34–40.

Rosenbaum, D., 2006. The limits of hot spots policing. In: D. Weisburd and A. Braga, eds. Police innovation: contrasting perspectives. New York, NY: Cambridge University Press, 245–266.

Sanders, C.B., Weston, C. and Schott, N., 2015. Police innovations, 'secret squirrels' and accountability: empirically study-ing intelligence-led policing in Canada. British journal of criminology, 55 (4), 711–729.

Schauer, F., 2003. Profiles, probabilities and stereotypes. Cambridge, MA: Harvard University Press.

Sherman, L.W., Gartin, P.R., Buerger, M.E., 1989. Hot spots of predatory crime: routine activities and the criminology of place. Criminology, 27 (1), 27–56.

Sklansky, D.A., 2015. Police and democracy. Michigan law review, 103, 1699–1830.

Tonry, M., 2014. Why crime rates are falling throughout the western world. Crime & justice, 43 (1), 1–63.

Uchida, C., 2013. Linking theory to practice: testing geospatial predictive policing in a medium-sized police agency – A pro-posal submitted to the National Institute of Justice. Unpublished document. Cited with permission.

Uchida, C., 2014. Predictive policing. In: G. Bruinsma, D. Weisburd, eds. Encyclopedia of criminology and criminal justice. New York: Springer, 3871–3880.

Verwer, S., Calders, T., 2013. Introducing positive discrimination in predictive models. In: B. Custers, T. Calders, B. Schermer, and T. Zarsky, eds. Discrimination and privacy in the information society: data mining and profiling in large databases. Berlin: Springer, 255–272.

Wang, X., Gerber, M.S. and Brown, D.E., 2012. Automatic crime prediction using events extracted from Twitter posts. In: J. Salerno, S.J. Yang, D. Nau, S.K. Chai, eds. Social computing, behavioral-cultural modeling and prediction. Berlin: Springer, 231–238.

Weisburd, D., 2008. Place-based policing. Ideas in American policing, 9, 1–16.

Weisburd, D., et al., 2016. Do stop, question, and frisk practices deter crime? Evidence at microunits of space and time. Criminology & public policy, 15 (1), 31–56.

Wilson, R.E., et al., 2009. Geospatial technology working group meeting report on predictive policing. Scottsdale, AZ: US Department of Justice.

Zarsky, T.Z., 2003. 'Mine your own business!' Making the case for the implications of the data mining of personal infor-mation in the forum of public opinion. Yale journal of law & technology, 5 (1), 1–56.

Zarsky, T.Z., 2013. Transparent predictions. University of Illinois law review, 2013 (4), 1503–1569.

Media articles and web-based sources

Adams, G., 2012. LAPD's sci-fi solution to real crime. Independent, 11 January, p. 32.

Bond-Graham, D. and Winston, A., 2013. All tomorrow's crimes: The future of policing looks a lot like good branding [online]. SF Weekly. Available from: http://www.sfweekly.com/sanfrancisco/all-tomorrows-crimes-the-future-of-policing-looks-a-lot-like-good-branding/Content?oid = 2827968 [Accessed 2 February 2016].

Olesker, A., 2012. White paper: Big Data solutions for law enforcement [online]. CTOlabs.com. Available from: http://ctolabs.com/wp-content/uploads/2012/06/120627HadoopForLawEnforcement.pdf [Accessed 2 February 2016].

Cushing, T., 2013. 'Predictive policing' company uses bad stats, contractually-obligated shills to tout unproven 'successes' [online]. Techdirt. Available from: https://www.techdirt.com/articles/20131031/13033125091/predictive-policing-company-uses-bad-stats-contractually-obligated-shills-to-tout-unproven-successes.shtml [Accessed 2 February 2016].

Friend, Z., 2013. Predictive policing: Using technology to reduce crime [online]. FBI Law Enforcement Bulletin. Available from: https://leb.fbi.gov/2013/april/predictive-policing-using-technology-to-reduce-crime [Accessed 2 February 2016].

Ibrahim, M., 2013. Crunching numbers to catch bad guys; UCLA researchers develop mackie algorithm to predict crime. *Edmonton journal*, 11 May 2013, p. C3.

Mitchell, R., 2013. *Predictive policing gets personal* [online]. Computerworld. Available from: http://www.computerworld.com/article/2486424/government-it/predictive-policing-gets-personal.html [Accessed 2 February 2016].

PredPol® website results section: http://www.predpol.com/results/ [accessed 18 December 2014] – product promotion.

Rich, S., 2011. *Predictive policing reduces crime in Santa Cruz, Calif.* [online]. Government Technology. Available from: http://www.govtech.com/public-safety/Predictive-Policing-Project-Reduces-Crime-Santa-Cruz-Calif.html [Accessed 2 February 2016].

Turner, G., Brantingham, J., and Mohler, G., 2014. *Predictive policing in action in Atlanta, Georgia* [online]. The Police Chief. Available from: http://www.policechiefmagazine.org/magazine/index.cfm?fuseaction=display&issue_id=52014&category_ID=4 [Accessed 2 February 2016].