Why Current Predictive Policing Methods Simply Cannot Work

Rick van der Kleij, Selmar Smit, Hans van Vliet, Freek van Wermeskerken

Predictive policing is about planning specific (geographical) policing activities based on predictions created by data analytics, whereas traditional policing operations rely on human risk assessment. Sophisticated algorithms to crunch big data will enable the police to predict crimes. Predictive policing enables the police to take preventive measures to mitigate or reduce crime risks, e.g. by being present at locations where crime probability is the greatest. An increasing number of police departments are now adopting it, although there are still questions about its effectiveness.

This paper focuses on the data and algorithms in the predictive policing feedback loop and shows how the current approaches and algorithms actually have two fundamental flaws, namely the 'reporting bias' and one we named the 'back-to-the-future paradox'. We propose a new approach for the predictive policing feedback loop to counter this flaw and show by means of simulation that it is able to overcome these flaws.

1. INTRODUCTION

Predictive policing is about planning specific (geographical) policing activities based on predictions created by data analytics, whereas traditional policing operations rely on human risk assessment. Sophisticated algorithms to crunch big data will enable the police to predict crimes (Smit et al., 2016). Predictive policing enables the police to take preventive measures to mitigate or reduce crime risks (Saunders et al., 2016), e.g. by being present at locations where crime probability is the greatest. Predictive policing was fielded in 2008 at the LA Police Department (Greengard, 2012, Perry et al., 2013).

An increasing number of police departments are now adopting it, although there are still questions about its effectiveness. Mohler et al. (2015) claim that near real-time Epidemic Type Aftershock Sequence (ETAS) crime forecasting models predict up to two times better than dedicated crime analysts, reducing crime on the basis of predictive policing algorithms (under realistic law enforcement resource constraints). However, the results and conclusions of their experiment seem ambiguous (Section 5).

Predictive policing is still relatively new. There are ethical and social aspects, e.g. related to systematic bias in predictions and transparency and accountability of algorithm-generated decisions. Bennett Moses and Chan (2016) listed the potential problems connected to predictive policing, including systematic bias in police-recorded data sets (risking discriminatory policing). Also, even accurate predictions do not necessarily reduce crime. Predictions may not be actionable if the time or location is not sufficiently specific (Saunders et al., 2016), or police departments ignore or do not intervene correctlys, or even use the wrong interventions for the type of predictions. Mohler et al. (2015) conclude that more research is needed on the relationship between police activity on a specific location at a specific time and the occurrence of crime: “patrol times and locations may be important variables to include in predictive policing models to close the feedback loop between officers, criminals and the algorithms, linking them together” (p. 32).

This paper focuses on the data and algorithms in the predictive policing feedback loop. We show that the current approaches and algorithms actually have a fundamental flaw. Our new approach to this form of predictive policing shows what exactly needs to change to improve the effects of algorithmic predictions. This will also improve the allocation of scarce available resources (officers, cars, etc.) to police interventions. The ethical aspects or faulty assumptions were already described by Bennet Moses & Chan (2016).

First, we introduce current predictive policing algorithms. Second, we zoom in on the issues with forecasting algorithms through the use of data and how these problems increase by actually using the predictions for operations. Third, we propose a fundamental new approach to predicting crime, showing how it behaves differently on biased data. Finally, we sketch the road ahead.

2. CURRENT PREDICTIVE POLICING ALGORITHMS

Predictive policing algorithms are commonly based on a combination of diverse sources of information, e.g. environmental/demographic, date-/time-related, and most importantly, information on past crimes. For example, the simplest forecasting model imaginable states that the predicted number of incidents in a geographical area (box) equals last week’s average number of reported incidents in this area. This model may actually work already fairly well, but is rather limited.

The number of burglaries in a given area is equal to 2 times the number of burglaries last week, minus 1 time the number of burglaries exactly two weeks ago.

This model can include a very simple trend: if there were 10 burglaries last week, and 8 the week before, the new forecast will be 12. Adding more data results in fairly complex formulas. The challenge is then how to determine the weighting factor of each variable. In the previous example the weighting factors (2 and 1) were easy to determine, since they have a specific meaning. But how do you determine the following formula?

The number of burglaries in this box next Monday is 0.257 times the humidity, plus 1.56 times the number of burglaries last week, minus 0.46 times the number of burglaries normally committed on a Monday, plus 0.12 times the number of burglaries in an adjacent box last week.

Such weighting factors can be estimated using historical data. If you know the exact last year's humidity at a specific time, the reported number of burglaries in the previous week, the normal Monday number of burglaries and the number of burglaries in adjacent areas, you can determine the specific weight of each variable, resulting in the smallest 'mean squared error'. One complex formula could fill an entire book. However, the terms related to past crimes are generally the most important; they are essentially trend analyzers. Hence, these formulas estimate the likelihood of the occurrence of an incident at a specific time and location, based on the trend in past reported crimes.

PredPol, HunchLab, PRECOBS, The Daily Forecast, Crime Anticipation System (CAS) and many others (Smit, 2015) use these, or similar approaches.

3. REPORTING BIAS

To be able to prevent crime we need to know the crime potential, the amount of crime that would happen without barriers to offenders. This is the result of a complex interaction of diverse factors coinciding at the point of crime commission (Natarajan, 2011). Many factors influence crime potential. They may differ per crime type. Some are more remote, e.g. the education of a potential offender. Others are more proximal, e.g. the presence of a motivated offender in a suitable situation to commit crime. Basically, a criminal event happens with the right conjunction of criminal opportunity.

A problem with current predictive policing software is the use of police data for predicting crime. Predictions are not based on the true crime potential, but rather on records of crimes reported (see also Bennett Moses & Chan, 2016). For some crime types (e.g. burglary) the reported crime rate fairly matches the total amount of committed crime. However, for other crime types (e.g. bike theft) the reported crime and committed crime probably diverge. People often choose not to report such petty crimes: most bikes are cheap and the chance of getting it back by reporting the theft is quite low, the process of reporting often cumbersome. Another complication is that the reporting of crimes varies across geographical areas; most software assumes this variable is constant.

Furthermore, the presence of more police officers in a particular area or hot spot may spur the recording of crime in that area. Policing practice could stimulate the public to report committed crimes, for instance, because people feel safer then, and less afraid of repercussions.

This is particularly true when data is gathered as byproduct of policing activity, e.g. in case of drug dealers or drunk driving. Most records on such phenomena will be generated by the police itself. On spots where no police officer has been no crime is reported either. But that does not mean that no crime has been committed. (Dutch) Data scientists at police agencies are often aware of this (Rienks, 2015), but the question rises what would happen if the usage of predictive policing becomes more common and is tailored to local needs.

So the data fed into predictive policing software does not reflect reality (Bennett Moses & Chan, 2016). Moreover, the amount probably varies per crime type and geographical area, creating biased predictions.

There is more. Police data does also not include crimes prevented. The crime potential for a certain location could be high, but due to good policing or other factors, e.g. citizen patrols, crime is prevented. An area under police scrutiny, which therefore has a low number of reported crimes, will be underrepresented in the prediction model.

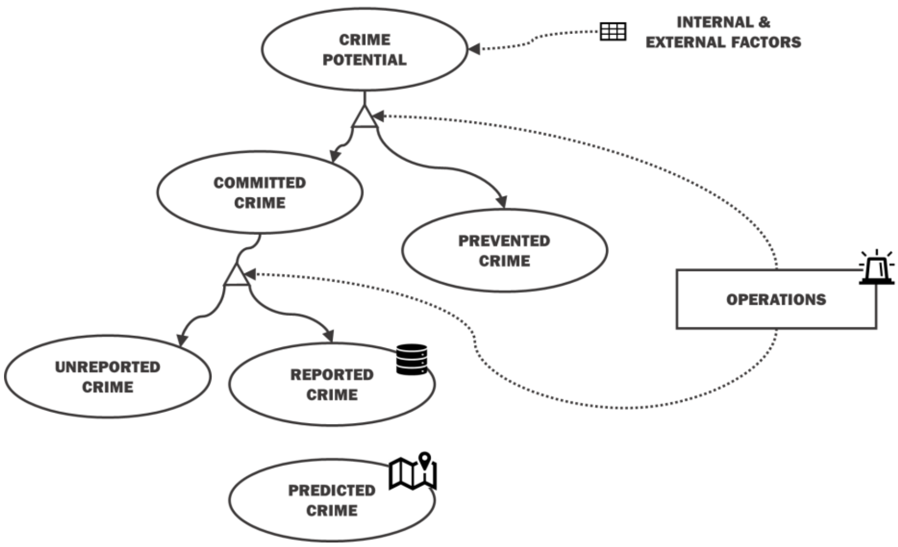

Current predictive policing approaches based on reported crime are therefore necessarily limited. They base predictions on an unrealistic set of data that does not include unreported and prevented crime, with strong geographical variations. In other words, the data used for making predictions does not resemble the true crime potential for a given location and is consequently not very accurate, for it (1) focuses on reported crime (not unreported crime); (2) does not consider prevented crimes, and (3) does not take into account the (geographical) variations that influence 1 and 2. Predictions in themselves can therefore not be very accurate in predicting criminal activity in a given area.

4. BACK-TO-THE-FUTURE PARADOX

An important precursor for criminal events is the absence of crime preventers, either real or anticipated (Natarajan, 2011). Crime preventers can intervene during the event, e.g. when the alarm sounds of an anti-theft system or when driving by a house that is being burgled. The practice of predictive policing itself affects the data collected. When we start using predictions for directing operations and interventions, we create a feedback loop:

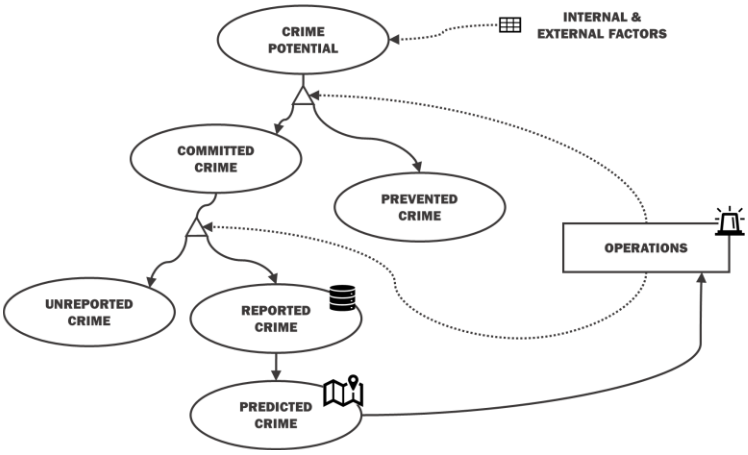

To wit, reported crime is used to predict crime; predicted crime is used to direct operations, and operations affect the number of reported crimes. In control theory such feedback loops are known to cause oscillations or chaotic system behaviour: the effects of a small disturbance on a system include an increase in the magnitude of the perturbation (Zuckerman & Jefferson, 1996).

Rather than correcting for this positive feedback loop in police data, current algorithms reinforce these biases by making flawed predictions possibly amounting to discriminatory policy (see also Bennett Moses & Chan, 2016). Thus, predictions can accordingly become self-affirming, leading to a self-perpetuating feedback loop, potentially resulting in observed stability of crimes, locations and individuals monitored by police despite potential changes in the actual crimes being committed that were invisible to the absent officers (Bennett Moses & Chan, 2016).

However, an often overlooked feedback loop is that policing activities concentrated at specific spots may prevent crime through deterrence, thereby invalidating the predictions: the back-to-the-future paradox. The paradox is that predictive policing itself affects the data collected and will render predictions to be incorrect. In time travelling, when you change your behaviour based on knowledge about the future, you create a temporal paradox. In other words, it is impossible to know if an officer prevented a crime or wasted time at a safe location. Likewise, if predictions turn out to be correct, then interventions weren't working, otherwise the officer that was sent would have prevented them. This effect is not only limited to the geographical box the operations were concentrated on, but due to deterrence and displacement effects surrounding locations may either have a decreased or increased crime potential, depending on the distance and the motivation of offenders.

This feedback loop through ‘prevented crime’ will cause oscillations in the predictions. An area with a high crime potential will have a high predicted number of crimes, causing police operations to be concentrated in that area. As a response, criminals will move to neighbouring locations, raising the crime potential in those areas, while the number of reported crimes in the original area will drop massively. After a week this number will be so low that the algorithm will direct all officers to the other locations, so that criminals will return. This is not a cat-and-mouse game, but a dog chasing its own tail.

Identifying the magnitude of these feedback loops is difficult. Reporting bias, and deterrence and displacement effects will be difficult to assess. Displaced crime will often fall outside the areas and crime types being studied or be so dispersed as to be masked by global trends, weather and other external factors influencing the behaviour of criminals. A pre-test post-test randomised controlled experimental design would be required (Robson & McCartan, 2016) to isolate this effect. Crime locations that received hot spot policing interventions need to be compared to similar locations receiving no policing at all for a prolonged period of time. Due to legal and ethical considerations, however, it is impossible to deny areas from policing interventions just for the purpose of research.

5. EVALUATION STUDIES

If all current approaches are flawed, how did predictive policing get so popular? First, there is an almost religious belief in the power of big data, even if solid evaluation of predictive policing is difficult, as indicated by the 2014 RAND report (Perry, 2013) on the use of predictive policing in Shreveport. It stated that this might e.g. be due to incomparable control areas and treatment groups, ineffective interventions, etc., not because it simply wouldn’t work. Concluding: it is more convenient to find reasons why the evaluation is flawed than to do the same for the algorithm.

Second, it appears to work in other experiments in which current predictive policing methods reached much higher levels (Moher, 2015) than human analysts. But was this a fair competition? Most likely the analysts were asked to predict future crime (potential), whereas the prediction accuracy was evaluated based on reported crime. Since the algorithm was trained to predict crime reports rather than crime potential, its higher accuracy is a logical outcome. Also, the areas selected by the analysts were arguably very similar to the ones normally used to direct the default policing routine, so crime was already suppressed in those areas, influencing prediction accuracy. Therefore, a high prediction accuracy is not an indication of effectiveness.

One could argue that the opposite holds true when comparing traditional policing methods with predictive policing. If predictions are used to direct interventions a lower prediction accuracy is to be expected. Crimes are prevented, and therefore the predictions are wrong. Brantigham, Moher et al. (Moher, 2015) found evidence for this shortfall, stating that when the police started using their predictions for interventions, prediction accuracy dropped. When using the predictions of analysts, no noticeable effect was found. This should have alerted the authors, but it did not. Police activity does have an effect, and the 'silent' prediction accuracy was larger than 0, hence crime should have been prevented, and prediction accuracy should have gone down. The only possible explanation is that policing those areas in 'available time' had no effect because they were already saturated with regular police activity. Another indication that the analysts predicted crime potential, not committed crimes.

This also indicates that reporting bias and back-to-the-future-paradox can have a great impact, especially when the areas identified by predictive policing algorithms are used for planning regular activity instead of 'available time'. This requires an algorithm that is immune to them.

6. A NEW APPROACH

An important first step towards breaking the feedback loop causing instability in predictive policing is making predictions differently. Currently, data collection methods focus on the wrong criminal response: reported criminal acts are fed in the algorithms. However, predictive policing is not about predicting the number of files, it is about preventing crime: hence, it is about unobserved criminal acts, or, to put it differently, about the deterrent effects of policing. We propose to use the much more stable crime potential for predictions rather than reported crimes. However, we cannot determine the crime potential without knowing how much crime is prevented, and how much the reported crime is affected by operations. And we cannot determine those without knowing the ground-truth of crime potential. Therefore, to break the feedback loop we need the ground-truth of crime potential:

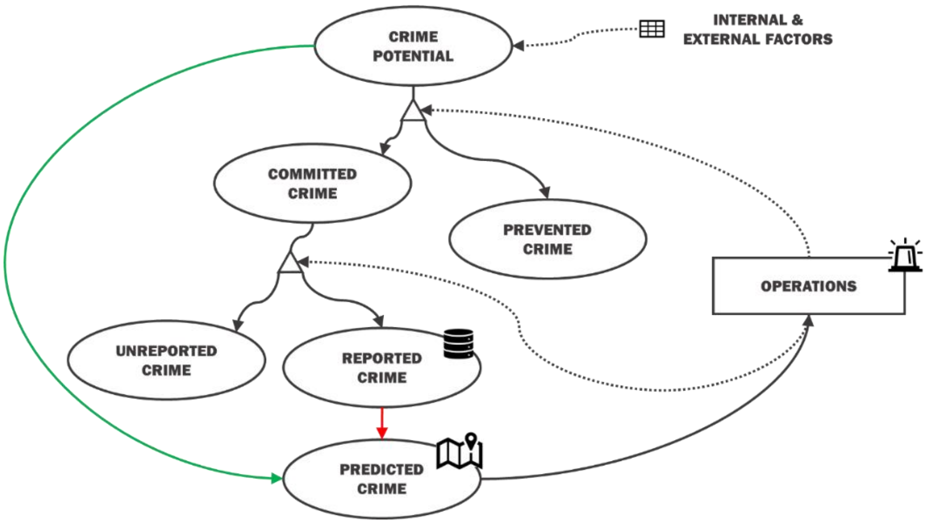

The new model incorporates the effects of interventions and thereby is able to estimate a ground-truth of crime potential. By taking into account the interventions and estimating their effects, the model is no longer biased by the unknown police interventions (I).

Current predictive policing algorithms model potential crime as a function of reported crime. Our approach differentiates between potential crime (*P*), committed crime (C) and reported crime (*R*):

$$P(t,l) = f_1(P(\tilde{t},l),\tilde{t},l), \text{with } \tilde{t} < t,$$

$$C(I,t,l) = f_2(P(t,l),I,t,l),$$

$$R(I,t,l) = f_3(C(t,l),I,t,l).$$

This looks complicated. However, f_1, f_2 and f_3 are just generic functions that indicate the (temporal) relation between the current potential and the past potentials (f_1) in time (t) and space (l), the relation between potential and committed (f_2) or committed and reported (f_3). For example, tomorrow's crime potential (which is a Monday) has some relation with the potential crime on other Mondays. And similarly, the crime potential in this specific area has some relation with neighbouring boxes. Other possible functions are provided later, but they can easily be replaced according to own ideas.

The basic idea is that potential crime is only influenced by time and location, and changes only gradually, because potential crime is essentially the overlapping area of '*likely offender*' and '*suitable target*' in the Routine Activity Theory (Natarajan, 2011).

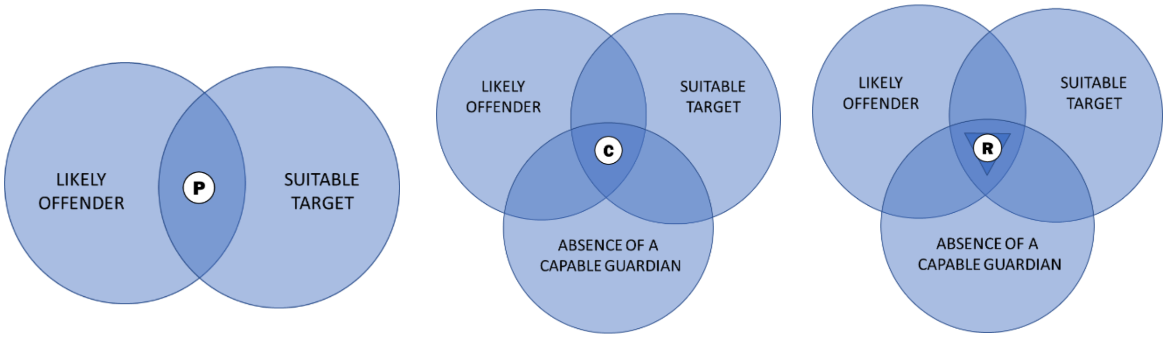

Figure 1: Routine Activity Theory on Crime Potential (P), Committed Crime (C) and Reported Crime (R)

The offenders taken into account here operate in a specific area. Unless they move, or are locked up, the number of 'likely offenders' is relatively stable. The same holds for 'suitable targets'; people will not move instantaneously, nor will they change their suitability (e.g. by buying new locks against burglars) in large quantities in a short period of time. Hence, f_1 will be some trend analysis of previous ($\tilde{t}$) potential crimes, but can also include month-, day- or time-specific terms, such as the average historical crime potential on Mondays.

Prescriptive Policing

Besides improved predictions on criminal activity, the new approach has an important additional benefit, namely f_2, the formula that indicates the effect of the vicinity of a capable guardian, also indicates the effectiveness of each intervention, as this is the difference between potential and committed crimes. Hence, this formula can be used to optimise operations. For example, if we know that in a specific area the effectiveness (in percentage) of 'patrol by foot' is this:

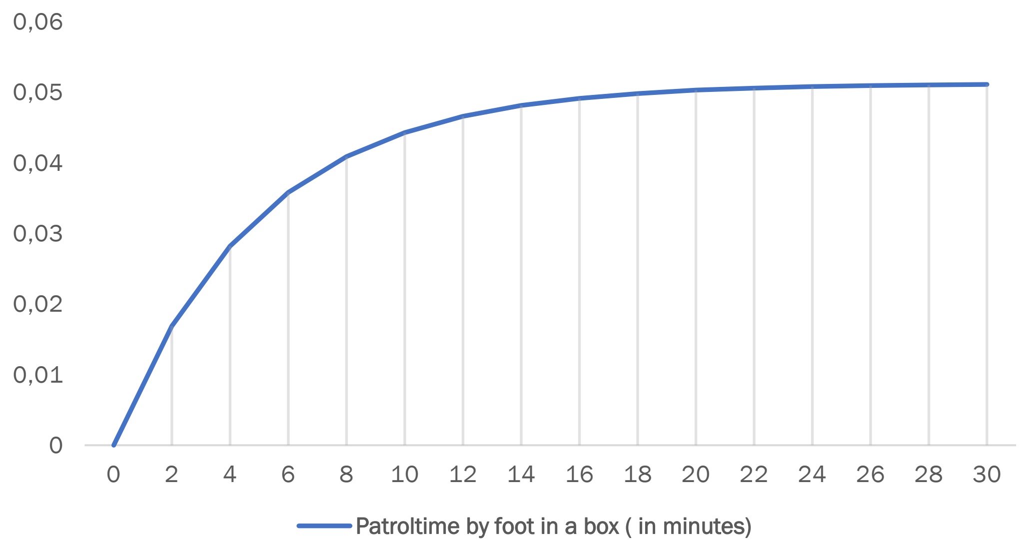

Figure 2: Patroltime Effect

We can optimise the patrol time by foot in a geographical box by multiplying the effectiveness with the (predicted) potential crime. For example, 10 minutes of patrol by foot might prevent 0.044 x 2.0 crimes (the potential) from happening. When the potential crimes for all boxes are determined, optimising available resources and expected prevented crimes is straightforward.

The formulas

The (mathematical) estimation of potentials and effects requires some assumptions. For simplicity, we can e.g. assume a weighted linear relation between the next potential ($P(t, l)$) and historical potentials in time and space:

$$P(t, l) = \sum_{u \in U} \alpha_u g_u\big(P(\tilde{t}, l)\big), \text{with } \tilde{t} < t,$$

where g is the set of terms (as a function of $P(\tilde{t}, l)$), and α represents the weights of these terms.

For the committed crime equation f_2 we have chosen a function that mimics the graph in Figure 2:

$$C(I, t, l) = P(t, l) e^{-\sum_{i \in I} E(i)},$$

where E is the effect of an intervention i. Reasonably, this effect would not be a constant scalar, but will be influenced by environmental conditions, for example: the number of houses. Hence, we model the effect of an intervention as:

$$E(i) = \sum_{v \in V} \gamma_v h_v(l_i),$$

In which h_v is the characteristic *v* of the location l_i where the intervention *i* was committed.

For the currently used application areas, such as burglary and robbery, crime reports are filed by the victims, not the police, and in most cases it is relatively safe to assume that most crimes are reported. This means that f_3 can be replaced with the very simple function: $R = C$.

Although $f_1$and f_2 should already capture most of the complexity, the model can be extended to incorporate more complex effects. For example, we can include crime displacement by adding additional terms to C such that the number of committed crimes is increased when neighbouring boxes receive interventions.

Another improvement would be to add the repeat-crime effects as seen in most criminal activities. Successful criminals have an increased confidence, which simply leads to an increased potential. To add this we introduce $\overline{C(t,l)}$ which is the realisation of committed crimes in the previous day/week/month and add this term (and a weight factor) to the $P(t,l)$.

Finding the weight factors

Finding the weight factors is a technical process requiring much computing power, but is not necessarily difficult. For readability we have chosen to highlight only the essential steps in this process, but all details and an actual implementation can be found at http://www.policing.ai.

In order to find the weight factors in the proposed model, two things need to be done. First, an estimation of $P(t,l)$ for *t=0* has to be made, and second, the factors need to be fitted to the available data. When $P(0,l)$ is estimated and the factors in equation are fitted, the potential crime for all future times can be calculated.

It is expected that the accuracy of this potential crime prediction decreases when predicting further into the future. This is expected and simply means that the fitting should be done regularly when new data is available.

To estimate $P(0,l)$ we can invert the previously introduced formulas f_1, f_2 and f_3 and replace the predicted number of reports $R(t,l)$ with the actual number of crimes reported ($\overline{R}$) and insert the action realisation of interventions ($\overline{I}$):

$$\tilde{P}\left(\overline{I},t,l\right) = \overline{R(t,l)}e^{\sum_{i\in\overline{I}} E(i)},$$

Given an estimated potential at time 0, we can obtain the weight factors in f_1 and f_2 by solving the optimisation problem:

$$\min_{\alpha,\gamma} \sum_{t,l} | R\,(t,l) - \overline{R(t,l)}|^2,$$

something that can easily be done with stochastic optimisation techniques.

6.1. Simulation results

To show the impact of this new approach the new and traditional approaches to predictive policing were evaluated in a simulation model. Simulation has the benefit over real-world testing of total controllability and allows for the isolation of the methods itself. The simulation model is available for download at http://www.policing.ai and allows for custom predictive and prescriptive policing algorithms.

For these simulations we have assumed that a police department is able to invest 31 minutes a day of officer patrol time, and this will be assigned to the single box with the highest predicted risk of crime. Moher (2015) indicated that this intervention scheme would lead to a 7.4% decrease of crime in the box receiving the intervention. Furthermore, actual crime reportings are used to establish the variation of crime across boxes in a Dutch medium-sized city:

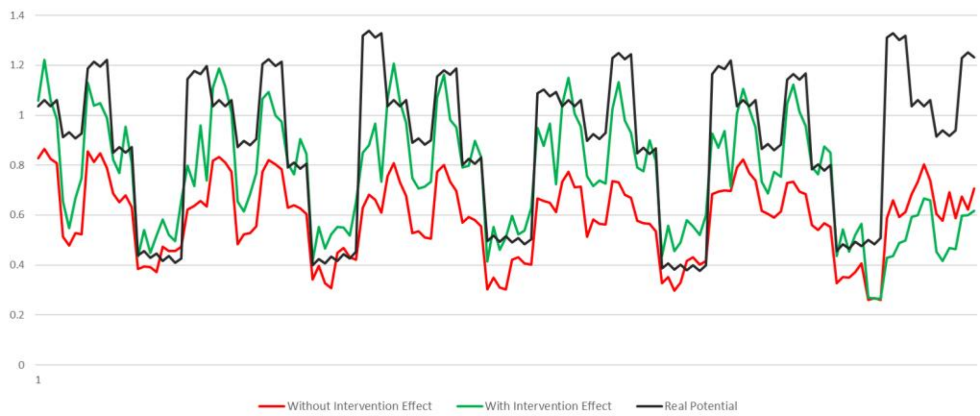

Figure 3: The real potential crime, the estimated potential crime in a traditional approach that does not include the possible effects of interventions (red), and the estimated potential crime using the presented approach (green)

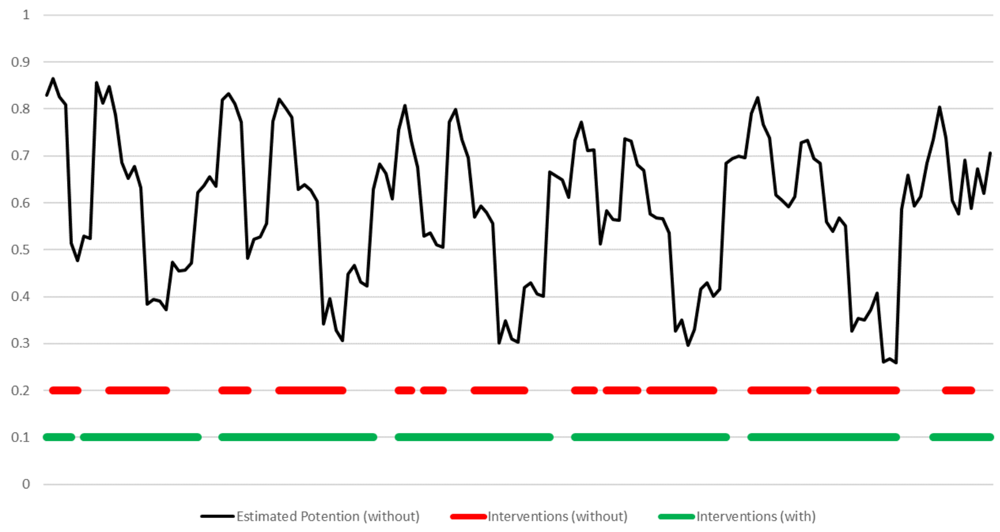

Figure 4: The estimated potential crime using a traditional appraoch, and the actual moments of intervention based on the estimation without (red) and with (green) intervention effects and a simpel prescriptive appraoch.

Figure 3 shows the predicted crime in the box with the highest potential for both an algorithm that does, and one that does not compensate for police activity (back-to-the-future paradox). As expected, the predictions that ignore police activity are much more volatile, and as Figure 4 shows, this is caused by a constantly changing 'box with the highest prediction'. The algorithm keeps assigning officers to a certain box, thus reducing crime and lowering the predicted crime for the next shift, and therefore directs officers to a different area.

This is not only very inefficient, it also affects police effectiveness. The method that does not take police activity into account has a total prevented crime of 47, while the new approach that includes this has a total prevented crime of 50. While a 7% difference in prevented crime may appear small, according to McCollister, French and Fang (2010), this means a 10M$ yearly reduction of costs on burglary alone in the Netherlands.

7. THE ROAD AHEAD

In this paper we proposed a new approach for the predictive policing feedback loop to counter the fundamental flaw in current approaches and algorithms, caused by reporting bias and the back-to-the-future paradox. The new model incorporates the effects of interventions and thereby is able to estimate a ground-truth of crime potential that is not biased by the back-to-the-future paradox. To show the impact of this new approach, both the new and the traditional approaches to predictive policing have been evaluated in a simulation model.

Improving the effect of the use of algorithmic prediction as a forecasting tool requires changes. An important first step towards breaking the feedback-loop causing instability in predictive policing is to include the actual actions taken, and to compensate for their effects.

To implement the new approach one must 1) adapt the current predictive policing approaches, and 2) measure the interventions applied. Therefore, one should invest in means and methods to record police interventions and to put them in context, while notifying the agents being tracked. It is not to put them under control, but to value the actions they take, especially those that otherwise might be invisible.

An additional benefit of this new approach is that it also enables the step to prescriptive policing: to predict the effectiveness of a particular deployment of police resources, given a particular situation, based on the knowledge of the effects of specific interventions. If predictive policing can predict where and when to be, then prescriptive policing can predict what the best likely course of action is for that geographical box and time period. If we can determine that something works in places with specific characteristics, but not in others, we can extrapolate that knowledge to other areas. Without having to determine the effectiveness of each intervention in that particular box we can already make an assessment of what is or is not likely to work (Smit, 2016).

Finally, even if this is done correctly, and predictive and prescriptive algorithms become more advanced and aware, the true value is in its use. Predictions have no value on their own, it is the actions taken by officers that can make them valuable. Hence it is important to research and invest in means of management, communication, explanation and transparency, in such a way that the results will not be overtrusted, or undertrusted

8. REFERENCES

Bennett Moses, L., & Chan, J. (2016). Algorithmic prediction in policing: assumptions, evaluation, and accountability. Policing and Society, 1-17.

Braga, A., Papachristos, A., & Hureau, D. (2012). Hot spots policing effects on crime. Campbell Systematic Reviews, 8(8), 1-96.

Greengard, Samuel (2012). Policing the future. Communications of the ACM. Volume 55, No.3; 2012.

Kathryn E McCollister, Michael T French, and Hai Fang. The cost of crime to society: New crime-specific estimates for policy and program evaluation. Drug and alcohol dependence, 108(1):98-109, 2010.

Mohler, George O., Martin B. Short, Sean Malinowski, Mark Johnson, George E. Tita, Andrea L. Bertozzi, P. Jeffrey Brantingham (2015). Randomized controlled field trials of predictive policing. Journal of the American Statistical Association. Volume 110, Issue 512; 2015.

Natarajan. M. (ed.) (2011). Crime Opportunity Theories. Routine Activity, Rational Choice and their Variants. Routledge

Perry, W.L., Brian McInnis, Carter C. Price, Susan C. Smith, John S. Hollywood (2013). Predictive policing: the role of crime forecasting in law enforcement operations. Santa Monica, CA: RAND.

Robson, C., & McCartan, K. (2016). Real world research. John Wiley & Sons.

Rienks, Rutger (2015). Predictive Policing: Taking a Chance for a Safer Future, Korpsmedia, PDC

Saunders, J., Hunt, P., & Hollywood, J. S. (2016). Predictions put into practice: a quasi-experimental evaluation of Chicago's predictive policing pilot. Journal of Experimental Criminology, 12(3), 347-371.

Smit, Selmar, Arnout de Vries, Rick van der Kleij ,Hans van Vliet (2016); From Predictive Policing to Prescriptive Policing - Beyond boxes. TNO.

Zuckerman, B. & Jefferson, D. (1996). Human Population and the Environmental Crisis. Jones & Bartlett Learning. p. 42. ISBN 9780867209662.