Masthead Logo

Brigham Young University Prelaw Review

Volume 33 Article 4

4-2019

Auditing Predictive Policing

Jeremiah Scanlan
Brigham Young University, miahscanlan@gmail.com

Auditing Predictive Policing

*Jeremiah Scanlan**

There is a young man lying on the floor, killed by two rounds fired into his back. A police officer is questioning the witness and gathering what information he can from her shaky recollections: approximate height, weight, and facial features. He takes note of the neighborhood's placement in between battling gangs. At the police station, he collects more information recorded by camera footage, gunshot detection systems, and other electronic sources to better understand the context of the crime. Until now, this is how we would expect a twenty-first century investigation to proceed.

However, times have changed, and these data points are not analyzed by the officer alone. Now this information is plugged into a complex algorithm. The algorithm sifts through the information and compares it with thousands of data points gathered over decades of traditional police work as well as data from social media. The algorithm then puts out a prediction: these are your most likely culprits. The prediction list rides with the officer on his patrols. The next time he sees someone from this list, he will take special care to watch the suspect's actions. The officer might even preemptively visit a name from the list in the suspect's home, bringing a warning of the consequences of breaking the law.

This technology is not from a science fiction movie like Steven Spielberg's *Minority Report*. These technologies have already been employed in cities such as Los Angeles, Chicago, and Kansas City, Mo.[1] "Predictive policing" algorithms are the tool of the future for police departments and more police departments are using them every year.[2] Meanwhile, the race to examine the legal implications is playing catch up.

One implication stands out: the possibility of bias working its way into a predictive policing algorithm's decision-making process. In 2016, a coalition of organizations headed by the ACLU expressed concerns about the possibility of racial bias and lack of transparency

---

* Jeremiah Scanlan is a Junior at Brigham Young University majoring in international relations. The author would like to thank Devon Allgood and Forrest Albiston for their great help and hard work in editing the paper, as well as Eric Jensen for his helpful advice and counsel.

[1] *Serve and predict*, ECONOMIST, May 5, 2018, at 26; Justin Jouvenal, *Police are using software to predict crime. Is it a 'holy grail' or biased against minorities?*, WASHINGTON POST (Nov. 17, 2016), https://www.washingtonpost.com/local /public-safety/?utm_term=.0362fc8f95bd (search for article title).

[2] CMTY. ORIENTED POLICING SERVS., U.S. DEP'T OF JUSTICE, FUTURE TRENDS IN POLICING 2-7 (2014) https://www.policeforum.org/assets/docs/Free_Online_Documents/Leadership/future%20trends%20in%20policing%202014.pdf.; Cory Doctorow, *Is this the full list of US cities that have bought or considered Predpol's predictive policing services?*, BOING BOING (Oct. 30, 2018, 6:00 AM), https://boingboing.net/2018/10/30/el-monte-and-tacoma.html; PREDPOL, https://www.predpol.com/about/ (last visited Jan. 9, 2019) (predictive policing company Predpol claims its software "is currently being used to help protect one out of every 33 people in the United States").

around the programs.[3] However, the possibility of bias has largely passed unnoticed by the public and policymakers.[4]

Some alert experts have offered solutions to police predictive policing,[5] but the ideas have largely been vague. If predictive policing is to be held accountable, there must be a model of legal mechanisms to allow for transparency and accountability. This article proposes that legislatures should design statutes that require frequent external audits of predictive policing software as used by police departments.

In Part I, I will give a brief history and an explanation of predictive policing. Part II will examine the potential for bias even when data or algorithmic decision-making is not coded to be explicitly biased. In Part III, I will explain how audits are a practical solution to problems of transparency and accountability, as well as how these audits should be designed. Simply auditing the algorithm is not enough to inform the public, which is why Part IV will explain the need for publicly released audit reports and the information that should be included in them.

I. What is Predictive Policing?

Prediction has always been a part of the effort to prevent crime. In a typical criminal investigation, the police try to understand their targets better by putting together the facts and studying them for patterns. Predictive policing has the same goal, but with modern tools such as specialized statistical techniques, big data, and the computational power beyond that of a normal human being. As Andrew Ferguson puts it, this is "…more a shift in tools than strategy."[6]

The first efforts to use statistics to study crime can be traced back to the 1920s, when statistical methods were tested to predict how likely a given individual was to relapse into crime. Now, over half of US states use some form of statistical method to predict parole recidivism.[7]

New statistical techniques were then developed to work with geographic data. Police departments hired analysts to look at the data and identify crime "hot spots," which allowed departments to focus their resources more efficiently.[8] In the early 2000's, Chief William Bratton in Los Angeles worked with a research team at the University of California, Los Angeles to create a software program based on these techniques. This software eventually became PredPol, a leading prediction software package.[9]

---

[3] ACLU, PREDICTIVE POLICING TODAY: A SHARED STATEMENT OF CIVIL RIGHTS CONCERNS (2016), https://www.aclu.org/other/statement-concern-about-predictive-policing-aclu-and-16-civil-rights-privacy-racial-justice.

[4] Doctorow, *supra* note 2.

[5] Andrew Guthrie Ferguson, *Policing Predictive Policing*, 94 WASH. U. L. REV. 1109 (2017) [hereinafter *Policing*]; Andrew Guthrie Ferguson, *Big Data and Predictive Reasonable Suspicion*, 163 U. PA. L. REV. 327 (2015); Andrew D. Selbst, *Disparate Impact in Big Data Policing*, 52 GA. L. REV. 109 (2017).

[6] *Policing*, *supra* note 5, at 1123.

[7] BERNARD E. HARCOURT, AGAINST PREDICTION 39-41 (2007).

[8] KEITH HARRIES, U.S. DEP'T OF JUSTICE, MAPPING CRIME: PRINCIPLE AND PRACTICE 1-3, 112 (1999) https://www.ncjrs.gov/pdffiles1/nij/178919.pdf.

[9] WALTER L. PERRY ET AL., RAND CORP., PREDICTIVE POLICING: THE ROLE OF CRIME FORECASTING IN LAW ENFORCEMENT OPERATIONS 5 (2013).

Today, rather than just plotting out hot spots, many software packages focus on predicting individual offenders and victims.[10] Police departments across the country are now adopting these prediction programs or developing their own.[11]

Naturally, the first question to ask is whether predictive policing is effective. Unfortunately, there is still no clear answer. Police departments such as the Chicago Police Department tout lower crime rates, claiming that the reduction is a direct result of predictive policing.[12] However, the few independent academic studies on predictive policing have shown mixed results, and have not established that it has significant effect on reducing crime.[13] The software companies themselves have paired up with academics to perform studies that have shown that predictive policing has reduces crime, but this also raises conflict of interest issues in play that cast a shadow on the credibility of the results.[14]

Companies and departments have been reluctant to share any details of the algorithmic process to keep secrets away from competitors or anyone that would attempt to "game" the system.[15] Nonetheless, these concerns have not stopped companies from marketing predictive policing as a panacea for crime, nor stopped departments from enthusiastically adopting it because that might be true.[16]

Another concern is that predictive policing could exacerbate problems of bias among the police. Skeptical activist groups such as the ACLU have issued statements of concern that

---

[10] Maha Ahmed, *Aided by Palantir, the LAPD Uses Predictive Policing to Monitor Specific People and Neighborhoods*, INTERCEPT (May 11, 2018, 7:15 AM), https://theintercept.com/2018/05/11/predictive-policing-surveillance-los-angeles/; Monica Davey, *Chicago Police Try to Predict Who May Shoot or Be Shot*, N.Y. TIMES, May 23, 2016, at A11.

[11] *See supra* note 2.

[12] *See Serve and predict, supra* note 1

[13] PRISCILLA HUNT ET AL., RAND CORP., EVALUATION OF THE SHREVEPORT PREDICTIVE POLICING EXPERIMENT 49 (2014) https://www.rand.org/pubs/research_reports/RR531.html; Jessica Saunders et al., *Predictions Put into Practice: a quasi-experimental evaluation of Chicago's predictive policing pilot,* 12 J. EXP. CRIMINOLOGY 374 (2016) https://doi.org/10.1007/s11292-016-9272-0.

[14] G. O. Mohler et al., *Randomized Controlled Field Trials of Predictive Policing*, 110 J. AM. STAT. ASS'N 1399, 1408-1410 (2015) https://doi.org/10.1080/01621459.2015.1077710.

[15] Darwin Bond-Graham & Ali Winston, *All Tomorrow's Crimes: The Future of Policing Looks a Lot Like Good Branding*, SF WEEKLY (Oct. 30, 2013), https://archives.sfweekly.com/sanfrancisco/all-tomorrows-crimes-the-future-of-policing-looks-a-lot-like-good-branding/Content?oid=2827968&showFullText=true; Ali Winston, *Palantir Has Secretly Been Using New Orleans To Test Its Predictive Policing Technology*, THEVERGE.COM (Feb. 27, 2018, 3:25 PM), https://www.theverge.com/2018/2/27/17054740/palantir-predictive-policing-tool-new-orleans-nopd; Yana Kunichoff and Patrick Sier, *The Contradictions of Chicago Police's Secretive List*, CHI. MAG., Aug. 21, 2017 https://www.chicagomag.com/city-life/August-2017/Chicago-Police-Strategic-Subject-List/.

[16] Ellen Huet, *Server and Protect: Predictive Policing Firm PredPol Promises To Map Crime Before It Happens*, FORBES (Mar. 2, 2015), https://www.forbes.com/sites/ellenhuet/2015/02/11/predpol-predictive-policing/#9432d1b4f9bf, *see also supra* note 15.

predictive policing could be unfairly biased against minorities groups.[17] One non-peer-reviewed study by ProPublica found evidence of racial bias in related parole recidivism algorithms in Florida, but this is so far the only study on the subject.[18] Newspaper reporters have also begun to mention the possibility of bias when reporting on predictive policing.[19]

The legal literature on predictive policing is growing, but still sparse. Some authors have offered solutions to problems of bias and accountability, but these solutions have often been general suggestions rather than concrete specifics.[20] The lack of legal literature is mirrored by a lack of attention by reporters and the general public.

Some researchers have written about the possibility of algorithmic bias outside of predictive policing, for instance in credit scoring or hiring algorithms.[21] Scholars and experts continue to debate the best methods to discover and counteract algorithmic bias. Yet because algorithms are so complex and the field is relatively new, there is still no consensus on what the best methods are. However, many experts agree that something must be done to prevent bias.

Predictive policing is not just the future, it is the present; the hypothetical proposed in the introduction could happen today in Chicago, or tomorrow in the next city to buy predictive policing software. As the technology continues to advance, algorithms assist more in the decision-making process, and will perhaps make decisions on their own. These decisions may disproportionately affect communities and individuals based on skin color, gender, age, or other characteristics, even when that is not the algorithm's designed intent.

Predictive policing is here, but its effects have not been adequately measured and problems of transparency and accountability have not been addressed. This is the heart of the problem that this article seeks to address and the hole it seeks to fill. The effects of predictive policing cannot be measured adequately if companies and police departments make it difficult for future experts to review what predictive policing algorithms do. Thus, this article proposes regular audits of predictive policing algorithms in order to catch and prevent bias.

II. The Potential of Predictive Policing Bias

The potential for bias is not unique to predictive policing. Activists, technology experts, and policy makers have discussed the possibility of bias in algorithms, big data, and machine learning broadly, as well as specific applications in areas such as credit score reporting and employment. To understand how algorithms may be biased, we first need to understand something of how they are created.

Computer programmers create a series of decisions that an algorithm uses to come to a certain conclusion. If the programmer is designing an algorithm that will decide whether to sell

---

[17] *See supra* note 3.

[18] JULIA ANGWIN ET AL., PROPUBLICA, MACHINE BIAS (2016), https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing.

[19] *See* Jouvenal, *supra* note 1.

[20] *Supra* note 5, *see also* Ric Simmons, *Quantifying Criminal Procedure: How to Unlock the Potential of Big Data in Our Criminal Justice System*, 2016 MICH. ST. L. REV. 947 (2016) (these authors have done important work in identifying and outlining the problem, while specific solutions are still forthcoming).

[21] *See* Danielle Keats Citron & Frank Pasquale, *The Scored Society: Due Process for Automated Predictions*, 89 WASH. L. REV. 1 (2014).

lemonade on a certain day, the decisions might look something like: If it's sunny outside, then sell, if it's rainy outside, then don't, and if it's cloudy outside, look at the temperature. The algorithm's conclusion is a sum of the decisions that the programmer makes; after all, she is the one who knows that when it's sunny outside it will be better for business. This is perhaps the simplest way of approaching an algorithm. Thus, the first place to worry about bias entering the system is at the very beginning, with the programmer herself.

However, it isn't likely that the programmer will write bias into the code. Although it would be convenient to find the line of code that says, "if suspect is black, then arrest," this is an obvious issue that companies and police departments would want to prevent. In fact, programmers may not even include race in the code. For example, leading predictive policing company PredPol claims that it only inputs three variables into its algorithm: type of crime, crime location, and time of the crime.[22]

Yet the fact that a company does not code bias into an algorithm or include racial factors does not mean that the algorithm is free of bias. It is possible for an algorithm to use factors that instead act as proxies for bias.[23] For example, poor socioeconomic conditions in many neighborhoods are correlated with race.[24] In the context of criminal data, race could also be correlated with gang membership and community affiliations. Thus, an algorithm may take racial factors into account even when the programmers did not originally intend it.

The problem of discovering bias is further compounded by the fact that not even the programmers themselves are entirely positive of what is going on under the hood of their algorithms, especially when they incorporate machine learning. Machine learning techniques create algorithms that are difficult to understand.[25] For example, a predictive policing algorithm that uses machine learning may not be trained to "solve" the problem of "When will crime happen?" Instead, the machine learning algorithm might be presented with a series of hypothetical scenarios – say time of day, season, and nearby landmarks. It would also be given the thousands and thousands of data points past and present about various factors – the variables already mentioned, perhaps past crimes, and others. Then the algorithm connects these data points, at first almost at random, trying to predict whether a crime will occur in the hypothetical scenario. If the algorithm does not predict the outcome correctly it goes back to the drawing board and starts making those connections again. If it predicts the outcome more precisely, the connections that helped it make that decision are incorporated into the code while the rest of the algorithm goes back to the drawing board again.

---

[22] PREDPOL, *Predictive Policing Technology*, https://www.predpol.com/technology/.

[23] BIG DATA WORKING GROUP, BIG DATA: A REPORT ON ALGORITHMIC SYSTEMS, OPPORTUNITY, AND CIVIL RIGHTS 7-8 (May 2016), https://obamawhitehouse.archives.gov/sites/default/files/microsites/ostp/2016_0504_data_discrimination.pdf.

[24] RUTH D. PETERSON & LAUREN JOY KRIVO, DIVERGENT SOCIAL WORLDS: NEIGHBORHOOD CRIME AND THE RACIAL-SPATIAL DIVIDE 53-57 (2010) (about one-third of residents in the average African-American neighborhood lived in poverty in 2010, compared with less than a tenth of residents in the average white neighborhood).

[25] Danny Sullivan, *How Machine Learning Works, As Explained by Google*, MARTECH TODAY (Nov. 4, 2015, 1:12 PM), https://martechtoday.com/how-machine-learning-works-150366 (for a non-technical overview of machine learning).

This process is repeated thousands upon thousands of times until the machine learning algorithm begins to successfully predict what happened in these hypothetical scenarios. This is somewhat akin to the idea of "brute force" discovering a pin number or password. If "1111" doesn't work, then maybe "1112" will, and if that does not work then it's time to try "1113," and so on. Computers can process these calculations much faster than a human brain, using datasets that contain millions of data points. However, the result is that the original programmers – and even the algorithms themselves – aren't necessarily aware of *why* the algorithm decides what it does, only that it does. These "black box" algorithmic systems are difficult to completely unravel.[26]

This is where the potential for algorithmic bias is the most dangerous, in large part because it is difficult for even technical experts to understand what is happening. The algorithm comes to its conclusions mostly through processing the data. This means that if the data is biased the decisions the algorithm comes to will be biased, without anyone being the wiser. As one researcher put it, "bias in, bias out."[27]

Predictive policing algorithms also have a greater potential to be biased than other machine learning algorithms for a variety of reasons. First, criminal justice data is often incomplete. Police department records may not be extensive, especially in smaller cities, which may not have the resources to collect and store robust data.[28] Even in big cities, historical data may not extend very far, and much of the data will have been collected in recent years.[29] As a report commissioned by the White House put it, "criminal justice data is notoriously poor."[30] Because these algorithms often require massive datasets, using incomplete data and small sample sizes will certainly leave holes in the final analysis.

Additionally, criminal justice data often reflects an implicit bias. The fact that police interact with and even arrest racial minorities more often than members of other groups means that these actions will be recorded as data. Similarly, if police focus their resources on patrolling high-crime areas, those areas will show up in the dataset more often and in more detail. There is the possibility that this creates a feedback loop where the police patrol the same areas, because more crime happens there, which leads to more data being collected, which leads to more police being sent to the same area.[31]

These problems are aggravated because of the credibility lent to the idea of "data." Companies and police departments are all too willing to accept the data and the conclusions

---

[26] Joshua A. Kroll et al., *Accountable Algorithms*, 165 U. PA. L. REV. 633, 651 (2017).
[27] Sandra G. Mayson, *Bias in, Bias Out*, 128 YALE L. J. (forthcoming 2019).
[28] *See Policing, supra* note 5, at 1148 ("National crime statistics exist, but they cannot provide a relevant database necessary to predict local crime patterns because the information is not localized. The result is that the existing data may be of limited value for predictive validity in the vast majority of jurisdictions and only useful in large urban cities with significant crime data collection capabilities.").
[29] Daniel W. Rasmus, *Why Big Data Won't Make You Smart, Rich, or Pretty*, FAST COMPANY (Jan 27, 2012), https://www.fastcompany.com/1811441/why-big-data-wont-make-you-smart-rich-or-pretty ("We must remember that all data is historical…Every model is based on historical assumptions and perceptual biases.").
[30] *See* BIG DATA WORKING GROUP, *supra* note 23 at 21.
[31] Kroll et al., *supra* note 26, at 680-681 (using current stop and frisk policies and racial profiling as an example).

reached by the algorithm as "the truth" because data is so often thought of as objective truth.[32] The officer on duty may not hesitate to question how the data was gathered or why the algorithm is programmed to suggest to suspect person x, and rightfully so, as that isn't necessarily her job. More poor data is fed into the system, leading to a reinforcement of biased decision making, yet these interactions are recorded as facts rather than the decisions of implicit or explicit police bias.

Policy makers have become aware of these problems and in many states have enacted laws requiring audits of police department data. However, these laws are weak in practice and have largely been ignored by police departments.[33] Additionally, these laws were not designed with predictive policing or big data machine learning algorithms in mind.

In summary, the potential for bias in predictive policing systems comes not only from how the code is written but from the data that is fed into it. Poor data, or even good data poorly chosen, can lead algorithms to learn bias even though this is not what the algorithms' creators intended. I propose audits of the algorithms that police departments use as a check against these potential biases.

III. Predictive Policing Audits

Predictive policing has been adopted without much resistance or a careful examination of the potential consequences. Companies and police departments have held onto their secrets while activists have failed to get the attention of lawmakers.[34] A more transparent process would increase trust in predictive policing methods and enable concerned citizens and policymakers to examine its effects. I propose that Congress and state legislatures should pass legislation that require external audits of the predictive policing algorithms that police departments use. These audits would examine algorithms for disparate impact by using the most up-to-date tools developed by big data and machine learning experts.

A. Audits Provide Transparency and Accountability

---

[32] Kate Crawford, *The Hidden Biases in Big Data*, HARVARD BUSINESS REVIEW (Apr. 1, 2013), https://hbr.org/2013/04/the-hidden-biases-in-big-data ("The hype becomes problematic when it leads to what I call 'data fundamentalism,' the notion that correlation always indicates causation, and that massive data sets and predictive analytics always reflect objective truth.").

[33] Wayne A. Logan & Andrew Guthrie Ferguson, *Policing Criminal Justice Data*, 101 MINN. L. REV. 541, (2016).

[34] Robert Brauneis & Ellen P. Goodman, *Algorithmic Transparency for the Smart City*, 20 YALE J. L. & TECH 103 (2018) (authors requested predictive policing records from police departments but were largely unsuccessful – only three of eleven departments provided documents about PredPol, for example); *but see* STOP LAPD SPYING COALITION, BEFORE THE BULLET HITS THE BODY – DISMANTLING PREDICTIVE POLICING IN LOS ANGELES (May 8, 2018), https://stoplapdspying.org/wp-content/uploads/2018/05/Before-the-Bullet-Hits-the-Body-May-8-2018.pdf (activists successfully requested documents on PredPol); *and Brennan Ctr. For Justice at New York Univ. Sch. Of Law v. New York City Police Dept.*, No. 160541/2016, 2017 N.Y. Misc. LEXIS 5138, (N. Y. Sup. Ct. Dec. 22, 2017) (Brennan Center successfully requested predictive policing documents under New York's freedom of information laws).

Transparency and accountability are essential checks against institutions such as the police. In this context, transparency means sharing information on government processes as broadly as possible without compromising citizens' safety or peaceful private interests. Regular audits by technical experts will provide the transparency and accountability that currently does not exist for predictive policing.

i. Transparency and Accountability are Essential for Democracy and the Police

Democratic societies need transparent institutions to flourish. Healthy democracy needs an active and engaged citizenry that understands the issues that it faces, which enables citizens to make informed decisions about governance.[35] This is particularly vital to preserve the rights of minorities and marginalized groups that may not wield as much political influence. Transparency provides essential knowledge of government actions so that the citizenry can make these evaluations. Without the knowledge provided by transparency, citizens will not be able to have a say because they will not even know what questions to ask about how the government behaves. Certainly, the ideal of a fully engaged citizenry is not realized in real life, but there are still thinkers, activists, and policymakers who are engaged. Transparency can, at the very least, provide these key groups with the information they need to find and debate appropriate solutions.

Transparency and accountability also increase trust in institutions. It is no secret that public trust in the police has become a hot topic in recent years.[36] When the public and the police do not trust each other, relations and cooperation break down. The public becomes more fearful, which makes it harder for the police to do their jobs.

Secretive predictive policing methods will not improve this relationship. Transparency would improve public trust in predictive policing, which would then improve the public's relationship with the police. If the public and police have a mutual understanding and respect, they will be able to cooperate to solve the problems that face communities.

ii. Audits Provide Transparency and Accountability by Allowing Experts to Examine the Algorithms

Transparency and accountability are particularly important for topics which are as obscure and technical as predictive policing algorithms can be. The effects and potential biases of predictive algorithms are likely to be understood fully by only a small number of experts in the field. If transparency does not exist for these algorithms, there is little chance that anyone will have the information needed to decide whether they are harmful or how they should be used and regulated. In fact, as noted above, many people still aren't even aware that predictive policing techniques exist.

If enough transparency for algorithms exists, experts who are interested in finding solutions to the problem it will be able to do so. Many experts are already very concerned about the potential problems predictive algorithms could pose.[37] However, at this point many of their

---

[35] Tal. Z Zarsky, *Transparent Predictions*, 2013 U. ILL. L. REV. 1503, 1530-1531 (2013).
[36] Jim Norman, *Confidence in Police Back at Historical Average*, GALLUP (July 10, 2017), https://news.gallup.com/poll/213869/confidence-police-back-historical-average.aspx (noting that confidence in the police hit a historic low in 2015).
[37] *See* CATHY O'NEIL, WEAPONS OF MATH DESTRUCTION (2016); *and supra* notes 17-19.

opinions are still limited to conjecture and generalization because of a lack of transparency. Once experts can analyze the algorithms, they will then be able to educate the public and inform a broader debate. Predictive policing may be a benefit to society, or it may be harmful, but there is not enough information yet to have much of a debate at all.

Requiring regular audits of predictive policing algorithms used by police departments would go a long way towards establishing this sorely needed transparency. Audits could not only catch bias red-handed but motivate a measure of preventative caution.

The most obvious benefit of an audit would be to catch algorithmic bias in the act. Again, these biases may creep in completely unexpected, even out of data that would, on its face, seem to be completely unbiased. The audit would act as a kind of peer review, checking for problems that not even companies and departments had thought of.

Another, perhaps more important effect of auditing would be to motivate companies and police departments to be on their toes and actively search for these biases. If a police department knows that it will be audited, it will most likely want to perform well on that audit. This provides an incentive to be on the lookout for problems in the algorithm. This can lead to more fundamental changes than if an audit simply finds the algorithm to be biased because companies and police departments would be motivated to look at practices that have so far escaped scrutiny. For example, police departments would need to examine their practices from beginning to end, from the data collection process to the implementation by officers on the job. In doing so, departments and companies may find solutions that would not be prescribed by the audit. The threat of audits would therefore have the possibility of preventing bias by motivating predictive policing users to examine the problem and find solutions. As the old adage goes, "An ounce of prevention is worth a pound of cure."

Requiring frequent audits is also important because algorithms are in a state of constant change. New data is added into the system every day and analyzed for patterns.[38] It is possible that an algorithm that does not exhibit bias today may, with addition of new datasets and information, exhibit bias in six months. Legislatures should consult with experts to establish audit requirements at intervals that are frequent enough to adapt to new problems but not so frequent that they strain department resources. A small police department in a rural town that uses location-based predictive policing in a limited context will not experience the same amount of flux and change in its algorithm as a police department in a metropolitan area that is flooded with new data daily. In this example, it may be more appropriate for the small police department to be required to contract an audit only once every two years, while the department in the big city should be audited every six months.

Of course, these fields are still developing, and an audit may not even find that there is a clear answer to whether a predictive policing algorithm is biased. This is not a problem because the goals of the audits—transparency, accountability, and catching offenders—are only one part of the larger objective to prevent bias. Additionally, part of the broader goal is to facilitate discussion, and that discussion will best be fueled with better information. The audit process will necessarily require more dialogue between departments, companies, auditors, those they report to, and the general public. Once auditors can reveal just how predictive policing is implemented,

---

[38] Rob Kitchin, *Thinking critically about and researching algorithms*, 20 INFO. COMM. AND SOC'Y 14, 18 (2017) ("[algorithms] are ontogenetic in nature (always in a state of becoming), teased into being: edited, revised, deleted and restarted…").

these communities can discuss its positive and negative impacts and how best to create rules for its usage.

iii. Objections to Audits

Transparency skeptics may point out that establishing an auditing system for predictive policing in the name of transparency may set a precedent that would reach far beyond law enforcement. This precedent may spur legislation that seeks to examine the inner workings of all sorts of algorithms, not just for bias but for any number of reasons. Companies such as Google and Facebook are understandably protective of the closely-kept secrets of the algorithms that have made them very wealthy. Trade secrets will be addressed more fully later in the article, but it is important to address the argument that this precedent could eventually lead to a sort of algorithmic witch hunt.

In response, it is important to ask whether this would be a bad thing. Algorithms govern more and more of our daily lives, from giving us suggestions on where to eat, to determining who might employ us.[39] The extent to which we interact with algorithms will only increase from this point forward. We may be long overdue for a systematic process to ensure the quality of these algorithms and examine their impact.

More importantly, the recommendation of this article is specific to law enforcement due to the important social, government, and constitutional issues involved. The police have become the focal point of much criticism lately for a variety of reasons, but a common theme among them is the desire to prevent the abuse of power by government offices that we have entrusted with our own protection. The possibility for predictive policing algorithms to cause harm is proportional to the great responsibilities and powers that the citizens and governments have given to the police.

The flip side of these responsibilities is that the police are, at least in theory, ultimately answerable to the public through local government participation. These methods of addressing the problem are fundamentally a public governmental process. Thus, audits are a public governmental solution. Likewise, transparency is a necessity for good governance and for democracy.[40] The methods for addressing the problems presented by Google's algorithms would be different and would not be rooted in this process, so the precedent set in this case would not immediately extend to private companies. Auditing predictive policing algorithms is a solution for the problem of predictive policing, which is why I only recommend audits for police departments. This article does not necessarily recommend audits in other areas.

Skeptics might also object to the costs that audits would impose on police departments, which would also be costs for taxpayers. Managing audits could also put extra strain on the federal government or state governments that oversee some of the process.

The benefits already listed far outweigh the potential costs of creating or reinforcing discriminatory practices among the police. Additionally, audits and the frequency of audits would necessarily be tailored to the circumstances of the department. Small departments with fewer resources to spend on an audit will also be spending fewer resources on these kinds of programs and data collection in the first place, so the expense of the audit would be proportional

---

[39] *Id.* at 15-16.

[40] *See supra* Section III.A.1.

to the expense of the programs. Bigger departments with more data and more expensive techniques will likewise be expected to require more expensive, more frequent audits.

It is worth questioning whether audits are the best of any available method for transparency and accountability. First, it is unreasonable to expect companies or police departments to release the code for public review. Not only would companies lose trade secrets they have worked hard to develop, but the public in general would most likely not understand the code. Auditing provides a middle ground for experts to analyze the data then release a report to the public while protecting information that should not be released.

Second, although governments could require companies to design algorithms with checks against bias, we should hesitate to give governmental bodies such a mandate. No matter how well-intentioned policy makers might be, such technical solutions are best left to experts who are reviewing the technology first-hand. It is important to give companies and departments the flexibility to find appropriate solutions to these problems. Although I am comfortable with the possibility of government stepping in to audit more algorithms, I am not comfortable with precedent that would allow government to change algorithms at a whim and dictate more broadly what they should or should not do.

Finally, although there may be problems with predictive policing algorithms, there is no reason to simply ban their use. Predictive policing has the potential to have a positive impact in enhancing law enforcement performance and efficiency and may even be used to correct bias among officers.[41]

Before we can see the benefits of predictive policing, it needs to be transparent and accountable. If experts can assess algorithms through audits, the public can be more reasonably assured that they are not biased.

B. The Auditing Process

This article does not attempt to outline every facet of the auditing process. That will best be left to technical professionals and legislatures to decide. However, there are some key points specific to predictive policing algorithms that the auditing process should include. The audit should test algorithms as used by police departments for disparate impact against a variety of different groups.

i. External Audits Examine Algorithms Used by Police Departments

Police department audits, such as financial or performance audits, are usually handled by audit offices within the departments themselves. Depending on the topics being audited, audits may be handled by other agencies with the proper oversight, such as a city government, a state attorney general, or even the Department of Justice.[42] These audits are not always published or made accessible to the public. This article encourages external audits rather than internal audits to achieve the goal of transparency.

If an audit is performed internally by a police department, it is reasonable to expect some public skepticism about its validity. Indeed, police departments have been known to misrepresent

---

[41] BIG DATA WORKING GROUP, *supra* note 23, at 19-20.

[42] Allan Y. Jiao, Police Auditing: Standards and Applications 49 (2d ed. 2015), EBSCOhost [hereinafter Jiao].

statistics to dishonestly diminish negative practices or enhance positive outcomes.[43] Of course, this problem is not unique to police departments; some suspicion would certainly be reasonable towards any internal audit. I also do not mean to ignore the efforts of internal department auditing offices established in good faith and operating exactly as expected. Nevertheless, this does not change the fact that audit reports published by an internal office, no matter how well done or honest, may fall under suspicion.

External audits appear to provide a more reliable, objective alternative. Audits have been performed at times by local and state governments. Some cities have created an Office of the Independent Police Auditor charged with the specific task of overseeing the city's police department.[44] These auditors appear to be more credible because they do not answer to the departments themselves. Audits could also be performed by external firms and contractors. At least one enterprising young company already offers to audit algorithms for potential discrimination, and its certification is seen as something akin to an "organically grown" sticker for algorithms, showing that the algorithm is free from unrecognized biases or other harmful effects.[45]

ii. The Audit Analyzes Protected Groups

The main purpose of the audit is to analyze predictive policing algorithms for bias against protected groups as defined by federal law. The ultimate consequences of an audit would be to give the Department of Justice the necessary facts to bring a pattern and practice lawsuit against the department or company. A sample of protected groups that may be discriminated against include those that have been set out in federal law: race, religion, sex, national origin,[46] age,[47] or disability.[48] Other groups not specifically covered in federal law would LGBTQ groups.

The audit would focus on discrimination against these groups because there are usually legal mechanisms that can protect them. However, this does not mean that legislatures should prohibit auditors from examining additional groups of concern. An audit that seeks to discover potential racial biases is likely to uncover other problems in the process – for example, patrols that are assigned more heavily than necessary towards poorer parts of cities, which may also correlate with race. Auditors should have the freedom to examine other problems and possible solutions.

iii. The Standard of Disparate Impact

Determining whether there is bias against these groups is a difficult question to answer. The federal law mentioned above typically defines discrimination using the standard of

---

[43] Jeff Morganteen, *What the CompStat audit reveals about the NYPD*, N. Y. WORLD (July 3, 2013), http://thenewyorkworld.org/2013/07/03/compstat/.
[44] Jiao, *supra* note 42, at 135-147.
[45] Jessi Hempel, *Want to Prove Your Business is Fair? Audit Your Algorithm*, WIRED (May 9, 2018, 8:00 AM), https://www.wired.com/story/want-to-prove-your-business-is-fair-audit-your-algorithm/.
[46] Civil Rights Act of 1964, 42 U.S.C. § 2000e (2017).
[47] Age Discrimination in Employment Act of 1967, 29 U.S.C. § 621 (2017).
[48] Americans with Disabilities Act of 1990, 42 U.S.C. §§ 12112, 12132 (2017).

"disparate impact." Disparate impact was first established in the context of employment in the Civil Rights Act of 1964,[49] and its use has been upheld by the Supreme Court in other contexts, such as housing.[50] An action that causes disparate impact leads to unequal, harmful treatment of members of a protected group. This impact does not have to be crafted intentionally.[51] Audits should use the standard of disparate impact to evaluate whether an algorithm is biased.

The process of assessing disparate impact of algorithms has been described as: "[One,] Statistical evidence demonstrating a disproportionate adverse impact caused by a policy or procedure; [Two], Assessing whether the policy or procedure serves a valid purpose and the extent thereof; and [Three,] Assessing whether there are alternative policies or procedures that would achieve the legitimate objective with less of a disparate impact."[52] Following this process would help satisfy statutory and judicial definitions of disparate impact.

The applications to a predictive policing audit follow naturally. First, the audit would have to show that the algorithm causes statistically observable biases of police conduct against a given group. For example, the algorithm might suggest that a city's Latino sector should be patrolled in a statistically disproportionate manner. This could cause harm by increasing the number of arrests among this population in a similarly statistically disproportionate manner.

Second, the audit would examine whether the algorithm is accomplishing what it needs to. Is it predicting crime accurately? Are the streets safer? In the hypothetical case above, the algorithm may be hurting more than it is helping.

Finally, the audit would examine whether there are better alternatives to predictive policing. The department could return to using more traditional methods. Or, the audit could suggest that the department improve its data collection methods, change how the officers use the algorithms, or even suggest how to tweak the code itself to correct biased predictions.

The audit's purpose would not be to punish police departments or demand binding changes. The goal is transparency, and auditors would not be given powers beyond that of making recommendations, functioning more as a peer review of the technology. This is part of the reason why auditors should be given the authority to examine a wide range of topics and potential biases.

Of course, proper government officials can take action when necessary. For example, if an audit that reveals that an algorithm exhibits clear disparate impact and this has led to a negative impact on a city's protected groups, it could be appropriate for the Department of Justice to bring a pattern or practice lawsuit against the police department.[53]

---

[49] Civil Rights Act of 1964, 42 U.S.C. § 2000e-2(k)(1)(A) (2017).

[50] *Tex. Dep't of Hous. & Cmty. Affairs v. Inclusive Cmtys. Project, Inc.*, 135 S. Ct. 2507 (2015), *see also* Mark MacCarthy, *Standards of Fairness for Disparate Impact Assessment of Big Data* Algorithms, 48 CUMB. L. REV. 67, 80-82 (2017).

[51] CMMTY NAT'L STAT., MEASURING RACIAL DISCRIMINATION 40 (Rebecca M. Black et al. eds., 2004).

[52] Mark MacCarthy, *supra* note 50, at 83 (paraphrasing PUB. POL'Y DIVISION, SOFTWARE AND INFO. INDUSTRY ASS'N, ALGORITHMIC FAIRNESS (Sep. 22, 2016), http://www.siia.net/Portals/0/pdf/Policy/Algorithmic%20Fairness%20Issue%20Brief.pdf).

[53] U.S. DEP'T OF JUSTICE, A PATTERN OR PRACTICE OF DISCRIMINATION, https://www.justice.gov/crt/pattern-or-practice-discrimination (updated Aug. 6, 2015).

iv. Auditing Tools and Methods

The audit will need to test for statistical evidence of disparate impact by using a wide variety of techniques and procedures developed by the technological community. Predictive policing algorithms, like other algorithms, come in many different flavors. Even when considering a single type of algorithm, there is no consensus, despite plenty of debate, on the best method to test it for bias. Thus, the auditors should be prepared to use several different methods before being satisfied that they have examined the algorithm thoroughly.

The first step, naturally, is to look at the code of the algorithm. This may give auditors a glimpse of the full picture, but it is very different from running the algorithm. Especially with machine learning, an algorithm needs to be fed data before auditors can understand how it works.[54]

This is why auditors should also use "black-box" styles of testing. Black-box methods observe both the inputs and outputs to an algorithm. By tweaking the different inputs and observing how the algorithm processes them, auditors can learn a lot about how the code works.[55] Other related techniques include running dummy data through the algorithm, using "scraping" and "sock-puppet" algorithms, running experiments with the algorithm using real world users ("crowdsourcing" audits), and others.[56]

This list is meant to establish a minimum standard for the breadth of tests that should be required to satisfy auditors and, consequently, the public. It is not an exhaustive list. Each method has its own strengths and drawbacks. Experts are aware of these differences and vigorously discuss them. They will continue to develop best practices, and I certainly hope that the discussion will continue.

Additionally, audits should look at the entire process from beginning to end. Because the inner workings of algorithms are often inscrutable, it is important to examine how an algorithm is used within its full context. It might not even be appropriate in many cases for the audit to make recommendations on changes to the code itself, but on the processes leading up to and occurring after the algorithm's decision.

In particular, the audits should take special care to examine the data and the data collection process because this is a possible cause of bias. Additionally, it would certainly be appropriate to revive current and past legislation on database audits that have fallen into disrepair.[57] This would provide a valuable complement to audits of predictive policing.

Predictive policing algorithms should also be required to keep an audit log, a timestamped record of actions taken by the algorithm. This audit log could be required as part of legislation on obligatory audits or in another piece of legislation. Audit logs are already required

---

[54] *Supra* note 38.

[55] *Supra* note 26, at 660-661.

[56] CHRISTIAN SANDVIG ET AL., AUDITING ALGORITHMS: RESEARCH METHODS FOR DETECTING DISCRIMINATION ON INTERNET PLATFORMS (2014) http://www-personal.umich.edu/%7Ecsandvig/research/Auditing%20Algorithms%20--%20Sandvig%20--%20ICA%202014%20Data%20and%20Discrimination%20Preconference.pdf, Deven R. Desai & Joshua A. Kroll, *Trust but Verify: A Guide to Algorithms and the Law*, 31 HARV. J. LAW & TEC 1, (2017).

[57] *See supra* note 33.

in other technology sectors and would be invaluable to auditors who are attempting to verify an algorithm's processes.[58]

Auditing with a variety of different methods can catch disparate impact in predictive policing algorithms and inspire change. However, change may not come if only the police and the auditors see the results. Thus, legislatures should mandate a final step in the auditing process: an audit report.

IV. Audit Reports

The audit report is an essential part of the auditing process. None of the benefits of transparency will be available if the public does not have access to the audit's findings. Police departments will naturally be worried about revealing sensitive information to the public that could be used to compromise law enforcement purposes, and companies will not want to reveal trade secrets. These are valid concerns. This article does not suggest that everything that an audit discovers should be available to the public. However, audit reports should not be kept secret; every audit report should be released to the public. The question is then: What should be included in the audit report?

Thankfully, there is precedent in this area that can guide legislatures to find a reasonable compromise between transparency and the interests of police departments and software companies. I suggest that the basic inputs and outputs of an algorithm should be made available in all public audit reports, as well as details on how the algorithms are implemented in daily procedure and assessments on the algorithms' potential bias. These suggestions are based primarily on provisions provided by the Freedom of Information Act.

A. The Freedom of Information Act: Basics

The Freedom of Information Act (FOIA),[59] created in 1967, is designed to "[keep] citizens in the know about their government" by making federal agency records public.[60] It has been amended several times and subjected to various executive orders, both to provide more transparency and to create more exemptions for agencies wishing to keep their records out of the public eye.[61] Nonetheless, it has remained the most important legislation of its type in the United States and provides a certain standard for transparency that we can apply to these audit reports.

FOIA is a federal law, but most state freedom of information laws are patterned after it.[62] In this section we will focus on the federal FOIA because there is not enough space to explain all the intricacies of the various state laws. The discussion here can be taken as a broad template that can be applied to more individual circumstances. It is worth noting that at least one lawsuit in

---

[58] *See* Devai and Kroll, *supra* note 56 at 40-41.

[59] Freedom of Information Act, 5 U.S.C. § 552 (2012).

[60] Off. of Info. Pol'y, U.S. Dep't of Justice, *What is FOIA?*, https://www.foia.gov/about.html (last visited Jan. 14, 2019).

[61] U.S. DEP'T OF JUSTICE, *Introduction* to DEPARTMENT OF JUSTICE GUIDE TO THE FREEDOM OF INFORMATION ACT 5-10 (2013), https://www.justice.gov/oip/doj-guide-freedom-information-act-0.

[62] Justin Cox, *Public Interest Practice Session: Maximizing Information's Freedom: The Nuts, Bolts and Levers of FOIA*, 13 N.Y. CITY L. REV. 387, 412-413 (2010).

New York successfully secured the release of predictive policing documents under New York's freedom of information law.[63]

B. What the Report Includes

FOIA allows the public to request documents from government agencies. This original mandate is broad and sweeping; under it, citizens are entitled to request whatever document they desire. However, the statute also includes nine exemptions that limit this mandate. Exemption 7 exempts the records of law enforcement agencies from being requested under certain circumstances.[64]

I will now discuss the bare minimum that a publicly-released audit report should include, using these exemptions as a guide. I will also discuss other statutes when relevant. However, it is valuable to remember that these are exceptions to the rule, rather than the rule itself, which requires broadly that agencies give up their records.

i. Data and Inputs

As has been discussed extensively above, the data that enters the algorithm is the first point of potential bias.[65] Thus, any public discussion of possible bias and possible remedies to bias in predictive policing models needs to begin at an understanding of what data police departments are recording and feeding into the model. Knowledge of this data could also foster a discussion that continues beyond the realm of predictive policing, as communities attempt to examine root causes of the kinds of behavior that departments are ultimately attempting to correct through law enforcement. However, this is also the part of the report that will likely need to be the least specific and the most heavily redacted so that police departments can continue to fight crime effectively.

First, personal information cannot be released to the public. This is covered in exemption 7(c) of FOIA,[66] as well as the Privacy Act of 1974,[67] which prohibits federal agencies from releasing any personal information. In the context of predictive policing data, this would include names, addresses, phone numbers, and any other personal information. This also applies to criminal records. However, this does not prevent a report from detailing the types of criminal records that are used by the algorithm. In that case, personal information would not be revealed.

Second, exemption 7(e) also applies to records that "would disclose techniques and procedures for law enforcement investigations or prosecutions, or would disclose guidelines for law enforcement investigations or prosecutions if such disclosure could reasonably be expected to risk circumvention of the law."[68] The language of the statute and the subsequent interpretation by courts have been rather broad in deciding what could "reasonably be expected" to

---

[63] *See Brennan Ctr. For Justice at New York Univ. Sch. Of Law v. New York City Police Dept.*, *supra* note 34.
[64] Freedom of Information Act, 5 U.S.C. § 552(b)(7) (2012).
[65] *Supra* Part II.
[66] *Supra* note 59.
[67] Privacy Act of 1964, 5 U.S.C. § 552a (2012).
[68] *Supra* note 59.

compromise law enforcement activities,[69] including a wide range of activities, both known and unknown to the public.[70]

Using this exemption as a guideline, the public audit report should exempt any data that is gathered using methods that, if they were known, would "risk circumvention of the law."[71] Using examples that have appeared in court rulings, this would mean that a report would not include data gathered through undercover surveillance techniques or operations. Federal databases have also been protected by FOIA, and so should not be disclosed in a public report.[72] Generally, it could be said that the more sensitive the data, the more likely it is that the exemption would apply.

It may appear that in the end the report would not disclose much of the data the algorithm uses, but the truth is that much of this data does not fall under these exemptions. Each predictive policing algorithm uses different kinds of data, but here are some ideas as to the data that should be released in the report: socioeconomic factors, "hot spot" locations, including neighborhoods, weather, type of crime, types of historical criminal records, and time of the crime. These types of data are not controversial, do not reveal personal information, and in many cases are already public. Police department objectives will certainly not be compromised if the public knows that the algorithm considers whether the day of the crime is cloudy or sunny. Of course, the public might not care either, as weather probably has nothing to do with bias. Nonetheless, the point is that transparency means that police departments should release as much data to the public as possible rather than the bare minimum.

Consider how the public could react if it knows that the department feeds into the algorithm a list of locations in the city that are considered to be high-risk, such as banks, schools, and restaurants.[73] The public could examine whether or not these locations are associated or correlated with problematic biases – for instance, particular restaurants in a low-income area of the city. Then the public could demand that police departments ensure that these factors are not weighted disproportionately. The public could also shine a light on how to solve these issues through community action. Knowledge of the data that the algorithms use is the first, vital step to creating transparency and accountability for predictive policing.

ii. Outputs and How Algorithms Are Employed

Transparency begins with data, but it must also include the algorithms' outputs and recommendations. Therefore, it is vital for an audit report to release these outputs to the public, as well as an assessment on their potential problems and details on how the results are used in the day-to-day of law enforcement. A report that did not include algorithmic outputs would not allow for transparency because, in the end, these outputs are what police departments act on.

---

[69] Including in some rulings that such an expectation not even be necessary.

[70] U.S. DEP'T OF JUSTICE, *Exemption 7(E)* in DEPARTMENT OF JUSTICE GUIDE TO THE FREEDOM OF INFORMATION ACT 1-15 (2013), https://www.justice.gov/oip/doj-guide-freedom-information-act-0 [hereinafter DOJ GUIDE].

[71] *Supra* note 59.

[72] *Supra* note 64.

[73] AZAVEA, HUNCHLAB: UNDER THE HOOD 16 (2015), https://cdn.azavea.com/pdfs/hunchlab/HunchLab-Under-the-Hood.pdf (a report released by HunchLab gives examples of the kinds of locations fed into the predictive policing algorithm).

For location-based predictive policing models, the report should reveal the specific areas that the algorithm has suggested require special attention. If the public knows what neighborhoods or city blocks are receiving attention, it will be able to better understand how the algorithm makes decisions and whether those decisions are biased. For instance, if it is revealed that an algorithm recommends that most police patrols be allocated to the Latino sectors of a city, the citizens could then have a discussion as to whether there is any bias involved.

Police departments might worry that revealing this information could endanger officers working in those areas or "risk circumvention of the law." However, criminals likely already know where patrols are allocated, so this information would not likely do them any good and would not endanger officers any more than they already expect.[74]

However, the report should not reveal algorithmic outputs that identify individuals. Not only would this violate individual privacy, it would certainly run a greater risk of endangering officers and "risk circumvention of the law." If an individual knows that an algorithm has added her to a list of potential offenders, and if that individual breaks the law, she will most likely take extra precautions against law enforcement catching her.

Instead of revealing outputs that identify individuals, the report should use hypothetical data to illustrate how the algorithm works. Creating such hypotheticals would almost certainly already be part of the audit process. Take an example where an algorithm attempts to predict who is more likely to be involved in gang violence. The audit report could include the algorithm's analysis of two hypothetical individuals that are identical except for their ethnicity. If one of the individuals is Hispanic, does the algorithm rate his threat level differently? The public could then evaluate whether the algorithm's decision is justified given the data it uses. This process would preserve privacy while also giving the public an honest look at predictive policing analyses.

The report must also include details on how the algorithm is used in the day-to-day of law enforcement. Is it used by command staff to make top-down decisions or is it used by patrol officers? How often and why do officers act on the algorithm's suggestions while on patrol? How are they trained to use the algorithms? This information is the vital link between the theory of predictive policing and how it impacts the lives of citizens.

iii. The Final Assessment and Recommendations

Reporting on the outputs is closely linked to an assessment of the outputs. The report must include an assessment of whether and how the predictive policing instrument is biased. This is the heart of the report, the part that the public is most interested in hearing. Discussions of data and outputs are important, but the technical layman will be most interested in whether experts give the program a thumbs up or a thumbs down.

Often, the nature of such complex systems will not be easy to explain or be fully analyzed. The experts may not come to a solid conclusion of "yes, it is biased" or "no, it is not." Thus, the report must be honest and balanced. It must include every pertinent detail and conclusion possible while still adhering to the guidelines already established. The report need not limit itself to examining disparate impact; if other issues are discovered, they must be addressed. The report must also include suggestions for possible solutions or alternatives for any problem

---

[74] *See Brennan Ctr. For Justice at New York Univ. Sch. Of Law v. New York City Police Dept.*, *supra* note 34 (the Brennan Center successfully made this argument in this case).

that the audit uncovers. This assessment would jump-start public discussion and likely prompt earnest soul-searching for police departments.

The report's recommendations make up the primary call to action that police departments should act upon. Recommendations would be tailored to the needs of the department and the circumstances of its jurisdiction. Recommendations could include methods to clean and update databases, selections of data that should or should not be used, changes in department policies that break feedback loops, changes in how officers on patrol use algorithmic predictions, reallocation of resources away from specific neighborhoods, or others. Although these recommendations would not be binding, pressure by the public, interest groups, and government officials would motivate departments to enact them.[75] If that fails, the public could demand that local government oblige departments to change their practices.

Additionally, it is important to emphasize here that the focus is on algorithms as used by departments and not as created by companies. Audits can certainly find problems and make recommendations on algorithms straight out of the box. However, the problems really begin when algorithms are given police department data to chew on. This is why the audit would primarily give recommendations to police departments. Companies would receive notification of these recommendations to understand how to make improvements. When it is possible to distinguish the effects of the police department's data from the code created by the company, then the report would make recommendations directly to the company.

iv. Trade Secrets

The companies that have created this predictive policing software may fear that this report would reveal valuable information about how the algorithm is created that could then be copied by competitors. However, there is no reason to fear that the report will reveal any such trade secrets. The report will only reveal algorithmic inputs and outputs. It does not reveal the inner workings or code of the algorithm. Neither is it practical for any enterprising citizens to reverse-engineer the algorithm from a knowledge of these statistics, especially considering the complex nature of many machine learning algorithms already mentioned above.[76]

C. Report Summary

The report is an essential part of the audit process. Using guidelines such as FOIA, a report that details algorithmic inputs and outputs as well as an assessment by the auditors follows legal precedent while allowing the public to know as many of the facts as possible. These facts will give the public enough knowledge to be able to advocate effectively and demand specific change, but without compromising police departments' objectives.

V. Conclusion

Predictive policing may discriminate against protected groups. This bias may be intentional or unintentional. Because companies and police departments are reluctant to release

[75] *See* Tal Z. Zarsky, *supra* note 35, at 1534-1536.
[76] *Supra* note 38, at 24.

information about how predictive policing works, policymakers and the public do not know enough to adequately debate its benefits and drawbacks.

Nonetheless, police departments in cities across the nation will continue to implement predictive policing. This will fundamentally change the way that we are policed, whether we like it or not. Without transparency and accountability, potential discrimination by predictive policing algorithms may continue unnoticed and unchanged. We need changes in the present to protect the cities and neighborhoods of the future.

I propose that legislatures should solve this problem by providing transparency and accountability through regular audits of predictive policing technology used by police departments. Audits can discover and then report on discrimination in the algorithmic process. Regular auditing would also give companies and police departments an incentive to find and solve these problems, either on their own or using the audit's suggestions.

The audits would review every step of the predictive policing process using a variety of statistical techniques. Discrimination would be assessed using a standard of disparate impact. An audit report would then be released to the public, including appropriate information and recommendations to companies and departments on how to improve their algorithms. These reports would also recommend legal action by proper authorities if necessary.

If the public and policymakers know what is happening, they can better debate what action should be taken. Predictive policing, properly used, has the potential to make law enforcement more effective and efficient. However, this should not be achieved without understanding what it does and whether it is discriminatory. Legislatures can solve this problem by requiring regular audits of the algorithms that police departments use. These audits can be a powerful tool to enhance transparency and accountability for predictive policing.