eLaw
Working Paper Series

No 2021/00ť - ELAW– 2021

\~{ruxuzs,mzp,\~qpuo¡u{z

O t mxxqzsq ,uz,O¦nq~o~u y q,^q qm~o t ,
P m¡mruom¡u{z

Custers, B.H.M.

Universiteit
Leiden
eLaw

4

Profiling and Predictions: Challenges in Cybercrime Research Datafication

Bart Custers

Introduction

Due to its online nature, each form of cybercrime typically leaves digital traces. Such data can be useful for criminal investigations, but also for cybercrime researchers. When large amounts of data, such as data on the characteristics, actions, behavior and whereabouts of suspects, witnesses, and victims, or data on *modi operandi* is available in digital formats, for instance, in police records, court cases, and other sources, it becomes possible to search for patterns in these datasets. Such patterns may reveal novel and unexpected insights from the raw data.

In order to keep at least some control and overview over such large amounts of data, new data analytics tools have been developed, such as data mining and machine learning, that allow for automated data

B. Custers (✉)
Leiden University, Leiden, Netherlands
e-mail: b.h.m.custers@law.leidenuniv.nl

A. Lavorgna and T. J. Holt (eds.), *Researching Cybercrimes*,
https://doi.org/10.1007/978-3-030-74837-1_4

analyses. Such techniques, categorized in supervised learning and unsupervised learning, may result in profiles, which are a characteristic or a collection of characteristics of people. This can concern individual profiles, such as for terrorists or serial murderers (Chifflet, 2014), or group profiles, such as risk profiles for aggression or violence (Harcourt, 2007; Schauer, 2003). All patterns discovered, particularly profiles, may contribute to making predictions, for instance, on who committed a crime, who may intend to commit a crime, which people are at increased risk of becoming crime victims, at which locations crime is more likely to happen, and so on.

The automated data analyses, and the profiling and predictions that may result from it, provide tremendous opportunities to gain new criminological insights. Furthermore, this can be done much faster and at a significantly larger scale than when human researchers have to go through these large amounts of data. At the same time, however, a strong focus on such quantitative, data-driven research methods may involve some challenges from methodological, practical, and ethical perspectives. This chapter will examine these challenges of big data research in cybercrime, particularly challenges related to profiling and privacy.

Big Data in Cybercrime Research

Automated Data Analyses

A traditional approach in criminology is to use a hypothesis-driven or theory-driven approach, in which a hypothesis is formulated on the basis of existing theory that is verified or falsified with relevant data, that sometimes has to be collected first. With the exponential increase of available data, very large amounts of data are available nowadays, allowing for a different approach, namely a *data-driven approach*. In this approach, comparable to explorative data analysis, the focus is not on specific hypotheses, trying to get confirmation or rejection of what is expected, but on looking at what the data are telling. Particularly when researching cybercrime this may be relevant, since due to its online nature, each form of cybercrime typically leaves digital traces, resulting in useful research data.

The very large amounts of data are often referred to as big data (see also Chapters 10 and 11), but big data is not only defined by its sheer volume, many terabytes, or even petabytes of data (Gandomi & Haider, 2015). Other challenging characteristics include its velocity, i.e., the fact that many data are real time or nearly real time, and variety, i.e., the fact that the data comes in many types and formats, such as text, numbers, images, videos, and sound (Laney, 2001).

Very large amounts of data usually do not allow for finding patterns via human intuition or overview. For that reason, many different tools for automated data analytics, usually based on algorithms, have been developed that can be used to disclose hidden patterns in large datasets (Calders & Custers, 2013). Typically, machine learning and data mining are such tools, of which many different forms exist, that allow for the automated extraction of patterns from large datasets, a process usually referred to as knowledge discovery in databases (KDD) (Fayyad et al., 1996). Data mining is an automated analysis of data, using mathematical algorithms in order to find new patterns and relations in data. Machine learning is the way in which computer algorithms improve themselves automatically through experience, i.e., on training data. Both these and other tools allow machines to show intelligence, which is why they are also important in the field of artificial intelligence (see also Chapter 12).

The use of these tools for automated data analyses may yield new knowledge extracted from the data (Adriaans & Zantinge, 1996). This knowledge usually consists of disclosing patterns and relationships in the data, sometimes novel and unexpected. What kind of patterns may be discovered depends on the types of tools used. These tools can be distinguished in supervised and unsupervised learning techniques (Bloch, 2019). The most important types of tools are regression, classification, and clustering tools. Classification is a supervised technique that requires the availability of pre-existing classes or categories, whereas regression and clustering are unsupervised techniques, with regression aiming to describe particular trends or regularities in the data and clustering aiming to build models by dividing datasets into clusters of homogeneous data records.

Profiling and Predictions

Regression, classification, and clustering techniques can all be used for profiling, i.e., the process of ascribing one or more characteristics (attributes) to individuals or groups of people (Custers, 2013; Hildebrandt & Gutwirth, 2008). Such profiles may contain typical attributes and forms of stereotyping offenders, suspects, witnesses, and victims. For instance, it may reveal that money mules are used in laundering cybercrime profits are typically from Eastern European countries (UNODC, 2014) or that people with low self-control and people participating in online forums are at increased risk of becoming consumer fraud victims (van Wilsem, 2011), or that people generally find it hard to assess which online protections actually safeguard them against online fraud (Jansen & Leukfeldt, 2016).

Risk profiles may not only relate to people or groups of people, but also to objects and locations of crimes. For instance, in the fight against cyberterrorism, risk profiles for critical infrastructure are commonly used to make threat assessments (Brown et al., 2006). Typical examples of profiling crime locations are heat maps for crime mapping (Khan et al., 2019; Weisburd & McEwen, 2015). Such information is often subsequently used for patrolling strategies (Mastrobuoni, 2017), a typical example of predictive policing (Ferguson, 2019).

Apart from the usefulness of risk profiles in preventing crime or identifying suspects, risk profiles are also commonly used on parole and probation assessments (Dressel & Farid, 2018; Hudson & Bramhall, 2005; Wright et al., 1984). Based on personality characteristics and characteristics of the crimes committed, risks can be assessed. A key element in this is predicting recidivism (Skeem & Lowenkamp, 2020). Typically, when models for such assessments are based on large amounts of data, this can yield more objective parole and probation decisions, pushing back any personal opinions, prejudice or subjective notions of the decision-maker that may play a role, or at least further inform decision-makers (Ball, 2011).

Another area in which profiling and predictions may be relevant is assessing evidence in courts (Vlek et al., 2015). When courts have to decide whether a suspect is guilty, there may be probabilities to deal with.

Typical examples are matches for fingerprints or DNA. The methods used are usually good but not perfect, meaning there may be false positives (the test result shows a match, but in reality there is no match) or false negatives (the test result shows no match, but in reality there is a match). In cybercrime, fingerprints and DNA are usually not part of the evidence that needs to be assessed in courtrooms, but similar issues may apply to identifying suspects. For instance, how likely is it that a suspect has used a particular computer from which malware was sent, or how likely is it that a particular e-mail or IP address belongs to a suspect. Courts are supposed to take such probabilities into account when determining guilt, but this is not an easy task. This brings in assessment problems that humans, including judges or juries, may have when dealing with probabilities and risks, including the so-called prosecutor's fallacy and the defense attorney's fallacy (Thompson & Schuman, 1987).

Challenges

The use of profiling and predictions brings along several challenges. Here we discuss three categories of challenges, i.e., methodological, practical, and ethical challenges. The methodological challenges focus on the creation of profiles and prediction models, the practical challenges focus on the use or usability of profiles and predictions, and the ethical challenges focus on any moral or societal concerns profiles and predictions may cause.

Methodological Challenges

From a methodological perspective, there are several challenges that the use of profiles and predictions may pose when researching cybercrime. These methodological challenges can be related to the data collection and preparation, to the profiling process, or to making any predictions.

It can be hard to collect data for cybercrime research. Although there should be data on each cybercrime, as it always leaves digital traces, it can be hard to locate and access such data. Data may be located on servers in other countries. Governments and private companies (such as hosting providers) may not be willing to provide data for various reasons. Apart from jurisdictional issues, there may also be language issues. But even before making any data requests, an initial problem may already be on which door to knock, as it may not be clear who has the relevant data. On top of this, obviously cybercriminals do not want their data to be disclosed and may have taken measures to prevent this, such as encrypting the data, splitting data over many different locations (e.g., in cloud computing), or deleting data to the extent they are able to do so.

If research data is available, another issue may be that the data is from different sources and in different formats. Any data that is to be used in automated analysis may need some form of preparation before it can be processed. Typical pre-processing techniques include discretization, missing value imputation, dimensionality reduction, and feature extraction and construction (Calders & Custers, 2013). Also the velocity of data (e.g., real time, streaming) may require specific tools.

The advantage of using large amounts of data is that typical issues normally encountered in sampling do not occur. Using big data analytics, all data can be used ('N = All' according to Mayer-Schönberger & Cukier, 2013), without any need for sampling. So questions on the representativeness of a sample and questions on minimum sample sizes can easily be avoided.

However, the profiling process can yield other challenges. In essence, profiling is a form of building models, and if done automatically, it can result in too few or too many patterns. This can easily be compared with using online search engines like Google or Yahoo: ideally, any search yields 3–5 answers, but if there are zero or a few thousand search results, this is undesirable. Apart from too many results, it may also be problematic that some results are non-novel or trivial. For instance, any correlation between size and weight of people may not be remarkable, nor is finding that people driving under influence are over 16 years old.

Another challenge is that of overfitting (Mooney & Pejaver, 2018). Overfitting is the production of a model or profile that corresponds too closely with the dataset or the people in it. If the model is as detailed as the number of people in the dataset, it is not a generalization. The model then contains more parameters than justified by the dataset. This may typically occur if datasets are too small, models too detailed, or both.

This makes clear that the choice of automated data analytics tools needs careful consideration before getting started. Some tools may be a better choice than others, depending on the data and the intended goals of any automated analysis. It could be argued that, if in doubt about the right choice of tools, several tools can be used (subsequently or in parallel), but this may significantly add up the required computing power and times required. Also, using more tools may yield (many) more results.

Practical Challenges

Although the use of profiles and predictions can be very valuable in many aspects when researching cybercrime, there can also be practical challenges regarding their usefulness and effectiveness. Perhaps the most important challenge is that profiles and predictions are never absolutely correct. They are models with a limited accuracy. Even though that accuracy can in some cases be very high, it is never perfect. Therefore, there may be reliability issues (Custers, 2003). Typically, each profile may yield false positives and false negatives when it is applied.

False positives are people in a profile, that should not actually be in the profile. For instance, when a profile shows that ransomware campaigns are ran by Eastern European cybercriminals, it does not mean that all Eastern European cybercriminals run ransomware campaigns. False negatives are the opposite, namely people that are not in a profile, that should actually be in the profile. In the same example, it would be cybercriminals from other countries than those in Eastern Europe running ransomware campaigns.

Limited accuracy can lead to incorrect conclusions, which may result in ethical issues like bias, prejudice, and discrimination toward particular groups of people (see next subsection). From a practical perspective, the main challenge is to determine which levels of accuracy are acceptable. This may depend on the context. For instance, in advertising, profiles that increase outreach to target groups with only a few percent already can make a huge difference. However, in most criminal law and criminology contexts, accuracy has to be high to avoid stereotyping, incorrect generalizations, and false accusations. For instance, when people are denied boarding a plane because they (incorrectly) match a terrorist profile, this may cause significant unnecessary trouble. In practical contexts, like law enforcement and criminal investigations, accuracy also needs to be high in order to be effective.

This leads to another practical challenge, namely that profiles can quickly become outdated. To illustrate this, suppose a risk profile is created suggesting that terrorists are young males with black beards wearing a djellaba. Apart from the fact this profile would be too general and inaccurate to be practical (and could be considered discrimination), actual terrorist could easily avoid matching this profile by shaving their beard and wearing something different. In fact, terrorist groups have even reverted to training female terrorists, to avoid such profiles (Jacques & Taylor, 2009). Apart from people adjusting their behavior, another reason why criminological profiles may 'run empty' and therefore become less effective is because people get caught. For instance, if a profile shows that people trafficking drugs wear white tennis socks under their suit and this profile turns out to be effective, border police may catch them. Once in prison, these people will no longer turn up in border controls, rendering the profile ineffective for border police after some time.

In order to deal with outdated profiles (and predictions based on them), it is imperative that profiles are continuously updated. Since building profiles is based on data and data analytics tools, it means that both the datasets and the tools for analysis need to be revised from time to time. In practice, many organizations tend to focus on building profiles, but have limited attention for updating them, which may result in tunnel vision and low-quality profiles and predictions. From a research perspective, it would be good to add expiry dates to profiles (or at least qualify limited validity), perhaps similar to confidence intervals in statistical data.

Since automated data analysis is data-driven, the focus is on statistical relationships. These may be indicative for causal relationships, but obviously not all statistical relations are causal relations. A data-driven approach may reveal novel patterns, which can be highly interesting, but focuses less on theory and causality. For this, additional work, using other research methods, may be necessary.

From a practical perspective, profiles based on a combination of automated data analytics and human expertise seem to be the most effective. Introducing expert knowledge into data analytics can be helpful to avoid too many or non-novel profiles, but too much expect knowledge may result in looking for assumed patterns in the data rather than looking at what the data is telling. Obviously, it can be challenging to find the right balance here.

A common phenomenon is that data-driven models are built on large numbers of parameters, but that does not always guarantee better results for accurate profiles and predictions. A typical example in this respect may be COMPAS, the decision support tool used by US courts to assess recidivism. Research has shown that despite COMPAS's collection of 137 features, the same accuracy can be achieved with a simple linear classifier with only two features (Dressel & Farid, 2018). In addition, COMPAS is no more accurate or fair than predictions made by people with little or no criminal justice expertise.

The use of profiles and predictions can also be challenging in courts. Whereas in criminal investigations, reasonable suspicion or probable cause may be sufficient to act, in courts convictions have to be beyond reasonable doubt. Reasonable suspicion and probable cause may be based (at least partially) on statistical evidence and, as such, can by their nature go hand in hand with probabilities. However, when dealing with evidence, particularly if the criterion of beyond reasonable doubt has to be met, there is obviously friction with statistical concepts like probabilities. Unless accuracy is very high, risk profiles may not weigh heavily as actual evidence when convicting suspects.

Ethical Challenges

Generally speaking, all datasets, particularly large datasets, contain errors. Parts of the data may be incorrect or incomplete. Furthermore, data may be biased, for instance due to the ways in which is it collected. Obviously, this may reduce the accuracy of profiles and predictions (according to the adage garbage in = garbage out). Apart from accuracy issues discussed above, this may also lead to ethical issues regarding equal treatment, privacy, and fairness that will be discussed here.

A typical example of bias is data collector bias (Berk, 1983), which is also common in criminological datasets. In most countries law regulates that the police can only collect data on suspects and convicted criminals (consider, for instance, EU Directive 2016/680 regulating the use of personal data in criminal law for all EU member states). This already creates a bias, as no data on non-suspects is available. If, for instance, the goal is to find profiles for perpetrators, this can only be done by contrasting characteristics of suspects and non-suspects, which is difficult if no data is available on the latter group.

This may also lead to self-fulfilling prophecies. A typical example of this may occur when surveillance of law enforcement agencies focuses on neighborhoods with ethnic minorities. The probable result of such a policy would be that law enforcement databases get filled with people from these ethnic minorities. This is a form of selective sampling. When these law enforcement databases are subsequently used to find patterns on which people are more prone to show criminal behavior, it may not be surprising to discover that people from these ethnic minorities may be profiled as showing increased levels of criminal behavior. However, since the data was biased, this is a mere self-fulfilling prophecy.

A related issue here is that some of the data analytics tools may be self-reinforcing, resulting in the amplification and further entrenchment of patterns. These effects may amplify existing bias and inequality in datasets when deriving profiles from it, undermine democracy, and further push people into categories that are hard to break out (O'Neil, 2016).

In fact, the data analytics tools may be biased themselves, in the way they are designed (Barocas & Selbst, 2016; Hutchinson, 2016) or by the training data provided to them. For instance, face recognition software still struggles to recognize black faces: even top performing facial recognition systems misidentify blacks at rates five to ten times higher than they do whites (Simonite, 2019).

All this may have profound impact from an ethical perspective. Any kind of bias in datasets or the design and use of data analytics tools may propagate and even amplify prejudice and discrimination. Although the use of profiles and predictions may avoid prejudice and discrimination of law enforcement on the ground (i.e., using evidence-based objective profiles may neutralize personal preferences and misjudgments), this will not work if the profiles are prejudiced and discriminating themselves.

Even without errors and bias, discrimination may be an issue when using profiles. For instance, particular attributes may appear in risk profiles that are not acceptable or even violating antidiscrimination laws, particularly when these criteria are used for decision making. This may concern particularly sensitive attributes like religion, political preferences, sexual preferences, criminal records, and gender. Research has shown that even when these sensitive attributes are not included in the datasets, they may appear by proxy (Calders et al., 2013). A typical example of such indirectly discriminating profiling is so-called redlining, in which characteristics are ascribed to people on the basis of their zip codes, whereas zip codes may be a strong indicator for someone's ethnic background. There are discrimination-aware data mining tools that can be used to avoid these discrimination issues (Zliobaite & Custers, 2016).

Apart from discrimination issues, the use of profiling and predictions can also be highly invasive for personal privacy. Typically, these methods can be very helpful in predicting characteristics of people they are not willing to disclose (Custers, 2012). For instance, Kosinski et al. (2013) show that, based on Facebook likes for movies, music, games, comments, etc., reliable predictions can be made about a person's gender, ethnic background, sexual orientation, religion, happiness, substance abuse, parental divorce, intelligence, etc. Furthermore, big data analyses may even predict attributes of people that they do not even know, such as their life expectancy, their risk to attract cancer, and so on. When trying to predict whether people will become criminals becomes more and more sophisticated, this may get close to dystopian perspectives like depicted in the 2002 movie Minority Report. Nevertheless, predictive policing is currently changing law enforcement (Ferguson, 2019).

It is often suggested that preserving privacy can be achieved by properly anonymizing datasets. However, removing key attributes such as name, address, and social security number of data subjects is insufficient to guarantee privacy; it is often still possible to uniquely identify particular persons or entities from the data, for instance by combining different attributes (Ohm, 2010).

The use of profiles may also yield chilling effects in society (Büchi et al., 2019). When particular profiles have become publicly known, people may want to avoid matching unfavorable profiles. This may not only concern criminals, but people in general if they do not want to be associated with particular groups. Also the fact that data is being collected on them and the fact people may be monitored may affect their behavior. To the extent this prevents people from committing crimes, this may be ok, but when it affects their rights and freedoms, it may be worrisome.

Profiles can be discriminating or privacy-invasive, but they can also be unfair in other ways (La Fors et al., 2019). Typically, profiles and predictions used in a law enforcement context can indicate increased likeliness for particular (groups of) people to commit crimes or having committed a specific crime. Such risk profiles can be considered accusatory and stigmatizing, in the sense that they cast suspicion on specific groups of people. For these groups of people, it may be hard to defend themselves against such practices for several reasons. For instance, it may not be clear on which data the profiles were built, the tools for analysis may be complex and therefore hard to challenge, and the subsequent decision-making processes are not always transparent. For these reasons, the use of profiles for such decision-making processes has been likened with the novel The Trial by Franz Kafka, in which the protagonist is arrested by government officials without knowing on the bases of which accusation, evidence, or underlying information (Solove, 2004).

Also the right to a fair trial can be under pressure. If a profile indicates increased crime risks for particular people, it may influence any existing unprejudiced, open-minded perspectives of law enforcement officers and judges and juries in courts. Apart from discrimination and stigmatization issues discussed above, profiles and predictions may put the presumption of innocence under pressure. Instead of assuming a person is innocent until proven guilty, assessors primed with such profiles may start off with bias and prejudice, even unintentionally and unaware of this.

Conclusions

Profiling and predictions can be very strong, useful, and effective tools in researching cybercrime. Based on all kinds of available data, criminological profiles can be built, for instance, on who committed a crime, who may intend to commit a crime, which people are at increased risk of becoming crime victims, at which locations crime is more likely to happen, and so on. Given the sometimes very large amounts of data, big data analytics (e.g., data mining, machine learning) is increasingly applied for the automated finding of hidden patterns. When applied to prospective characteristics, profiles may also be used as predictions. This may reveal novel patterns in datasets that can be highly interesting.

In this chapter, challenges of big data research in cybercrime, particularly with regard to profiling and predictions, were mapped. Three categories of challenges were identified: methodological, practical, and ethical challenges. The methodological challenges focus on the creation of profiles and prediction models, the practical challenges focus the use or usability of profiles and predictions, and the ethical challenges focus on any moral or societal concerns profiles and predictions may cause.

Methodological challenges can be related to the data collection and preparation, to the profiling process, or to making any predictions. Data collection can be hard due to an international context (e.g., due to unwillingness to share data) or due to technology (e.g., encryption, cloud computing). Data preparation can be complicated due to the large volumes, velocity (e.g., streaming data), and variety (e.g., different formats). The profiling process is in essence a delicate modeling process, which caveats like overfitting and finding merely non-novel, trivial results.

Practical challenges include the accuracy of profiles and predictions, which if limited can considerably reduce their usefulness and effectiveness. Limited accuracy can cause false positives and negatives, stereotyping, incorrect generalizations, and, in a law enforcement context, even false accusations. Profiles need to be updated continuously, as they can become outdated quickly. Since data-driven approaches focus on statistical relations, additional work may be needed for establishing causal relations and further development of criminological theories. Introducing human intuition and expert knowledge may significantly enhance the usefulness and effectiveness of profiles and predictions and reduce overly complex models. In courts, evidence criteria like 'beyond reasonable doubt' can be in tension with statistical concepts like probabilities and error margins.

Ethical challenges when dealing with profiles and predictions typically are equal treatment, privacy, and fairness issues. These can be caused by bias in the data or data analysis tools, which may propagate and even amplify existing prejudices. People may be pushed into categories that are hard to break out. If the profiles are based on particularly sensitive attributes, this may be considered discrimination or stigmatization. If data analytics are used to predict characteristics of people they are not willing to disclose, this may interfere with their privacy. Profiles may also yield chilling effects, if people want to avoid matching less favorable profiles. Fairness as a value can also get under pressure when the use of profiles and subsequent decisions is not transparent, making it hard for people to defend themselves. This may even interfere with the presumption of innocence in practice.

Altogether, it can be concluded that the use of profiles and predictions can be very valuable in law enforcement and criminology, particularly in cybercrime research. At the same time there are many caveats. On the one hand, this means that these new approaches and methods do not invalidate or set aside existing tools, but essentially are an addition to the criminologist's toolbox, providing new research opportunities. On the other hand, this means that these new tools should only be used after careful consideration and preparation, assessing their pros and cons before deciding they are the most appropriate tools to use in a given context. If not applied in the right way, profiles and predictions are better not used at all.

References

Adriaans, P., & Zantinge, D. (1996). *Data mining*. Addison Wesley Longman.

Ball, W. D. (2011). Normative elements of parole risk. 22 *Stanford Law & Policy Review, 395*.

Barocas, S., & Selbst, A. D. (2016). Big Data's disparate impact. 104 *California Law Review, 671*.

Berk, R. A. (1983). An introduction to sample selection bias in sociological data. *American Sociological Review, 48*(3), 386–398.

Bloch, D.A. (2019). *Machine learning: Models and algorithms*. Quantitative Analytics.

Brown, G., Carlyle, M., Salmeron, J., & Wood, K. (2006). Defending critical infrastructure. *Interfaces, 36*(6), 530–544.

Büchi, M., Fosch Villaronga, E., Lutz, Chr., Tamò-Larrieux, A., Velidi, S. & Viljoen, S. (2019). *Chilling effects of profiling activities: Mapping the issues*. Available at: https://ssrn.com/abstract=3379275.

Calders, T., & Custers, B. H. M. (2013). What is data mining and how does it work? In B. H. M. Custers, T. Calders, B. Schermer, & T. Zarsky (Eds.), *Discrimination and privacy in the information society.* Springer.

Calders, T., Karim, A., Kamiran, F., Ali, W., & Zhang, X. (2013). Controlling attribute effect in linear regression (pp. 71–80). In *Proceedings of 13th IEEE ICDM*.

Custers, B. H. M. (2003). *Effects of unreliable group profiling by means of data mining*. In G. Grieser, Y. Tanaka, & A. Yamamoto (Eds.), *Lecture notes in artificial Intelligence, Proceedings of the 6th International Conference on Discovery Science (DS 2003),* (Vol. 2843). Springer-Verlag.

Custers, B. H. M. (2012). Predicting data that people refuse to disclose. How data mining predictions challenge informational self-determination. *Privacy Observatory Magazine*, p. 3.

Custers, B. H. M. (2013). Data dilemmas in the information society. In B. H. M. Custers, T. Calders, B. Schermer, T. & Zarsky (Eds.), *Discrimination and privacy in the information society.* Springer.

Chifflet, P. (2014). Questioning the validity of criminal profiling: An evidence-based approach. *Australian & New Zealand Journal of Criminology*. https://doi.org/10.1177/0004865814530732.

Dressel, J., & Farid, H. (2018). The accuracy, fairness and limits of predicting recidivism. *Science Advances, 4*(1).

Fayyad, U. M., Piatetsky-Shapiro, G., & Smyth, P. (1996). The KDD process for extracting useful knowledge from volumes of data. *Communications of the ACM*, *39*(11).

Ferguson, A.G. (2019). Predictive policing theory. In T. R. Lave & E. J. Miller (Eds.), *The Cambridge handbook of policing in the United States*. Cambridge University Press.

Gandomi, A., & Haider, M. (2015). Beyond the hype: Big data concepts, methods and analytics. *International Journal of Information Management, 35*, 137–144.

Harcourt, B.E. (2007). *Against prediction: Profiling, policing and punishing in an actuarial age*. Chicago University Press.

Hildebrandt, M., & Gutwirth, S. (2008). *Profiling the European citizen.* Springer.

Hudson, B., & Bramhall, G. (2005). Assessing the "Other": Constructions of "Asainness" in risk assessments by probation officers. *The British Journal of Criminology*, 45(5), 721–740.

Hutchinson, Y. (2016, August 23). Biased by design. *MIT Technology Review*.

Jacques, K., & Taylor, P. J. (2009). Female terrorism: A review. *Terrorism and Political Violence, 21*(3), 499–515.

Jansen, J., & Leukfeldt, E. R. (2016). Phishing and malware attacks on online banking customers in the Netherlands: A qualitative analysis of factors leading to victimization. *International Journal of Cyber Criminology, 10*(1), 79–91.

Khan, M., Azhar, R. & Rauf, A. (2019). *Hotspot analysis of crimes using GIS: A case study of district Abbottabad*. Available at: https://ssrn.com/abstract=3312540.

Kosinski, M., Stillwell, D., & Graepel, T. (2013). Private traits and attributes are predictable from digital records of human behaviour. In *Proceedings of the National Academy of Sciences* (PNAS). Available at: www.pnas.org/content/early/2013/03/06/1218772110.

La Fors, K., Custers, B. H. M., & Keymolen, E. (2019). Reassessing values for emerging big data technologies: Integrating design-based and application-based approaches. *Ethics and Information Technology, 21*(3), 209–226.

Laney, D. (2001). *3D data management: Controlling data volume, velocity and variety*. Gartner. META Group.

Mastrobuoni, G. (2017). Crime is terribly revealing: Information technology and police productivity. *Review of Economic Studies* (online first).

Mayer-Schönberger, V., & Cukier, K. (2013). *Big data: A revolution that will transform how we live, work and think*. Harcourt Publishing Company.

Mitchell, T. M. (1999). Machine learning and data mining. *Communications of the ACM, 42*(11).

Mooney, S. J., & Pejaver, V. (2018). Big data in public health: Terminology, machine learning, and privacy. *Annual Review of Public Health, 39,* 95–112. https://doi.org/10.1146/annurev-publhealth-040617-014208.

Ohm, P. (2010). Broken promises of privacy: Responding to the surprising failure of anonymization. *UCLA Law Review, 57*, 1701.

O'Neil, C. (2016). *Weapons of math destruction; How big data increases inequality and threatens democracy*. Crown

Schauer, F. (2003). *Profiles, probabilities and stereotypes*. Harvard University Press.

Simonite, T. (2019, July 22). The best algorithms struggle to recognize black faces equally. *Wired*.

Skeem, J., & Lowenkamp, C. (2020). Using algorithms to address trade-offs inherent in predicting recidivism. *Behavioral Sciences & the Law, 38*, 259–278.

Solove, D. (2004). *The digital person; Technology and privacy in the information age*. New York University Press.

Thompson, W. C., & Schuman, E. L. (1987). Interpretation of statistical evidence in criminal trials: The prosecutor's fallacy and the defense attorney's fallacy. *Law and Human Behavior, 11*, 167–187.

UNODC. (2014). *Basic manual on the detection and investigation of the laundering of crime proceeds using virtual currencies*. Available at: https://www.imolin.org/pdf/imolin/FULL10-UNODCVirtualCurrencies_final.pdf.

van Wilsem, J. A. (2011). Bought it, but never got it: Assessing risk factors for online consumer fraud victimization. *European Sociological Review, 29*(2), 168–178.

Vlek, C., Prakken, H., Renooij, S. & Verheij, B. (2015). Constructing and understanding bayesian networks for legal evidence with scenario schemes (pp. 128–137). In *Proceedings of the 15th International Conference on Artificial Intelligence and Law*. ACM Press.

Weisburd, D. L., McEwen, T. (2015). *Introduction: Crime mapping and crime prevention*. Available at: https://ssrn.com/abstract=2629850.

Wright, K. N., Clear, T. R., & Dickson, P. (1984). Universal applicability of probation risk-assessment instruments. *Criminology, 22*, 113–134.

Zliobaite, I., & Custers, B. (2016). Using sensitive personal data may be necessary for avoiding discrimination in datadriven decision models. *Artificial Intelligence and Law, 24*, 183201.