

Ethics Committee Briefing Note

Project Reference: WMP DAL Predicting Violent Crime

Source of analytical question / hypotheses to be examined:

The project was requested by the Senior Responsible Officer (SRO) for Project Guardian, Superintendent Ed Foster. The aim of Project Guardian is to reduce serious violence, in particular between young people in public spaces (see Context below).

As this project is at the proposal stage and is presented to the committee 'in principle' in order that any immediate concerns can be raised, the finer details of the methodology will not be determined until after the exploratory data analysis (EDA) phase. Once the analyses have been completed the projects will be presented to the Committee again so that findings and methodology can be examined in more detail.

Purpose of data analysis

The purpose of this project is to develop a statistical model to predict the likely number and most likely locations of violent crimes committed over a four week period. It is similar in nature to and is intended to complement the Knife Crime project submitted to the Committee in March 2020.

Therefore, the analyses focus on serious violent offences but exclude Domestic Abuse offences. In the performance measures for Project Guardian, the Home Office defined the following offences as 'serious violence': Homicide, Attempt Murder; Threats to Kill; Robbery; Assaults (causing injury); Firearm enabled Threats to Kill and Firearm enabled offences.

Context:

There has been a notable increase in violence within several of the UK's urban areas, and within the West Midlands in particular. In this region, gun crime has increased by 33%, and instances of knife crime have increased by 85% since 2012 and violent crime against the person is up 32% in the last year.¹

Figure 1 shows Violence with Injury offences (excluding Domestic Abuse) since 2008/09 for victims of all ages; and for offences where the victim was a 'young person' (under the age of 25).

¹ WMP OPCC: SPCB Violence Reduction Unit paper 19/11/2019 Item 8b <https://www.westmidlands-pcc.gov.uk/strategic-policing-crime-board/agendas-minutes-reports/>

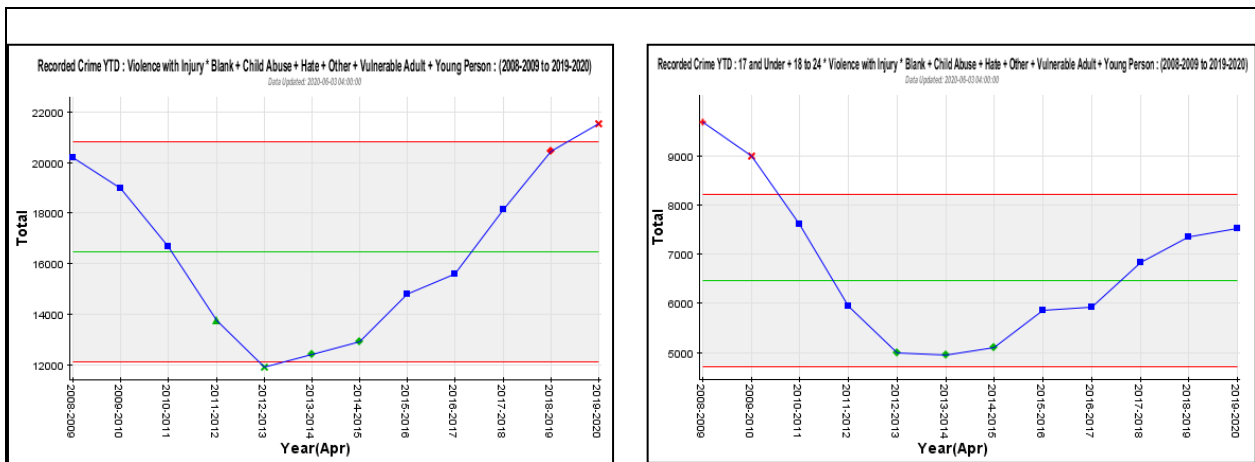


Figure 1: Violence with Injury offences since 2008/09 (excluding Domestic Abuse); victims of all ages and victims aged under 25

In addition, the rate per 1000 residents of Violence with Injury offences in the West Midlands is above average when compared to the average for England and Wales and compared to our most similar force of Merseyside.²

In April 2019, the Home Office gave West Midlands Police (WMP) £7.62 million in police surge funding with the mandate that it is to be used to reduce serious violence in public spaces, with a focus on reducing knife crimes among young people. This was followed by a further £4.9 million in 2020/21. The Force’s response was to create a two year project, known as Project Guardian. As required by the Home Office, the key focus is on prevention and enforcement activity within hotspot areas. In addition, WMP is providing an additional £1.2 million towards the activity in hot spots in 2020/21. Therefore, the purpose of this DAL proposal is to identify (in advance) the hotspots of violent crime for Project Guardian. The funding profile for this year is shown in Figure 2.



Figure 2: Project Guardian surge funding profile 2020/21

² Office for National Statistics Crime in England and Wales: Police Force Area data tables year ending Dec 2019: E&W 9.2; WMP 11.5; Merseyside 10.5; W. Yorkshire 11.8; no data for Greater Manchester.

<https://www.ons.gov.uk/peoplepopulationandcommunity/crimeandjustice/datasets/policeforceareadatatables>

Hotspot policing is a strategy that involves the targeting of resources and activities to those places where crime is most concentrated. The strategy is based on the premise that crime and disorder is not evenly spread within neighbourhoods but clustered in small locations. Focusing resources and activities in hotspots aims to prevent crime in these specific areas and potentially, reduce overall crime levels in the wider geographic area.

The College of Policing Crime Reduction Toolkit uses the EMMIE framework (Figure 3) to identify the best available evidence from systematic reviews of research on crime reduction interventions.³ The review of the evidence for hot spot policing suggests that it has reduced crime. Hotspot policing programmes that take a problem-oriented approach appear more effective than increased traditional policing, such as high visibility patrols. The evidence suggests that hotspot policing was more effective for drug offences, violent crime and disorder than it was for property crime. Hotspot policing can also lead to a diffusion of benefits to the areas immediately surrounding the hot spot.⁴ This is a tactic that has been used widely in WMP⁵ and other forces; for example most recently in London where Violence Suppression Units focus on ‘micro-beats’.⁶

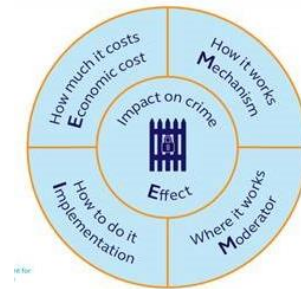


Figure 3: EMMIE framework

Intended activity resulting from the project

WMP’s 3 year strategy ‘This Work Matters’ requires us to act with precision and maximise the impact of every intervention to keep people safe, by equipping us with the information and tools we need to make the most effective decisions (Figure 4).

To support this objective, the DAL will provide forecasts over a 4 weekly period to enable the SRO for Project Guardian to plan resource allocation and activity.

This analysis would not influence the targeting of specific individuals.

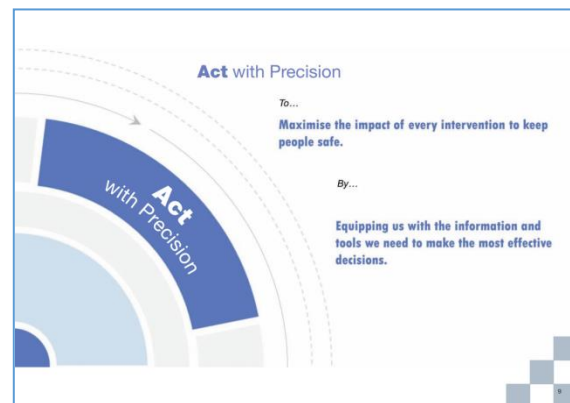


Figure 4: WMP Strategy ‘Act with Precision’

³ <https://whatworks.college.police.uk/toolkit/Pages/Welcome.aspx>

⁴ Braga, A., Papachristos, A. and Hureau, D. (2012) ‘Hot spots policing effects on crime’, Campbell Systematic Reviews, 2012:8, DOI: 10.4073/csr.2012.8. <https://whatworks.college.police.uk/toolkit/Pages/Intervention.aspx?InterventionID=46&src=%23filter%3D1>

⁵ https://whatworks.college.police.uk/About/News/Pages/Hot_spots_RCT.aspx

⁶ <http://news.met.police.uk/news/commissioner-sets-out-plans-to-continue-to-tackle-violent-crime-402325>; <https://www.theguardian.com/uk-news/2020/may/31/out-on-the-microbeats-where-police-hope-to-reclaim-londons-streets-from-gangs>

The activity that will be undertaken in the identified hotspots includes both prevention and enforcement, however, the focus is on the prevention of violent crime and reducing recidivism, rather than simply arresting offenders. As shown in *Figure 2*, almost two thirds⁷ of the Home Office surge funding allocated to the hotspot areas is being used to increase the capacity of our neighbourhood teams. These teams, who know the local area, focus on prevention activity such as engagement with schools and youth groups. The taskforce is an additional team of specially trained officers who work proactively at locations and times when the likelihood of violence occurring is greater. They use tactics such as knife arches to deter young people from attending events with the intention of causing violence.

In addition to the Home Office funding, WMP is part of the West Midlands Violence Reduction Unit (VRU) which takes a collaborative regional approach to addressing violence, vulnerability and exploitation. This multi-sector body convenes a range of activities and initiatives designed to embed a 'public health' approach to violence across the system, underpinned by the conviction that 'violence is preventable, not inevitable'. The VRU seeks to develop a system change approach, whilst commissioning some services directly to develop an evidence base to develop preventive practice.

Project Guardian is closely linked to the VRU which is the mechanism through which a partnership approach to violence is delivered.⁸ As seen in *Figure 2* WMP has allocated further funding for a team of 'violence interrupters' from partner agencies who offer diversionary tactics in youth violence hotspots. As discussed above, this partnership problem solving approach is identified in the research as the most effective use of hotspot policing.

Data to be used:

The units of analysis will be serious violent crimes as defined for Project Guardian by the Home Office (as described above) and intelligence reports which relate to violence within the WMP area.

Crimes data relating to individuals (as offenders or victims) will be the basis of this analysis. However, this will be aggregated and no individual will be identifiable from the analysis.

The geographic output would be bespoke spatial units as generated by the hotspots.

Level of analysis:

- Individual
 - Individuals aggregated?
 - Yes
 - No
- Specific Area:
 - Output Areas
 - Super Output Areas - Lower
 - Super Output Areas - Mid

⁷ £1,994,000 of £3.052,200

⁸ More information about the range of support via the VRU can be found in the paper presented to the SPCB in November 2019.

<https://www.westmidlands-pcc.gov.uk/strategic-policing-crime-board/agendas-minutes-reports/>

- Wards
- Districts
- Other
- West Midlands
- Other

Reliability of data:

An extensive exploratory data analysis (EDA) phase will be undertaken to examine the extent of any data quality issues, including processes to identify the presences of any bias, to ensure that no bias is built into the model.

Discussions with subject matter experts (SMEs) will be undertaken both to capture any extraneous requirements and to sense check the analyses.

Sample or entirety: Entirety

If sample: N/A

Method of sampling: N/A

Method of choosing sample size: N/A

Sample size: N/A

Type of analysis:

- Exploratory
- Explanatory
- Predictive
- Optimisation

Proposed methodology:

It is apparent that there are patterns to violent crime both over time and over space. Therefore it is envisaged that a spatio-temporal model would likely provide the best means of prediction, potentially also with a separate univariate time series model. For the purposes of spatial prediction, it is also envisaged that the WMP area would be broken into small spatial units in order to cater for analyses and predictions over areas that are distinct and small enough to allow operational activities to be planned.

This analysis is not intended to explain the reasons why violent crime occurs in certain locations and at particular times.

Will the project eventually be automated:

- Yes
- No

Means of evaluation:

The accuracy of the model would be assessed during its build phase through comparison to a test dataset. Following this beta testing of the model would be undertaken in order to assess its accuracy

over a period of time on new data.

ALGO-CARE considerations:

Advisory:

If applicable, are the outputs from the algorithm to be used in an advisory capacity?

The predictions would be used to assess levels of risk and potential levels of harm and threat. This would feed into operational planning for Project Guardian.

Does a human officer retain decision-making discretion?

The predictions would be for informational purposes only to feed into decisions that would ultimately be made by officers following operational protocols.

Lawful:

What is the policing purpose justifying the use of the algorithm (means and ends)?

Being able to predict where and when violent crime is most likely to occur, so as to plan prevention activity, supports the delivery of the government's Serious Violence Strategy⁹ which puts greater focus on steering young people away from a life of crime, while continuing to promote a strong law enforcement response. Similarly, WMP's Improvement Plan for 2020/21 includes the objective to '*act with precision in preventing serious violence affecting the under 25s*'. In addition, the Home Office requirement is that surge funding is used to target hotspots.

Is the potential interference with the privacy of individuals necessary and proportionate for legitimate policing purposes?

Whilst data regarding individuals will be processed, this would be to produce aggregated data (counts of crimes over certain time periods and in various locations) as the basic unit of analysis would be crimes and their location. Similarly, the use of Intelligence Reports would be to understand the volume of intelligence relating to geographical areas, rather than to focus on any individuals named. Therefore there would essentially be no interference with the privacy of individuals.

In what way will the tool improve the current system and is this demonstrable?

The temporal and geographical analysis of crime that has occurred in the past has traditionally been produced by Intelligence Analysts using mapping software (MapInfo) to produce hot spot maps and excel. The computing power and expertise of the DAL will enable a more sophisticated analysis of WMP crime data in order to make predictions about where and when violent offences are more likely to occur in the future, which is beyond the capability of the tools available to analysts.

Are the data processed by the algorithm lawfully obtained, processed and retained, according to a genuine necessity with a rational connection to a policing aim?

⁹ <https://www.gov.uk/government/publications/serious-violence-strategy>

The data are from WMP systems and collected to enable normal day-to-day operations.

Is the operation of the tool compliant with national guidance?

The analyses proposed would accord with the DCMS Data Ethics Framework 2018.¹⁰

Granularity:

Does the algorithm make suggestions at a sufficient level of detail given its purpose and the nature of the data processed?

Given the overall aim of predicting the number and locations of violent crime, using crime data aggregated over time and to bespoke spatial units would enable better decision making within WMP.

Are data categorised to avoid broad-brush grouping and results and therefore issues of potential bias?

Using aggregated counts would best suit the aims of the project, but would not involve any other form of categorisation (given that crimes would be the unit of analysis). For example, there would be no categorisation of individuals' protected characteristics.

Do the potential benefits outweigh any data quality uncertainties or gaps?

The project will include an extensive EDA element and this should highlight areas of heightened uncertainty in the data or where particular gaps exist. Should any such issues be identified, these would be addressed as a part of the project. Given the benefits of reducing violent crime it is not expected that any data quality issues would be of such a magnitude as to warrant not undertaking the project.

Is the provenance and quality of the data sufficiently sound?

The data have been gathered during the day-to-day investigative work of WMP and do enable analyses of the type envisioned for this project.

If applicable, how often are the data to be refreshed?

Currently envisaged as circa every four weeks.

If the tool takes a precautionary approach in setting trade-offs, what are the justifications for the approach taken?

If a predictive model can be built, its accuracy would be tested on separate data which had not been used in building and training the predictions. This can be done using historic data – where we know what happened and would tell us if the model would have predicted events or not. If the model is not deemed to be accurate enough, then the project would not be pursued. Ultimately any model developed would aim to maximise specificity whilst trying to gain as high a sensitivity as possible. This approach would mean that we could best allocate WMP resources whilst ensuring a minimisation of false positives.

¹⁰ <https://www.gov.uk/government/publications/data-ethics-framework/data-ethics-framework>

Ownership:

Who owns the algorithm and the data analysed?

WMP own the algorithm and data.

Does WMP need rights to access, use and amend the source code and data?

No

Are there any contractual or other restrictions which might limit accountability or evaluation?

No

How is the operation of the algorithm kept secure?

The data and the analyses are contained wholly within the WMP Hadoop system and the security measures employed therein.

Challenge:

What are the post-implementation oversight and audit mechanisms, e.g. to identify any bias?

Checks will be made as to the accuracy of the model on an on-going basis (overall accuracy, sensitivity, specificity, AUC, etc.) as well as any consistent patterns that may represent biases.

If the algorithm is to inform criminal justice disposals, how are individuals notified of its use?

Not applicable.

Accuracy:

Does the specification of the algorithm match the policing aim and decision policy?

The model would aim to produce information to aid decision making within WMP to support the Force's strategic aim of reducing violent crime.

Can the accuracy of the algorithm be validated periodically?

The productionisation of any model resulting from the project would include checking its accuracy on an on-going basis.

Can the percentage of false positives / negatives be justified?

Not yet known, however any model developed would aim to maximise specificity whilst trying to gain as high a sensitivity as possible.

How was the method chosen as opposed to other available methods?

Currently the broad approach has been identified due to the nature of the business question and the

data available.

What are the (potential) consequences of inaccurate forecasts?

The main issues arising from inaccurate forecasts would be (a) potential for actions for WMP that may not be necessary (including in particular locations) and (b) WMP resources being allocated ineffectively.

Does this represent an acceptable risk?

Any model arising from this project would seek to balance the advantages against the risks arising from inaccurate predictions partly via balancing the model's sensitivity and specificity and partly through assessing the types of decisions for which any such model would be effective for and any actions that would arise from these decisions. This would be subject to periodic review.

How are the results checked for accuracy and how is historic accuracy fed back into the algorithm for the future?

For any model that was productionised, it's accuracy would be assessed on an on-going basis via measuring its accuracy (sensitivity, specificity and AUC) as well as producing histograms of counts / estimated probabilities so that any degradation of the model could be tracked and the model rebuilt if necessary.

How would inaccurate or out-of-date data affect the result?

This is partly dependent on the nature of any model should one be capable of being built. Generally inaccurate or out-of-date data could detrimentally impact on the model's performance (in terms of accuracy) and lead to inefficient decision making and resource deployment.

Responsible:

Would the operation of the algorithm be considered fair?

During the development of any model, the presence of any biases in the underlying data or for predictions to produce biases would be fully examined and mitigated if the potential was present.

Is the use of the algorithm transparent (taking account of the context of its use), accountable and placed under review?

The details of any model arising from this project would be provided and, as mentioned above, when productionised there would be on-going checks as to model performance.

Would it be considered to be used in the public interest and to be ethical?

Enabling the reduction of violent crime would be the aim of the project and given the cost to society of violence, enabling more effective means of its reduction would likely be seen as positive.

There may be potential for this project to be seen as another instance of hotspot prediction that will be disadvantageous to various communities. This project would aim to use robust analyses and methods such that any such concerns could be mitigated. Once productionised, the accuracy of the models would

be kept track of in order to ensure that the model(s) do not degrade.

There are a number of existing mechanisms for reporting back to our communities in order to enable them to scrutinise our decision making. These include the OPCC Strategic Policing and Crime Board (SPCB); Independent Advisory Groups (IAGs) on each geographical area and scrutiny panels for Stop and Search and Use of Force. These existing arrangements could be considered as the channel for communicating the output of this project.

Explainable:

Is information available about the algorithm / decision-making rules and the impact of each feature?

Such information would be produced for any model arising from this project.

Glossary

WMP / Law Enforcement Terminology	
COP	College of Policing
DA	Domestic Abuse
DAL	Data Analytics Lab
EMMIE	The EMMIE framework developed by University College London used to rate the best available evidence on crime reduction interventions by evaluating their Effect; Mechanism; Moderators; Implementation and Economic cost. Used by College of Policing in their systematic reviews of research papers.
IAG	Independent Advisory Group
NPU	Neighbourhood Policing Unit
ONS	Office for National Statistics
OPCC	Office of the Police and Crime Commissioner
PCC	Police and Crime Commissioner
SME	Subject Matter Expert
SPCB	Strategic Policing and Crime Board
SRO	Senior Responsible Officer
VRU	Violence Reduction Unit
WMP	West Midlands Police
Data Science Terminology	
ALGO-CARE	All projects have used the ALGO-CARE to consider ethical implications: Advisory, Lawful, Granularity, Ownership, Challenge, Accuracy, Responsible, Explainable
AUC – ROC Curve	AUC stands for ‘area under the curve’ of a ROC (Receiver Operating Characteristics) curve. Essentially, ROC is a probability curve and the AUC tells us how good the model is at distinguishing between different groups within the data. This is a statistical test for the accuracy of the model that has been built.
DCMS	Department for Digital, Culture, Media and Sport – developed the Data Science Ethical Framework.
EDA	Exploratory Data Analysis

Productionise	To 'productionise' means that once we are satisfied that the model works well, we would automate the process of providing predictions on a regular basis.
Spatio-temporal model	Spatio-temporal models use patterns evident over space as well as patterns evident through time – “everything is related to everything else, but near things are more related than distant things” (Tobler’s first law of geography). This is often true in time as well as over space.
Sensitivity	Refers to the ability of the model to identify the ‘true positives’ as a rate. It measures the proportion of actual positives that are correctly identified as such. The greater the sensitivity of a model, the less the specificity will be.
Specificity	Refers to the ability of the model to identify the ‘true negatives’ as a rate. It measures the proportion of actual negatives that are correctly identified as such. The greater the specificity of a model, the less the sensitivity will be.
Univariate time series model	A time series that consists of single data points recorded sequentially over equal time increments.