in brief...

Tackling domestic violence using machine learning

Artificial intelligence could help to protect victims of domestic violence, according to research by **Jeffrey Grogger**, **Ria Ivandic** and **Tom Kirchmaier**.

In England, domestic violence accounts for one-third of all assaults involving injury. A crucial part of tackling this abuse is risk assessment – determining what level of danger someone may be in so that they can receive the appropriate help as quickly as possible. It also helps to set priorities for police resources in responding to domestic abuse calls.

In recent work, we show that machine-learning methods are far more effective at assessing which victims of domestic violence are most at risk of further abuse than conventional risk assessments.

Currently, the risk assessment is done through a standardised list of questions – the so-called DASH form (Domestic Abuse, Stalking and Harassment and Honour-Based Violence) – which consists of 27 questions that are used to categorise a case as standard, medium or high risk. The resulting DASH risk scores have limited power in predicting which cases will result in violence in the future.

Following this research, we suggest that a two-part procedure would do better both in prioritising calls for service and in providing protective resources to victims with the greatest need.

Method and data

To build our predictive models, we use individual-level records on domestic abuse calls, crimes, victim and perpetrator data from the Greater Manchester Police. We combine this with DASH questionnaire data in order to forecast reported violent recidivism for victim-perpetrator pairs. This is defined to be any domestic abuse incident involving violence with injury or a sex offence, occurring within one year of a preceding call from the same pair.

In our analysis, we use two sets of variables. The first are the DASH responses and risk score coming from the victim and recorded by the responding officer for each domestic abuse call. These enquire about the victim's emotional state and the abuser's behaviour.

The second are the criminal and domestic abuse history variables covering the two years before the respective call. These include crime and domestic incident counts for both victims and perpetrators.

Overall, our sample consists of over 165,000 calls for service, covering the period from April 2014 to July 2018. We then aim to predict whether or not violent recidivism occurred within the following 12 months.

Machine learning offers a way to improve how police responses to domestic abuse calls are handled

Our predictive models are random forests, which are a machine-learning method consisting of a large number of classification trees that individually classify each observation as a predicted failure or non-failure. Importantly, we take the different costs of misclassification into account.

Predicting no recidivism when it actually happens (a false negative) is far worse in terms of social costs than predicting recidivism when it does not happen (a false positive). While we set the cost of incurring a false negative versus a false positive at 10:1, this is a parameter that can be adjusted by stakeholders.

Results

The random forest model based on the criminal history variables together with the DASH responses significantly outperforms the models based on DASH alone. The negative prediction error – that is, the share of cases that would be predicted not to have violence yet violence occurs in the future – is low at 6.3% as compared with an officer's DASH risk score alone where the negative prediction error is 11.5%. This is based on the 'test set', which is the data that were not seen during the model training phase.

The results show that a machine-learning-based approach to risk assessment outperforms the classical approach based on officer risk assessment alone. We also examine how much each feature contributes to the model performance. There is no single feature that clearly outranks all others in importance, but it is the combination of a wide variety of predictors, each contributing their own 'insight', which makes the model so powerful.

What can we learn?

This research could help to improve the way in which police responses to domestic abuse calls are handled. At the moment, the incoming call is handled by a call handler, who asks questions and assigns the incident a priority score. The urgency of the response is determined by the priority score – and an officer responds according to that priority.

DASH data are only available after an officer has appeared on the scene, but information about someone's criminal history is available as soon as the call comes in and the call handler has identified the parties involved.

This approach would lend itself to a two-stage risk assessment process, in which the first stage would use the criminal histories at the call handling stage, followed by the second stage after the officer has arrived at the scene.

This means that an initial prediction of violent recidivism could be made while the caller is on the line. Indeed, it could be used to set the priority score of the call, which would mean that the police would have a better chance of responding quickly to high-risk cases.

Conventional assessments of which victims of domestic violence are most at risk are outperformed by machine-learning methods

This article summarises 'Comparing Conventional and Machine-Learning Approaches to Risk Assessment in Domestic Abuse Cases' by Jeffrey Grogger, Ria Ivandic and Tom Kirchmaier, CEP Discussion Paper No. 1676 (http://cep.lse.ac.uk/pubs/download/dp1676.pdf).

Jeffrey Grogger of the University of Chicago is a research associate in CEP's community programme. **Ria Ivandic** is a postdoctoral research economist in CEP's community programme and part of the policing and crime research group. **Tom Kirchmaier** is director of CEP's policing and crime research group.