Applications of Machine Learning to Pretrial Sentencing: A Literature Review

A Technical Report submitted to the Department of Computer Science

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree

Bachelor of Science, School of Engineering

Jalen Gregory

Spring, 2021

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Jundong Li, Department of Computer Science

Applications of Machine Learning to Pretrial Sentencing

A Literature Review

Jalen Gregory
School of Engineering and Applied Sciences, Computer Science
University of Virginia
Charlottesville, Virginia, US
jg6zy@virginia.edu

Abstract

No less beholden to the idea of efficiency as private entities, the public sector has turned to algorithms in the form of predictive analytic tools/instruments to expedite the decision-making process in criminal justice settings and save resources simultaneously. These algorithms, however, are often the proprietary property of the companies that develop them and thus their outputs are impossible to decode by the public. Much research has been done about big data policing as to whether the algorithms used can create fair outcomes or whether they simply reinforce the biases that they are built to eradicate. Still, the jury is out on the net impact from these tools: different researchers prioritize different factors and come to conclusions that are not necessarily mutually exclusive. This paper is a literature review of research concerning risk assessment tools in hopes of synthesizing the important points laid out by different researchers and finding the points where they agree work must continue. The main points of contention are whether the risk assessment tools provide a significant boost over the same process done by a judge, whether the tools are alleviating or exacerbating inequality, and if these tools are here to stay, what are the necessary steps to be taken to ensure that they can be trusted. The issues underlying the use of predictive analytic tools are some of the most important of our time, and as the worlds of criminal justice and technology continue to intermingle, citizens must have oversight lest what is supposed to be a democratic system quickly devolve into a technocracy.

Introduction

In the season 4 episode of Person of Interest "Pretenders", two computer scientists argue for their side of the issue under study at a conference on neuroevolutionary technologies. Beth Bridges, having sought out Harold Finch to accost him over the incorrect conclusions drawn from one of his papers, bristles at the idea that ethics should take precedence over innovation as it pertains to the field of predictive analytics. Noting her belief in the importance of proactivity, Finch responds that "precautionary principles create the only possible framework that allows us to see the implications of neuroevolutionary technologies". She retorts that "people have a right to science", and without a beat Finch answers back, "and to responsible scientists." [19]

This dichotomy is not new to computer scientists or engineers, but advances in information retrieval and processing techniques have brought it up in a new arena: using machine learning to enhance the predictive capabilities of various law enforcement entities. Despite its appearance as being created full cloth from the whims of the founding fathers, the American criminal justice system, among other entities, is constantly evolving in response to the ever-changing social landscape that gives it its credibility and from which it receives its orders. Like companies in the private sector, the government aims to leverage the most valuable commodity of our time, data, to create and improve systems that more efficiently and accurately present information to relevant parties [8]. For the police, these systems allow them to deploy their officers to not only respond quickly, but ideally to prevent violent crimes from taking place. For judges, machine learning tools attempt to compute the likelihood that the accused will go on to commit another crime if they were to be released pre-trial with or without bail. Despite the pervasiveness of algorithms in the daily lives of Americans, they go unnoticed, operating in the shadows, their creation and maintenance often left in the hands of the few instead of the many. The last decade has spawned much literature on the issue of using machine learning algorithms in criminal justice contexts, the review of which has raised three main issues to be settled before the affected stakeholders can reach any kind of closure on how to deploy these informational tools: a.) what inputs allow the algorithms to perform at their best, b.) how are the outputs intended to be interpreted, and c.) how do we define fairness and ensure we are holding ourselves to that standard? I hope that synthesizing the research in this area will provide a clear view of where we currently stand and what the next steps are.

Background

The underlying problems that risk assessment tools aim to alleviate (public safety concerns, deterrence of future crime, the inherent uncertainty in predicting human behavior) are not new at all. In other words, it is not the mission of the police or judges that has changed, it is the tools available to them. In the case of judges making bail decisions, a series of 20th century laws asserted that they may take into consideration not only the defendant's risk of flight but the harm that could be caused to society if they were to be released [7]. Existing in forms from simple linear regression tools to clinical evaluations to now self-learning algorithms, risk assessment tools take in information about the defendant and attempt to distill it into a probability value that indicates how at risk they are of recidivating. Many states have adopted risk assessment tools to use in pretrial decision making in response to the bail reform movement as well as budget concerns, hoping that their adoption will encourage non-carceral solutions that save taxpayers money [13].

Machine learning is not going anywhere, its benefits have been shown in various fields and the potential to unlock even greater knowledge lies ahead. For that reason, the idea that machine learning is used in general is not controversial, as much as its specific application in contexts where it may not be ideal. Their use in criminal justice settings does not call into question the science itself, only, as Finch noted, whether that science is being conducted responsibly. I believe a useful framework to discuss this is that put forth by Adebe et. al in their paper "Roles for Computing in Social Change". In this paper, the authors explain that there is a lack of precision involved when trying to bridge the gap between technical and normative concerns. To help in this field, they described 4 roles that computing can play as it is further engrained in non-technical areas of expertise: a.) Diagnostic (applying computation methods to measure flaws in present/future methods, b.) Formalizer (taking advantage of a computer's need for explicit specifications of inputs and outputs to better shape how problems are measured and understood), c.) Rebuttal (showing the inherent limitation of technical solutions on non-technical or even technical problems), and d.) Synecdoche (the tendency for advances in technology to shine new light on old problems) [1]. By viewing the current and future applications of machine learning algorithms in criminal justice through one or more of these roles, we can assess whether machine learning is being used as a positive force, or whether there is an appropriate application at all.

The Inputs

There are issues in relation to the input for risk assessment algorithms dealing with balancing the want to gain the most accurate information, and whether the information necessary to do that is constitutional. Richard Berk wrote *The Role of Race in Forecasts of Violent Crime* in 2009, addressing the effect of choosing to omit certain factors in the algorithm's classification process [2]. He argued that the algorithms should be given as much information as possible and be allowed to figure out what is pertinent as part of its training process [2, 3]. The problem is less machine than human: the purpose of using algorithms to predict who is most at risk of committing a crime instead of simply allowing judges to use their absolute discretion is because we believe that we can do better. If people are unable to accurately predict risk, it is fair to assume that a reason could be that they do not know what factors play into crime. Cleaning what we believe is useless data could serve to contaminate the result and/or lessen its accuracy. This was one of the conclusions drawn by Lin et. al, in showing that when making decisions on risk of recidivism people tended to ignore the base rate of recidivism for a group and make a ruling based on intuition [15].

One of the more popular tools for risk assessment, COMPAS, has come under fire for reasons particularly relevant to its applications of machine learning [4, 10, 11, 15, 17]. Examining the legality of risk assessment tools and the traits used to classify defendants, Sonja Starr called on precedent from the Supreme Court establishing that statistical discrimination is illegal and restricts the use of protected group identity to cases where it is materially relevant [17]. Abolishing statistical discrimination was done in service of protecting the American ideal that everyone is their own person and should not be judged by the actions of a group. It could be argued that the application of classification algorithms in criminal justice is incompatible with this ideal. Though her work was generally in response to check-list style risk assessment tools in use at the time, at a Penn Law Symposium on the issue she argued that a higher number of groups by which to analyze defendants could lead to a higher classification accuracy but did not solve the problem of needing individualized sentences [20]. When the Wisconsin Supreme Court read the case of Loomis vs. Wisconsin, they were forced to deal with some of these issues. The defendant argued that he was not being judged as himself, but as an amalgamation of people who fit his circumstances. In their ruling, they pointed out a few things that people would need to consider as these algorithms continue to gain prominence across the country: that people could fall prey to overestimating the strength of statements made by the software's output and that it was trained on national data and was not fully representative of the jurisdictions it was used in [10]. Corbett-Davies et. al elaborated on the latter point, showing that because state and local criminal data varied widely from national data, a simple linear regression could outperform COMPAS if trained on local data [6].

Regarding input, machine learning first serves as a diagnostic tool for the relevance of classification factors. For centuries criminologists have attempted to link various factors to the likelihood of criminal behavior. Immediately, through machine learning we can comb through scores of data, detecting links between factors that never would have been discovered by lesser statistical methods. This is a strength of machine learning as a tool, we can allow the algorithm to determine what it needs to know, or we can tune it in accordance with our values to ensure it sees what we want it to see [4]. A key example of this is exhibited by Kleinberg et. al in their investigation of the hidden, unobservable variables, or the idea that maybe judges are using intuition based on factors that are not put into the data. What they discovered comparing the results of judge decisions against the output of their regression is that whatever factors judges are observing that the model is not, do not add to accuracy, but detract from it [12]. This is exactly the kind of insight that machine learning can give, pointing out which are important factors in finding a result and which are superfluous. The importance of the inputs, however, means that it is important to utilize computing as rebuttal as well. The problems at hand are not problems that technology or statistics can solve alone. A long conversation is necessary about how much information will be fed and sifted through to find connections to possible criminal behavior. If the path is taken of predicting all crime through machine learning algorithms, it may be difficult to omit certain details as disparate beliefs of what makes a criminal come into play.

The Outputs

Aside from the legality of whether a truly individualized risk assessment can ever be attained, the outputs given from risk assessment tools have caused contention in several areas. The interdisciplinary nature of using machine learning to predict outcomes in criminal justice bands together experts from disparate fields that have little in common to one another. This is problematic because while algorithms can be tuned to produce relevant outputs to different fields, it remains in the hands of those experts to determine what to take from those results [13]. This has already proven to be an impediment in the adoption of risk assessment tools. In analyzing the impact of laws that mandated the use of risk assessment scores in pretrial release decisions, Megan Stevenson and the team of Konig and Wenzelburger came to similar conclusions [13,18]. Despite the significant gains in classification accuracy over human intuition, the net impact on crime have proven much less substantial.

A common fear among researchers turns out to be the opposite of what has shown to bear out in practice: that the legal profession's standard deferrals to "expert testimony" could lead to a series of decisions that are made almost without any human insight at all, blindly following algorithmic outputs. This is a major contention of Professor Starr, who fears that this deferential tendency could lead to the veneration of machine learning methods as infallible when they are far from it [17]. Technical experts once again would be wise to use computing as a tool of rebuttal in this case, explaining the limits of classification as a tool and warning against the possible implications of its use. As awe-inspiring as these algorithms can be, they are bounded in their accuracy by the data available to them. What is a strength in the flexibility in manipulating inputs can be a weakness as it pertains to outputs [4, 10]. Connections drawn between inputs can be difficult to reverse engineer, making it difficult to assess the algorithms in contexts outside of the accuracy of their output. A common fear in legal contexts is that in situations where protected group identities are not in use, proxies for them are substituted. Berk and Bleich highlight this as a weakness of the intersection of technical and normative concerns when they speak of ridding data of potentially valuable data points as a need to satisfy societal values [3]. Classification accuracy can drop significantly as more relevant factors are dropped from the pool, and people have been shown incapable of understanding the extent to which all these factors relate. The cause and effect between inputs and outputs are easier to comprehend, but the way that we shape the discussion of outputs has as much of an effect on the inputs.

Interpreting the output is more of a societal issue than a technical one, but computing can play a useful role as formalizer. The input to a risk assessment tool is a list of details about a defendant that may be useful in predicting their likelihood to commit another crime, finding the appropriate inputs can be difficult but the idea is simple: output is the opposite. The threshold at which one decides whether it is advisable to return to society varies from person to person. Kleinberg et. al defined the judge's decision to release as a function of how they weigh crimes committed relative to the costs, on behalf of the public through incarceration costs or on behalf of the defendant mentally, of incarceration [12]. Taking this into account was an important consideration in their methodology because they were testing both the classification accuracy of their algorithm and reasons for the difference in outcomes between judges and the algorithm.

Recognized as a significant barrier in the final acceptance of these tools by the public, transparency is an issue rampant in both computing and criminal justice in general [10, 14]. The lack of transparency in computing can be caused by either failure to effectively communicate the meaning of results or by the outright refusal to allow the public to verify results, for the maintenance of trade secrets or otherwise. Josh Kroll defines true transparency as the full availability of source code to appropriate entities and the full disclosure of input data, and partial transparency as the ability to disclose certain facts to certain people [14]. The application of these tools in the public domain should mandate that full disclosure is reached; however, public entities are not often developing the software for these tools on their own, relying on proprietary software from the private sector. This combination of effects threatens to dampen the impact of risk assessments in the best case and call the entire practice into question. Using outputs that can be difficult to understand based on methods generated by proprietary software that cannot be examined could serve to bolster the profile of dubious technical methods while simultaneously hurting the populations the tools were designed to protect. Computing as rebuttal can play an important role here once again, once again showing the need for normative and technical concerns to be faced in concert.

While he acknowledges the obvious importance of being able to apply the results in societally acceptable ways, Berk fears that these concerns are overstated and can be exploited to spread fear and doubt about the capability of risk assessment tools [5,10]. While not directly an output of the algorithm, the weight of the decision to take one's freedom away is heightened by the concretization of their risk in a single number, as are the consequences of letting a prospective violent criminal back on the streets. After decades of protests and criticism against the results of less analytical methods, is it wise to sacrifice good in pursuit of perfection? In most situations, the answer is a resounding yes. However, this could result in a shift of responsibility from judges to teams of computer scientists and statisticians having outsized impact on decisions made on the freedom of others. Whether computer scientists would be willing to accept that role is an entirely different question than should they be trusted with the responsibility. Judges have direct expertise in deciphering legal issues but are human and have other concerns. The rate at which judges used risk assessments in different jurisdictions is often tied to whether they can afford to be held liable for a mistake in preparation for their next election campaign [13]. The distance that computer scientists have from the impact of risk assessment tools can be a positive in the sense that they have less directly applicable bias but can also serve to separate them from the concerns of those the tools will affect.

Fairness

Though machine learning algorithms can appear to be an impartial middleman between data and predicted outcomes from the outside, much research has been done on the issue of whether these algorithms are fair in the assessments that they provide. An immediate issue that arises is that there are multiple ways to define fairness, and each comes with its own impact on the output of the algorithms. One attempt at fairness sees risk scores measured within percentiles of the race of the defendant and then calibrating the algorithm so that the output aims to create a result where the proportion who reoffend is the same among different races [6]. This could help to alleviate racial disparities that plague sentencing as it stands, but not without cost. A simulation of this in Broward County resulted in higher incarceration rates for lower risk defendants and increases in violent crime because of the current imbalances in sentencing [6]. Calibrating the algorithms for more equitable but less effective outputs could result in the loss of public faith in them and end up wiping out whatever gains they could have had. Separating datasets by race and using the machine learning technique of random forests showed similar outcomes, as well as the impossibility of optimizing for either a combination of public safety and fairness or even different kinds of fairness [5]. Though the impossibility to achieve perfectly fair outcomes can be discouraging, it is at least encouraging that the extent to which these definitions of fairness impact one another and public safety can be diagnosed.

Berk et. al identified three potential solutions to the issue of developing fair algorithms: pre-processing, in-processing, and post-processing [5]. A pre-processing step would attempt to isolate problematic inputs and remove them from the data before the "random forests" procedure is run. This would solve the problem of having potentially unconstitutional or unsavory inputs being used in classification but could lower classification accuracy if they are relevant and affect outcomes. Another potential solution taken in pre-processing would be the calculation of the distance between original and transformed predictors to be used as new inputs to preserve as much information as possible. In-processing solutions would attempt to achieve statistical parity by identifying classifications that the software was less confident in and flipping the outcomes to maintain equality across races. Post-processing solutions could follow a similar process as in-processing solutions, flipping the algorithm's output essentially at random to maintain equality among groups. This implicit application of machine learning as formalizer shows the need to establish universal and applicable definitions of fairness if progress is to be made in reaching it.

Achieving fairness is made even more difficult when attempting to place algorithms in the context of the criminal justice system. Ben Green stated it perfectly in saying that "many aspects of society have been measured only in limited ways, other aspects of society resist quantification, and the data that exists reflects the biases and power dynamics that have led to certain aspects of society being measured at all" [9 This is important to consider because even if a consensus is reached allowing all data available to be treated as input in classification, it cannot be forgotten that the collection of that data also reflects the society that it was collected in. Until a machine like that in Person of Interest is developed that can predict without flaw, it will be up to computer scientists to make clear that the algorithms are going to draw conclusions from the data to make the best classifications that it can, even if they are not ones that it is strongly confident in.

Few virtues are shared universally, but in terms of risk assessments, fairness just might be. Criminal justice being a public sector entity means that algorithms and risk assessment tools can fall victim to the political discourse that surrounds their adoption and/or use. Risk assessments are an easy political target because they are not widely understood, and this allows non-experts the opportunity to define them in their own terms. In politics, these terms tend to focus either on the threat to public safety posed if someone is misclassified as low risk and commits a crime or the affront to fairness it is that opaque machines are making decisions on the lives of US citizens. As it turns out, in at least 3 states, attempts to strike down risk assessment legislation by referring to their outcome as letting criminals roam the streets seriously threatens their adoption, while framing the concern as one based on fairness leads to attempts at developing framework to make them fairer [13]. This desire to strive for improved fairness in political and legal outcomes give a point to rally around as research on risk assessment tools continues.

Looking Forward

Seen as experts as one of the most significant developments in the criminal justice system since the admission of expert evidence, machine learning algorithms will only serve to further engrain risk assessment tools in pretrial decision making [16]. While it is crucial to continue study of the impact of risk assessments on rates of recidivism across the country, the power of data makes it just as necessary to study the impact of these algorithms on legal and social ways of thinking. The persuasive power of algorithms lies not only in the ability to generate accurate predictions, but in their ability to mask non-technical factors and their role in these outcomes. Settling for good in pursuit of perfection can lead to consequentialist thinking that is then reinforced by the accurate outcome prediction. Risk assessment tools can be a great tool in the pursuit of reducing crime if used correctly, but what if reduction of crime through incarceration is not the ideal path? Starr questions this when she asks if machine learning algorithms can consider the effect that non-carceral solutions or treatment programs may have had on one's likelihood of reoffending, or the impact that incarceration will have on future crimes [17]. It is entirely possible that there is a solution to these problems that does not necessitate the removal of one's freedom. In this case, what would be touted as incremental gains could make it more difficult to achieve true progress by giving the illusion that the problem is already solved.

The limitations of computing may make it so that the most important choices made in the realm of risk assessments have little to do directly with applications of machine learning. Fairness is not an issue that can be solved through data, but only through discussions of pros and cons between stakeholders. First, it will have to be defined in a way that balances their various concerns. The most difficult step in this process according to Berk will be balancing calibrating the algorithm across groups so that different groups achieve similar outcomes and calibrating the tools within groups, and that any discussion of fairness must be understood strictly in relation to the output of the algorithm and not post-hoc as a reflection of outcomes [5]. Much work will have to be done in gaining the trust of the public to stand aside as judicial decisions are thrown more often into the area of machine justice, judges who are used to wielding considerable decision-making power in pretrial decision and may be loathed to relinquish it, and defendants who will surely claim that punishment for what a machine believes they are likely to do is not fair. These hurdles can only be cleared if risk assessment tools are refined to work in specific jurisdictions and the citizens within them can agree on how they want them to be used. Machine learning will serve in each role of computing in stages of this process.

Conclusion

Risk assessment tools and the machine learning algorithms that underly them are one of the greatest examples of the applications of technical solutions to normative problems and can thus guide computing in improving society. As a diagnostic tool, we can for the first time quantify things like the relationship between personal factors and judicial outcomes and simulate the outcome of favoring certain kinds of fairness or public safety. The nature of machine learning forces conversations to advance past vague rhetoric into action, for better or worse. Through them, people will be made to see that technology alone cannot solve all problems, and only through that understanding can those applications be sure to minimize harm. Lastly, risk assessment tools can change the calculus behind a life changing decision from a judge who has spent years gaining legal expertise to one that has been trained on decades of data, just in a very different way. The implications of allowing an algorithm to make such a weighty decision as removing one's freedom could potentially go to show the lack of structure by which pretrial decisions have been made in the past. As potent as the potential for scientific breakthroughs can be in their applications in different walks of life, they can only reach the public effectively and safely in the hands of responsible scientists.

References

[1] Rediet Adebe et. al, 2020. Roles for Computing in Social Change. *FAT* '20: Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, (Jan 2020), 252-260. DOI: https://doi.org/10.1145/3351095.3372871.

[2] Richard Berk, 2009. The Role of Race in Forecasts of Violent Crime. DOI: https://doi.org/10.1007/s12552-009-9017-z. Retrieved from https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.297.9852&rep=rep1&type=pdf.

[3] Richard Berk and Justin Bleich, 2014. Forecasts of Violence to Inform Sentencing Decisions. *Journal of Quantitative Criminology*, (Mar 2014), 79-96. DOI: https://doi.org/10.1007/s10940-013-9195-0. Retrieved from http://homepage.ntu.edu.tw/~r08a21037/print_pdf/paper/Machine%20Learning%20Forecasts%20of%20Risk%20to%20Inform.pdf.

[4] Richard Berk and Jordan Hyatt, 2015. Machine Learning Forecasts of Risk to Inform Sentencing Decisions. *Federal Sentencing Reporter 27, no. 4*, (Apr 2015), 222-228. DOI: https://doi.org/10.1525/fsr.2015.27.4.222. Retrieved from https://www.researchgate.net/publication/275352025_Machine_Learning_Forecasts_of_Risk_to_Inform_Sentencing_Decisions.

[5] Richard Berk et. al, 2017. Fairness in Criminal Justice Risk Assessments: The State of the Art. arXiv: 1703.09207v2. Retrieved from https://arxiv.org/abs/1703.09207v2.

[6] Sam Corbett-Davies et. al, 2017. Algorithmic decision making and the cost of fairness. *KDD '17: Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, (Aug 2017), 797-806. arXiv: 1701.08230. Retrieved from https://arxiv.org/pdf/1701.08230.pdf.

[7] Charles Doyle, 2017. *Bail: An Overview of Federal Criminal Law.* Congressional Research Service. Retrieved from https://fas.org/sgp/crs/misc/R40221.pdf.

[8] The Economist, 2017. The world's most valuable resource is no longer oil, but data. Retrieved from https://www.economist.com/leaders/2017/05/06/the-worlds-most-valuable-resource-is-no-longer-oil-but-data.

[9] Ben Green, 2018. "Fair" Risk Assessments: A Precarious Approach for Criminal Justice Reform. Retrieved from https://scholar.harvard.edu/files/bgreen/files/18-fatml.pdf.

[10] Danielle Kehl, Priscilla Guo, and Samuel Kessler, 2017. Algorithms in the Criminal Justice System: Assessing the Use of Risk Assessments in

Sentencing. Retrieved from https://dash.harvard.edu/bitstream/handle/1/33746041/2017-07_responsivecommunities_2.pdf.

[11] Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan, 2016. Inherent Trade-Offs in the Fair Determination of Risk Scores. arXiv: 1609.05807. Retrieved from https://arxiv.org/abs/1609.05807.

[12] Jon Kleinberg et. al, 2017. Human Decisions and Machine Predictions. *The Quarterly Journal of Economics* vol. 133, no. 1, (Feb 2018), 237-293. DOI: https://doi.org/10.1093/qje/qjx032. Retrieved from https://pubmed.ncbi.nlm.nih.gov/29755141/.

[13] Pascal D. Konig and Georg Wenzelburger, 2021. When politicization stops algorithms in criminal justice. *British Journal of Criminology.* Oxford University Press. DOI: https://doi.org/10.1093/bjc/azaa099.

[14] Josh Kroll, 2017. Accountable Algorithms. Video. (17 May 2017). Retrieved from https://www.youtube.com/watch?v=KfAL6ZaExrI.

[15] Zhiyuan "Jerry" Lin et. al, 2020. The limits of human predictions of recidivism. *Science Advances*. DOI: https://doi.org/10.1126/sciadv.aaz0652.

[16] Rik Peeters and Marc Schuilenburg, 2018. Machine justice: governing through the bureaucracy of algorithms. *Information Polity* 23, (2018), 267-280. DOI: https://doi.org/10.3233/IP-180074. Retrieved from https://marcschuilenburg.nl/wp-content/uploads/2018/08/MachineJustice.pdf.

[17] Sonja Starr, 2014. Evidence-Based Sentencing and the Scientific Rationalization of Discrimination. *Stanford Law Review 66, no. 4.* Retrieved from https://repository.law.umich.edu/articles/1190/.

[18] Megan Stevenson, 2017. Assessing Risk Assessment in Action. *Minnesota Law Review* no. 303. DOI: https://doi.org/10.2139/ssrn.3016088. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3016088.

[19] Steven Surjik. 2014. Pretenders. *Person of Interest* 4,6 (October 2014).

[20] University of Pennsylvania Carey Law School, 2016. Quattrone Spring Symposium 2016: Risk Assessment Tools: Capabilities, Benefits, and Risks. Video. (29 Jun 2016). Retrieved from https://www.youtube.com/watch?v=wehHGEQG8b4.