

# Improving Policing with Natural Language Processing

**Anthony Dixon**

School of law  
University of Leeds  
mm18acd@leeds.ac.uk

**Daniel Birks**

School of Law  
University of Leeds  
D.Birks@leeds.ac.uk

## Abstract

This article explores the potential for Natural Language Processing (NLP) to enable a more effective, prevention focused and less confrontational policing model that has hitherto been too resource consuming to implement at scale. Problem-Oriented Policing (POP) is a potential replacement, at least in part, for traditional policing which adopts a reactive approach, relying heavily on the criminal justice system. By contrast, POP seeks to prevent crime by manipulating the underlying conditions that allow crimes to be committed. Identifying these underlying conditions requires a detailed understanding of crime events - tacit knowledge that is often held by police officers but which can be challenging to derive from structured police data. One potential source of insight exists in unstructured free text data commonly collected by police for the purposes of investigation or administration. Yet police agencies do not typically have the skills or resources to analyse these data at scale. In this article we argue that NLP offers the potential to unlock these unstructured data and by doing so allow police to implement more POP initiatives. However we caution that using NLP models without adequate knowledge runs the risk of perpetuating existing, or introducing new, biases that have the potential to produce unfavourable outcomes.

## 1 Introduction

This article will first provide a brief overview of Problem-oriented Policing (POP) and demonstrate that it is an efficient crime prevention strategy. It will show that by implementing POP processes and reducing criminal opportunities less people are likely to commit crime and end up within the criminal justice system. It will then demonstrate that while POP has previously been successful the analytical burden it places on crime analysts is substantial and is an impediment for wider adoption.

Subsequently, we will argue that NLP methods have the potential to support efforts to overcome these challenges - enabling at-scale systematic extraction of insights from police free text data sets that can support the POP process. We will conclude by discussing several ethical challenges that must be overcome if NLP is to help deliver positive societal outcomes by supporting those who seek to reduce crime.

## 2 Problem-Oriented Policing

POP is a model of policing proposed in 1979 by Herman Goldstein ([Goldstein, 1979](#)) as an accompaniment to the traditional policing model. Traditional policing focuses resources on reactive response, investigations and arrests. Arrests lead to prosecution, court, prison and probation costs and the criminalisation of, mostly, young males. By contrast, POP seeks to re-balance this traditional reactive approach ([Goldstein, 1990](#)) to include preventative efforts which act before the crime or problem arises ([Tilley, 2008](#)).

To this end, POP seeks to prevent problems from reoccurring by analysing how previous similar events occurred then intervening in that generation process to prevent recurrence - see Fig 1 for a pictorial representation. In this regard, an essential element for conducting POP is understanding the conditions that allowed crime to occur in the first instance. POP is based upon understanding crime as a socio-physical process that occurs when three separate elements coincide. Much like a fire relies on a fuel, a spark and oxygen to occur, crime relies on the convergence of a motivated offender, a suitable target in a setting without a capable guardian ([Cohen and Felson, 1979](#)). POP seeks to understand how these elements, known as the crime triangle, coalesce and therefore how the triangle can be disrupted to prevent crime opportunities.




Figure 1: A schematic for POP. Reproduced from (Eck and Spelman, 1987)

POP is generally regarded as a successful police model when implemented correctly. There are systematic reviews that provide evidence for POP's increased effectiveness in crime prevention over the traditional model. A recent systematic review (Hinkle et al., 2020) found that POP was much more effective at preventing crime than traditional policing. A second review (Braga et al., 2014), also found that when targeted alongside another police tactic, hot-spot policing, POP was also more successful than traditional policing. Moreover, a number of randomized controlled trials have shown that POP is more effective at preventing crime than traditional policing approaches (Taylor et al., 2011; Braga et al., 1999).

From a social justice perspective POP has the effect of reducing opportunities for crime across communities, and thereby reducing the attractiveness of crime in areas where it is traditionally higher. With reference to the crime triangle, high crime areas may contain similar quantities of potential offenders to low crime areas, but lack capable guardians or security measures, thus creating more viable opportunities for crime. A decreased reliance on the criminal justice system also means less people are criminalised. In what follows we outline the POP process, provide some illustrative examples, identify some key criticisms and challenges associated with its application, and describe how NLP might

be used to overcome these and facilitate positive impact.

## 2.1 POP Framework - SARA

The POP analytical framework is typically based upon a four stage process - Scanning, Analysis, Response and Assessment (SARA):

1. **Scan.** Firstly the problem space is scanned for collections of incidents that represent a potential problem to be addressed. Typically this scanning is completed by the police in conjunction with the community, either directly or indirectly through received complaints. The *scan* is wide but analytically shallow. The output is a reduced collection of incidents that share the same characteristics indicating common underlying causes.
2. **Analysis.** After the problem space is defined, it is then analysed with the aim of identifying underlying conditions that might be manipulated to prevent the crime - these are often known as pinch points. This stage is typically the most arduous from an analytical perspective, as the details of each crime need to be thoroughly understood to allow common pinch points to be identified and understood. In comparison to the *scan* stage, *analysis* is much more focused delving deeper into the crimes selected.

3. **Response.** The third stage - *response* - the is aimed at the pinch points identified in the previous stage. By manipulating these pinch points the conditions for crime are altered, with the aim of making criminal opportunities less attractive, more risky, more difficult or removing them altogether.
4. **Assessment.** The final *assessment* stage seeks to assess the effectiveness of the intervention, capturing information that can enhance the response and inform future POP users.

We now illustrate this this framework by means of an example from an op rational policing environment in the UK.

## 2.2 POP Example

An example of POP implementation is included to demonstrate how the process operates and how success is achieved. The example is centred on residential burglary reduction in Durham UK.<sup>1</sup> Durham Constabulary, situated in Northern England, had experienced consistently high rates of residential burglary. Reliance on traditional policing methods had not addressed the problem with burglary rates remaining high even after offenders had been caught and convicted. In response a different approach was sought through POP.

1. **Scan.** Durham's burglary data from a number of years was analysed to identify the type of dwelling, items stolen and modus operandi (how burglary was committed) associated with residential burglaries. These factors were used to highlight areas where the same types of burglary occurred - that is the *scan* of the whole force area identified smaller areas where the same types of crime were being committed, thus allowing an investigation into the underlying causes. At this stage large volumes of crimes are analysed (typically there are around 4000 burglaries in Durham a year) in order to select a coherent manageable group of crimes for further analysis in the following stage.
2. **Analysis.** Once the areas for enhanced analysis had been determined, crimes were further explored to understand how and (where possible) by whom they had been committed. Combined analyses of police records and

intelligence data led to the identification of opportunistic as well as organised gang burglaries, and identified poor residential security as an underlying issue along with insufficient informal guardianship in selected areas.

3. **Response.** After *analysis* of the problems and a shift away from relying on the criminal justice system, the police garnered public support to change community behaviours. This made the areas less attractive to burglars by enhancing informal guardianship. In addition, the police provided home security packs to the most vulnerable residents. The result was a reduction in burglaries in the majority of the POP response areas, against a backdrop of rising burglaries across the region. Not only was this intervention cost effective relative to a traditionally criminal justice response, it also, more importantly, meant that significantly fewer residents had their homes violated.
4. **Assessment.** The *assessment* phase was conducted by comparing levels of crime in the intervention and control areas pre- and post-response. This was carried out using simple count data and tracked whether the POP initiatives had reduced crime in the target areas relative to control areas. While this approach was able to estimate the impact of the *response* in the target area, it still exhibited a key weakness, in that without further detailed analyses it could not provide insights into how offences had been prevented or how their nature may have changed as a result of the response. Consequently this assessment was of limited value for considering how such tactics might be improved or adapted for use in other areas.

## 2.3 Impediments to POP

Significant information that is required for POP is contained in textual data. Some of this will be in police generated crime notes - such as the modus operandi described above, witness statements, forensic reports or other sources such as complaints from the community. Analyses of these data is largely completed manually (Goldstein, 1990), and as such it is often a long and laborious task, and given resource pressures, the work often has to be completed selectively. Unlocking access to this information would enable analysts

<sup>1</sup><https://popcenter.asu.edu/sites/default/files/17-04.pdf>

and officers access to a much wider source of information with which to implement POP responses. In a guide to POP, [Scott and Kirby \(2012\)](#) cite the need to both get and train the right staff (Chapter 9) and the requirement for enhanced analytical support is highlighted at great length (Chapter 17). POP requires appropriate knowledge, skills and experience to be delivered effectively, but because these skills are not required for the traditional response policing model, they are often lacking in within police agencies.

To chronologically bookend this point, a lack of analytical skills was identified by Goldstein in 1990, ([Goldstein, 1990](#)), and was still seen as an issue in 2016 ([Scott et al., 2016](#)). A recent review of POP in England and Wales ([Sidebottom et al., 2020](#)) concluded that “recurrent weaknesses in the application of SARA...concerned the depth and quality of problem analysis.”, additionally they also found that “43% of survey respondents said they did not have access to information necessary to perform effective problem-solving”. Given that the crux of POP lies in the understanding of the problem at hand, yet the police agencies that want to implement POP do not have the necessary skills available in sufficient quantities, it is hardly surprising that POP usage is not widespread. However, it is encouraging to note that it would largely appear to be a resourcing issue, rather than a systemic POP problem as where analytical resourcing have been sufficient, often as a result of collaborations with academia, POP implementations have been more successful.

With these constraints in mind, it seems clear that if some components of the POP process could be supported through automation, then at least one obstacle to expanding POP implementations would be overcome. It is here that we believe modern NLP techniques have the potential to facilitate rapid exploitation of police free text information, in turn contributing to a significant lowering of the analytical burdens associated with successful POP implementation. Yet, to simply burden police analytical staff with yet another complex tool will likely not produce a desirable outcome. Instead, tools need to be simplified and packaged so that time-poor analysts without extensive training can leverage the technology even if that means not harnessing the full potential of NLP technologies.

### 3 Police Free-Text

In many countries, including the United Kingdom, the police have a legal requirement to record and document crimes. This documentation can vary depending on the severity of the crime and procedures within individual agencies. As can be seen from example texts in ([Birks et al., 2020](#)) and ([Kuang et al., 2017](#)) police free text includes misspellings and specialised vocabulary like acronyms and contractions. Police free text is also generally unedited, capital case rules are liberally applied and often there is little formal grammar. All this sets police free-text apart from the data sets that are generally used to train existing NLP models, suggesting that the nature of the text will require model adaptations to reach similar results to those achieved using the types of data sets existing models are trained on. Despite these differences, some preliminary experimental work carried out by the author has shown that existing models give sufficient coverage to the language without adaption. Work to understand the utility of part-of-speech taggers showed that using a universal dependency parser based on the English Web Tree Bank<sup>2</sup> ([Silveira et al., 2014](#)) an overall token accuracy of 90% was achieved when tested on Burglary Modus Operandi text, although that did mean that around 67% of sentences contained at least one error.

A further challenge is the sensitivity of police data. Police free-text data can contain personal information and so are often subject to local laws and regulatory frameworks (such as GDPR in the UK and EU). These protections present challenges. Police agencies, as we have previously discussed, typically do not have the expertise to conduct the detailed analyses in house and almost certainly do not have access to GPUs or other accelerators to build some of the more powerful models from scratch. At the same time, timely sharing of sensitive data in ways that facilitate academic research can present significant logistical challenges. This means that the NLP analytical engines will most likely have to travel to the data located in the police IT systems, unless systems can be developed to securely move and store the data. Any NLP implementation would, ideally, have low hardware requirements and be packaged so that it can be used by practitioners who may be quantitatively competent, but not be experts in NLP or machine learning tech-

---

<sup>2</sup>Work was completed in R using `udpipe` package with model `english-ewt-ud-2.4-190531`



niques. In order to overcome these data sharing obstacles we have initially adopted a very low risk approach with a partner agencies to release data for experimental research. This approach is characterised by the following methods:

1. Low risk data. Requests for data are designed from the outset to be low risk, we request modus operandi data which is designed to be shared with other parts of the criminal justice system and as such is not supposed to contain personal data.
2. In-house pre-processing. To add an additional level of security we have also developed a simple approach to further pre-process data in police systems prior to sharing. Our white-listing approach simply redacts all tokens that are not found within a list of commonly used words (circa 10,000). Crucially this list does not contain common names, again minimising the risk of disclosure of personal data. While this approach may be sub optimal relative to other methods it is deterministic and easily explainable.
3. Safe place. All data are held in modern secure environments. We have utilised a secure area (ISO27001 compliant) that can only be accessed by members of the project team.
4. Safe people. Members of the research team are vetted by the police force in question to ensure they meet necessary standards for data handling.
5. Shared insights. We agree to share all insights with our police partners. All publications detailing research are vetted by multiple parties from both police and academia prior to submission.

Clearly these approaches will have an impact on the data received and therefore the generalisation of NLP applicability to different types of data (e.g. witness statements). However this approach does offer a promising beginning to understand how and if NLP can be useful for POP processes.

#### 4 Related Work

Machine learning, text mining and data science have, unsurprisingly already been seen as useful tools by crime scientists (Marshall and Townsley,

2006). However, as a recent review into the intersection of crime and AI has shown (Campedelli, 2020), although some methods of AI and machine learning exist in the criminological literature, there is a general paucity of NLP related research compared to other areas. In this section we concentrate on analyses of free-text police data only.

Much of the existing crime free-text analysis is dominated either by unsupervised learning and revolves around the problem of crime linkage rather than crime reduction (Hassani et al., 2016). Crime-linkage seeks to identify crimes that are committed by the same individual(s), whereas POP typically requires crimes grouped according to enabling characteristics. Notable examples of unsupervised learning with Police Free-text data are Birks et al. (2020) and Kuang et al. (2017) who use unsupervised natural language processing to understand how crimes may be grouped relative to how they were committed rather than traditional crime classifications. Birks et al. (2020) completes this within a single crime classification and Kuang et al. (2017) conducted this across multiple crime classifications. This is referred to as crime topic modelling and seeks to understand crime from an ecological perspective.

In addition to the previous studies a pair of recent studies conducted with police data from Brazil, (Basilio et al., 2020, 2019) utilise unsupervised NLP techniques to cluster crimes with the hope of understanding what policing strategies will be suited to different areas of the city. The authors cluster crimes, then show police officers a representative sample of the clusters and ask them to nominate a suitable policing style (traditional, POP or hot-spot). They do not report if the styles were subsequently adopted or if they were successful.

Recently the complexities of models used with crime data has increased and there has been work to extract specific information directly from police free text data, see for example the work by Karystianis et al. (2018, 2019) who seek to explore relationships between mental health and types of domestic violence through rule-based information extraction. However, information extraction requires significant efforts to build rules and dictionaries, and whilst this approach is undoubtedly more effective than manually trawling through thousands of records it still likely represents an implementation hurdle that is too great for routine adoption.

For NLP to aid POP, algorithms need to be de-

veloped that can assist with the characterisation of crime events. Whether this is with known characteristics, such as presence of alcohol or type of victim-offender relationship, or perhaps unknown characteristics that are discovered through unsupervised learning. The extant research discussed above provides a foundation for further explorations into the utility of NLP, but to the authors' knowledge no current research focuses on characterising crime events for the purposes of aiding crime prevention, more so if one also considers the desire for such solutions to operate without the need of high performance computing. Thus, the focus of future research to enable POP should be on examining how existing NLP models can be utilised against police generated free text data, in a low resource environment, with the aim of enhancing the characterisation of crime events.

## 5 NLP Applications

Policing encompasses a diverse set of tasks and responsibilities. It is conceivable that NLP methods could be used to support a broad array of processes associated with these functions. This section will focus on those NLP applications that we believe may offer direct benefit to POP processes, in turn reducing the aforementioned analytical burden associated with their application in real world police settings.

### 5.1 Classification

Police agencies often flag crimes with keywords to help understand contextual factors associated with a particular offence. For instance, a common flag is to record if an offender is under the influence of alcohol or illegal drugs. Often these flags are not completed thoroughly (there may be hundreds of flags to select from) because police officers are under time-pressure to deal with the situation at hand. Classification algorithms can be used to check these flags and broaden the coverage where officers have described the presence of a flag but not separately recorded it, thereby giving police analysts a more complete picture of known factors. In reference to the Durham example highlighted above classification may have been used to understand if force had been used to enter a given residence or if the residence had been deliberately targeted for example to steal a high performance motor vehicle. This kind of classification can be very useful for the *scan* stage of POP, as enhancing the structured data with

additional and more complete crime characteristics from text data can assist in grouping crimes with a similar context or process to form the nucleus of a POP intervention.

### 5.2 Named Entity Recognition

Named Entity Recognition (NER) may be used by POP analysts to extract specific elements of a crime from crime reports, modus operandi or related intelligence data. For instance, it may be used in assault cases to extract a weapon type, or in domestic abuse cases to understand the relationship between the victim and the offender. Matching on key characteristics like this will facilitate better problem grouping, and will be an improvement on current information availability as quite often this level of detail is not included in a structured manner. In the case of the Durham example NER might have been used to further understand the method of entry - for instance distinguishing between entry methods such as smashing a window or the breaking of a particular type of lock. Crime prevention strategies work best when they are specific, for examples denying entry through snapping patio door locks requires a different strategy to that of combating burglars who exploit insecure properties. NER has the ability to extract this level of detail from crime reports and thereby vastly reducing the time spent in the *analysis* phase of the POP cycle where currently police analysts have to trawl manually through the detail to retrieve this information in order to form an appropriate POP response.

### 5.3 Clustering

The two previous techniques rely on searching for known characteristics. Unsupervised clustering may improve on this by allowing similar crimes to be grouped so that POP Responses (the R in SARA) can be targeted more efficiently. This would build on the work mentioned above (Kuang et al., 2017; Birks et al., 2020) enabling analysts to be free from the strictures of pre-existing administrative categories and pre-conceived notions of the main causal factors. This clustering can also be extended to encompass other variables, such as time and location information, enabling a richer scan for problems than would otherwise be the case. In the example of burglary, clustering may provide insights into the emergence of new modus operandi. In the past techniques such as hooking keys through letterboxes or snapping certain door locks have emerged and have only been tackled once in widespread use.

Unsupervised techniques could also be useful in the *assessment* phase of the POP framework, as understanding how criminals are adapting to POP responses is an important part of ensuring lasting impacts from POP interventions. The emergence or shift of crime clusters after a POP evaluation can indicate that perhaps new techniques are being used in order to overcome the POP intervention.

## 6 Ethical Implications

While NLP may offer a range of opportunities to police agencies, utilisation of free-text information from police activities will be subject to similar ethical considerations and biases as other usages of NLP. However, in the case of police usage the a key consideration must be the potential societal impact of biases.

There is a real risk that improper or careless uses of NLP may introduce or perpetuate biases that serve to undermine relationships with the communities that the police are there to serve, thus adding to problems rather than solving them. For this reason it is imperative that ethical considerations, particularly around potential biases are considered before implementation and at all stages of the utilisation, by those devising analytical solutions, analysts who apply them, and those officers that formulate the POP responses. Here we envisage three main areas where use of NLP may be effected by bias. Typically these areas are likely to produce resource allocation biases (Blodgett et al., 2020).

### 6.1 Data Coverage

Police do not know about all crime, in the UK it is estimated that only around 40% of crime is reported to the police (Tarling and Morris, 2010). The single biggest factor for reporting crime is the seriousness of the offence and in other research (Baumer, 2002) the level of disadvantage in a neighbourhood has correlated with lower reporting rates. This lack of coverage could lead to biases in areas where reporting of crime to the police is lower than in other areas (similar problems already exist when analysing structured police data). That is, NLP could bias resource allocation to areas where recording is more complete and POP implementations are therefor easier to implement, thus leading to an unfair distribution of resources.

### 6.2 Data Richness

When utilising free-text information the quality of the information extracted is wholly dependent on the information recorded in the first instance. If there are systematic imbalances in the detail of recorded crime across areas, communities or particular groups then these biases will be resident within the free-text data and are likely to be replicated into the available information for POP responses. These biases will need to be guarded against, and as part of the development of NLP for POP there will need to be research into the richness and overall quality of information that is recorded across victim characteristics and crime types. Failure to guard against these biases could see an uneven application of POP activities favouring areas where the police-community information flows are more efficient.

### 6.3 Algorithmic Bias

Crime is highly concentrated both in space and in relation to particular victims (Farrell, 2015). That is, we would expect different crime types to disproportionately affect different parts of society. Similar crimes are also likely to have similar written descriptions as they describe similar processes. The danger is that if the description of certain crimes are not well understood by certain models, (e.g. certain crime descriptions might use unusual language in the context of the original training data for pre-trained language models) then this will mean poorer information retrieval for certain crimes and therefore potentially for certain victim profiles. This is an example of algorithmic bias (Hooker, 2021) where model selection can effect the distribution and quality of model outputs. Consequently, it will be important to review all models in the context of the specific crimes for which they are to be utilised. This suggests that model applicability will need to be judged at a crime-specific level. This approach should allow metrics to be reviewed for each crime type to make sure that no crimes, and, in turn, victim types are misrepresented. Relatedly, biases in errors from models, perhaps reflecting some of the existing recording practices, will also likely need to be monitored to ensure that particular crimes and/or victims are not disadvantaged by particular models.

## 7 Pre-Trained Language Models

With the recent proliferation and success of large pre-trained natural language models (e.g. BERT (Devlin et al., 2018)), it is natural to ask whether any of these models can be utilised in the contexts described above. Not only have these models proven powerful across a range of NLP tasks and domains (Lee et al., 2020; Chalkidis et al., 2020; Beltagy et al., 2019), but they also reduce some of the pre-processing burden such as feature engineering and embedding generation. For example, Hugging Face have recently introduced an autoNLP<sup>3</sup> service that allows access to high powered NLP models with very little training. While pre-trained language models are good candidates for facilitating POP through NLP, the ethical challenges discussed above remain pertinent. Commercial offerings of pre-packaged auto-NLP have the potential to be successful within police agencies, and are likely to offer good general results with a relatively low training burden. However, as suggested above, the richness and completeness of the data and the selection and usage of particular models are all potential sources of bias. To combat against these biases, users of the system must be able to understand the models, or be partnered with an agency that can, so that the models can be leveraged in an appropriate fashion. Police will need to delve beneath the surface of potential headline metrics to ensure that the models are not creating new, or perpetuating existing biases. If the police are ill-equipped to do this then it is, in our view, the responsibility of the academic community to investigate these potential problems before systems are used in an operational settings.

## 8 Societal Implications

Authors submitting to the NLP for Positive Impact workshop were challenged to define what they felt positive impact meant to them in the context of their work. Positive impact for us would be, firstly, the more wide spread adoption of problem oriented policing. This would see more police agencies devoting more of their time to proactive activities and thus to crime prevention rather than focusing on reactive detection and arrest of offenders. The positive societal impact of this would be less people embroiled in the criminal justice system, as the conditions for crime would not manifest them-

<sup>3</sup><https://huggingface.co/autonlp>

selves as often, and so the opportunities to commit crime would be reduced (Felson and Clarke, 1998). These may seem lofty aims for an analytical technique, and perhaps they are, but in this instance NLP would serve as one part of a new approach to understanding crime. NLP can be the key enabler to unlock the latent potential in a policing technique that will allow a shift away from the contentious response-arrest based policing style to a more balanced system. A balanced system that promotes preventing people, often young and disadvantaged, from becoming criminalised. A system that values a problem prevented over an arrest made or a person incarcerated. In this context, a positive impact would see police agencies more aligned with their communities needs and more focused on preventing crime harms before they occur.

## 9 Conclusion

Problem-oriented Policing (POP) can be an effective method for reducing crime. Empirical evidence suggests it is more effective than the traditional response model in many situations. However, the key requirement of effective POP is an understanding of the crime event, information that is often stored but is too resource intensive to extract from police administrative free text data. Here we have argued that NLP has the potential to be applied in a range of ways that could lower the analytical burden of police who seek to take a POP approach, thus enabling it to be adopted more extensively. Widespread adoption of POP has the potential to have a positive impact on society. By reducing opportunities for crime, POP is capable of reducing the societal harms that stem from both victimisation and offending. Moreover, the preventative approach advocated by POP relies less heavily on traditional arrest-based response method of policing which can create tensions between the police and local communities that they serve alongside producing a range of social and economic costs downstream.

NLP is not, however, without its drawbacks, and chief among these are the technical knowledge required to utilise the models and a need to account for potential biases. This all means that the introduction of NLP to police agencies will have to be carefully considered, with biases understood, quantified and addressed in ways that minimise undue harm. Generally speaking, police agencies do not have the expertise to do this themselves, and pri-



vate providers who might offer such expertise often have a vested interests in protecting their technologies which in turn reduces transparency. As such, it is incumbent on the academic community to investigate how NLP might support such policing efforts and better understand how the aforementioned challenges might be met prior to them manifesting in negative outcomes. If applied correctly and with appropriate safeguards, NLP has the potential to unlock the power of prevention-focused policing techniques, thereby reducing crime and the diverse societal harms associated with its occurrence.

## Acknowledgements

We thank the anonymous reviewers for their insightful and helpful comments which have undoubtedly served to enhance the paper. This research was supported by the Economic and Social Research Council.

## References

- M.P. Basilio, G.S. Brum, and V. Pereira. 2020. [A model of policing strategy choice: The integration of the latent dirichlet allocation \(lda\) method with electre i](#). *Journal of Modelling in Management*, 15(3):849–891. Cited By 1.
- M.P. Basilio, V. Pereira, and G. Brum. 2019. [Identification of operational demand in law enforcement agencies: An application based on a probabilistic model of topics](#). *Data Technologies and Applications*, 53(3):333–372. Cited By 2.
- Eric P Baumer. 2002. Neighborhood disadvantage and police notification by victims of violence. *Criminology*, 40(3):579–616.
- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. [Scibert: A pretrained language model for scientific text](#). *arXiv preprint arXiv:1903.10676*.
- Daniel Birks, Alex Coleman, and David Jackson. 2020. [Unsupervised identification of crime problems from police free-text data](#). *Crime Science*, 9(1):18.
- Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. Language (technology) is power: A critical survey of “bias” in nlp. *arXiv preprint arXiv:2005.14050*.
- Anthony A Braga, Andrew V Papachristos, and David M Hureau. 2014. The effects of hot spots policing on crime: An updated systematic review and meta-analysis. *Justice quarterly*, 31(4):633–663.
- Anthony A Braga, David L Weisburd, Elin J Waring, Lorraine Green Mazerolle, William Spelman, and Francis Gajewski. 1999. Problem-oriented policing in violent crime places: A randomized controlled experiment. *Criminology*, 37(3):541–580.
- Gian Maria Campedelli. 2020. [Where are we? Using Scopus to map the literature at the intersection between artificial intelligence and research on crime](#). *Journal of Computational Social Science*.
- Ilias Chalkidis, Manos Fergadiotis, Prodromos Malakasiotis, Nikolaos Aletras, and Ion Androutsopoulos. 2020. Legal-bert: The muppets straight out of law school. *arXiv preprint arXiv:2010.02559*.
- Lawrence E Cohen and Marcus Felson. 1979. Social change and crime rate trends: A routine activity approach. *American sociological review*, pages 588–608.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- John E Eck and William Spelman. 1987. Problem-solving: Problem-oriented policing in newport news.
- Graham Farrell. 2015. Crime concentration theory. *Crime prevention and community Safety*, 17(4):233–248.
- Marcus Felson and Ronald V Clarke. 1998. Opportunity makes the thief. *Police research series, paper*, 98:1–36.
- H. Goldstein. 1990. *Problem-oriented Policing*. Temple University Press.
- Herman Goldstein. 1979. [Improving policing: A problem-oriented approach](#). *Crime and Delinquency*, 25(2):236–258.
- Hossein Hassani, Xu Huang, Emmanuel S Silva, and Mansi Ghodsi. 2016. A review of data mining applications in crime. *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 9(3):139–154.
- Joshua C Hinkle, David Weisburd, Cody W Telep, and Kevin Petersen. 2020. Problem-oriented policing for reducing crime and disorder: An updated systematic review and meta-analysis. *Campbell Systematic Reviews*, 16(2):e1089.
- Sara Hooker. 2021. Moving beyond “algorithmic bias is a data problem”. *Patterns*, 2(4):100241.
- George Karystianis, Armita Adily, Peter Schofield, Lee Knight, Clara Galdon, David Greenberg, Louisa Jorm, Goran Nenadic, and Tony Butler. 2018. Automatic extraction of mental health disorders from domestic violence police narratives: text mining study. *Journal of medical internet research*, 20(9):e11548.

- George Karystianis, Armita Adily, Peter W Schofield, David Greenberg, Louisa Jorm, Goran Nenadic, and Tony Butler. 2019. Automated analysis of domestic violence police reports to explore abuse types and victim injuries: text mining study. *Journal of medical Internet research*, 21(3):e13067.
- Da Kuang, P Jeffrey Brantingham, and Andrea L Bertozzi. 2017. Crime topic modeling. *Crime Science*, 6(1):12.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2020. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240.
- Ben Marshall and Michael Townsley. 2006. Needles or needless? Technical report, Jill Dando Institute, UCL.
- Michael Scott, John Eck, Knutsson Johannes, and Herman Goldstein. 2016. Problem-oriented policing. In Richard Wortley and Michael Townsley, editors, *Environmental criminology and crime analysis*, 2 edition, chapter 11, pages 227–258. Taylor & Francis.
- Michael Scott and Stuart Kirby. 2012. Implementing pop: leading, structuring and managing a problem-oriented police agency.
- A Sidebottom, K Bullock, R Armitage, M Ashby, C Clemmow, S Kirby, Gloria Laycock, and Nick Tilley. 2020. Problem-oriented policing in england and wales 2019.
- Natalia Silveira, Timothy Dozat, Marie-Catherine de Marneffe, Samuel Bowman, Miriam Connor, John Bauer, and Christopher D. Manning. 2014. A gold standard dependency corpus for English. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC-2014)*.
- Roger Tarling and Katie Morris. 2010. Reporting crime to the police. *The British Journal of Criminology*, 50(3):474–490.
- Bruce Taylor, Christopher S Koper, and Daniel J Woods. 2011. A randomized controlled trial of different policing strategies at hot spots of violent crime. *Journal of experimental criminology*, 7(2):149–181.
- Nick Tilley. 2008. Modern approaches to policing: community, problem-oriented and intelligence-led. *Handbook of policing*, 2.