UNIVERSITY of PENNSYLVANIA

Department of Criminology

Working Paper No. 2016-2.0

Machine Learning Risk Assessment at Preliminary Arraignments for Domestic Violence

Richard A. Berk

Susan B. Sorenson

Geoffrey Barnes

This paper can be downloaded from the
Penn Criminology Working Papers Collection:
http://crim.upenn.edu

Machine Learning Risk Assessment at Preliminary Arraignments for Domestic Violence*

Richard A. Berk**, Susan B. Sorenson,*** & Geoffrey Barnes****

Risk assessments have a long history in criminal justice settings. Since the 1920s, parole release decisions in many jurisdictions have been informed by projections of how inmates would perform under supervision once released. Over time, these assessments have relied more heavily on actuarial methods and increasingly sophisticated statistical procedures. Sentencing decisions are undergoing a similar evolution from lengthy pre-sentencing reports to legislatively mandated statistical forecasts of risk. Beginning in the 1980s, pre-trial decisions have in many jurisdictions been shaped by actuarial forecasts of flight risk and of new crimes committed while an offender is awaiting trial.

These historical trends share two themes. First, forecasts of "future dangerousness" have long been undertaken informally to help insure public safety, and can, in principle, be made more accurate, more fair, and more transparent with the proper use of actuarial methods. Second, the yardstick for these methods is not perfection — no criminal justice forecasting procedure will ever be error-free — but conventional practice. The goal is to do demonstrably better with the understanding that some forecasting errors are inevitable. The same kind of yardstick applies to fairness and transparency. Perfection is much too high a standard. The intent herein is to improve current practice.

In our recent paper summarized here, we develop actuarial risk assessment procedures to be used at preliminary arraignments for domestic violence cases. In the jurisdiction studied, standard practice requires that shortly after an arrest, an individual is arraigned. A magistrate reads the charges and then decides whether the offender can be released — sometimes on bond or subject to certain conditions — until the date of his or her next scheduled court appearance. Under the Bail Reform Act of 1984, threats to public safety can be crucial to that decision. In effect, magistrates are required, often quickly and with relatively little information, to make a forecast of future dangerousness.

Because many domestic violence offenses are serial, there is justified concern about victim safety. Inaccurate forecasts of future dangerousness can dramatically increase the chances of subsequent domestic violence and can subject offenders to unnecessary and even damaging criminal justice interventions. Improving the accuracy of future dangerousness forecasts has the potential to reduce the risks for domestic violence victims and make the treatment of offenders more just.

We make an important distinction in theory and practice between intimate partner violence and domestic violence. For intimate partner violence, the victim and offender are, or have been, romantically and/or intimately involved with one another. In criminal justice parlance, a much wider range of familial and household relationships qualify as domestic violence. For example, an altercation between a man and his adult son is considered domestic violence.

For both kinds of offenses, "violence" means attempts to cause or to actually cause bodily injury and to rape or attempt to rape. Placing an individual in reasonable fear of those acts is included as well. It may be important to note that domestic violence also can include actions that some may not see as violent such as burglary, destruction of property, threats of economic retribution, threats against pets, witness intimidation, and arguments over child custody, many of which can be charged as other kinds of crimes. Consistent with the jurisdiction's practice and the state statutes under which magistrates operate, we use the term “domestic violence” in this inclusive fashion.

Intimate partner violence risk assessments have been undertaken for decades. However, recent efforts usually focus on the needs of *victims* and how best to provide services that improve victim well-being and safety. Our risk assessment procedures for domestic violence target the *offender* and how best to apply criminal justice interventions to prevent or deter new domestic violence incidents. The data used, statistical procedures applied, and forecasts made are, by design, and as described in the following sections, very different.

Research Methods

We applied machine learning procedures to data from over 28,000 domestic violence arraignment cases in a major metropolitan area. In this setting, machine learning is a modern actuarial method that will forecast as least as accurately, and usually more accurately, than traditional actuarial methods. Random forests was the particular machine learning procedure applied because it has performed well in past criminal justice forecasting applications and because it can easily incorporate different relative costs of different forecasting errors. All forecasting errors are not created equal. For example, in the eyes of many stakeholders, it is far worse to release an individual who will commit a brutal act of domestic violence than to detain unnecessarily an individual who poses no such risk. It follows that the random forests algorithm should require very strong statistical evidence before a defendant is forecasted to be a good risk. Weaker statistical evidence is required before a defendant is forecasted to be a bad risk. Conventional risk assessment tools treat all forecasting errors the same so that equivalent statistical evidence is used for both. Badly misleading forecasts can result.

Most conventional risk assessment tools consider two outcomes: any arrest for a new offense or no arrest at all. These days, criminal justice officials typically favor more nuanced approaches. We forecasted three possible post-arraignment outcomes within two years of an arraignment release: (1) a domestic violence arrest associated with a physical injury, (2) a domestic violence arrest not associated with a physical injury, and (3) no arrests for domestic violence.

Using arrests as the re-offense measure means many domestic violence incidents were no doubt missed. A domestic violence incident becomes a domestic violence offense when police are involved and determine to their satisfaction that a domestic violence crime has been committed. Then, an arrest must follow, and sometimes the offender is not easily found. Nevertheless, we used arrests as the re-offense measure because unreported crimes are not captured in official crime statistics and because an arrest is what triggers an arraignment. It does no good to provide forecasting procedures that require data that do not exist.

The two-year window was established to capture new domestic violence offenses between the arraignment and the final resolution of the case. The forecasting procedure was developed with “training data” and evaluated with “test data” not used in the procedure-building process. Honest performance assessments follow when they are conducted in this “out-of-sample” manner.

Results

Among the full set of arraignments heard over the period from which the data came, approximately 13% were for domestic violence. With one exception, offenders arrested for domestic violence had similar ages and numbers of priors compared to offenders arrested for other crimes. For example, both groups had an average of about 23 prior charges overall. But, 15% of those arrested for domestic violence had priors for domestic violence whereas far fewer – 5% – of the other offenders had domestic violence priors. Both figures are probably underestimates because a substantial fraction of prior DV incidents may have been charged as other crimes. Offenders arrested for domestic violence in the past typically had been arrested for other crimes as well, some of which could have perhaps been charged as domestic violence.

The forecasts were necessarily based on features of individuals (e.g., age, gender) and past contacts with the criminal justice system (e.g., prior jail incarcerations) that were routinely available in electronic form. No doubt a variety of potentially useful predictors was not included, such as whether the offender was employed or had drug problems. Nevertheless, when machine learning forecasts of no post-arraignment domestic violence arrests were made, they were correct about 90% of the time.

Approximately 20% of the individuals released at arraignment were arrested within two years for a new domestic violence offense. If at arraignments of DV cases the machine learning forecasts were used by magistrates to determine which offenders to release, as few as 10% might be re-arrested for domestic violence. In other words, the arrest failure rate could be cut nearly in half. Over a typical 12-month period in the jurisdiction studied, perhaps well over one thousand post-arraignment arrests for domestic violence could be averted.

The six predictors most responsible for forecasting accuracy were: (1) the number of prior charges for violent crimes; (2) the number of days spent in jail in the past; (3) the total number of prior charges for any crime; (4) age; (5) the age at which the offender was first charged as an adult; and (6) the number of prior charges for domestic violence. The direction of these effects was as one would expect. For example, individuals who were young, had more priors for violent crimes in general and domestic violence in particular, and whose adult criminal records began at an earlier age, were more likely to be rearrested for domestic violence, including domestic violence leading to injuries. Perhaps the only surprise was that the risk of re-offending did not decline as rapidly with age as it did for street crime. For street crime, the peak years are late teens and early 20s, after which there is precipitous drop in risk. For domestic violence, there is a substantial risk into the 30s and 40s.

<u>Implementation</u>

A common concern about the use of machine learning in criminal justice settings is whether local IT expertise is up to the task. At least in large urban jurisdictions, ample IT support exists or is ramping up. In the jurisdiction studied, machine learning forecasts have been successfully and routinely used in the Department of Probation and Parole for over five years. The court system is a little behind, but thanks in part to a combination of city and philanthropic foundation support, motivation and resources exist.

The forecasts are meant to supplement, not replace, a magistrate's discretion at arraignment. How the supplemental information will be used in practice is, at the moment, unknown. One can imagine that in some circumstances, the forecasts will determine a magistrate's decision unless there is strong evidence contradicting the forecasts. Alternatively, there will probably be situations in which the reverse holds: the forecast will matter only when the other evidence is seen as inconclusive.

Also undetermined is what should be done with domestic violence offenders forecasted to be at high risk for re-arrest. Detention in jail is very expensive and potentially destructive for those detained and their families. Under consideration in the jurisdiction from which these data come are a variety of diversion programs designed to address the behavior of offenders while maintaining the safety of victims. There also is the hope that for detained offenders, a second round of forecasts can be made as more information about each is collected. It is likely the pool of offenders forecasted to pose a serious risk can be significantly reduced.

Conclusions

In criminal justice settings, machine learning currently creates the best available prediction tools. Machine learning has been used to help inform a variety of criminal justice decisions including release on parole, the intensity of probation supervision, and the level of security in which prison inmates are housed. In contrast to traditional actuarial methods, machine learning relies on algorithms rather than models to discover how to make accurate forecasts and can take directly into account stakeholder assessments of the tradeoffs between different kinds of forecasting errors and policy constraints.

This research shows that machine learning can be a valuable tool at arraignments when decisions about domestic violence cases are made. Other work we have done on arraignments in the same jurisdiction indicates it can be at least as effective with other kinds of crime such as drug possession. There too, re-arrests can be reduced substantially.

* This synopsis is based "Forecasting Domestic Violence: A Machine Learning Approach to Help Inform Arraignment Decisions," published in the Journal of Empirical Legal Studies, Volume 13, issue 1, 94-115, March 2016.
** Department of Statistics, Department of Criminology, University of Pennsylvania
*** School of Social Policy and Practice, University of Pennsylvania
**** Department of Criminology, University of Pennsylvania