Predictive Policing in the US: How Far is Too Far?

A Research Paper submitted to the Department of Engineering and Society

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

Alexander Matthew Pascocello

Spring 2022

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

Advisor

Kent Wayland, Assistant Professor, Department of Engineering and Society

In the U.S., how are advocates, law enforcement agencies, and tech companies competing to determine the proper extent and applications of predictive policing?

Introduction

"Please tell me how much confidence you, yourself, have in each one -- a great deal, quite a lot, some or very little?" Gallup has been asking Americans this question since 1973, polling the public's confidence in a range of public institutions since the Watergate scandal. In 1993, this question began being asked about the police, and in the summer of 2020, those with "a great deal" or "quite a lot" of confidence fell below the majority for the first time in the question's 27 year history (Ortiz, 2020; Brenan, 2020).

The ideal role of a police force is to fairly maintain a safe and orderly environment for their communities. Police would justly enforce laws, deter crime and disorder, and provide a sense of security and safety for everyone they watch over. Unfortunately, this is rarely the case. Decades of abuse of minorities at the hands of police officers have led to a strong and deep-seated distrust of police by those groups. In recent years, multiple prominent examples of this mistreatment have become highly publicized, causing growing outrage at policing measures and support of the targeted groups.

This issue is only exacerbated by the introduction of predictive policing software. Like many other predictive softwares, much of predictive policing relies on analysis of vast amounts of data to reveal relationships that would otherwise go unnoticed. This data is used to "predict where and when specific crimes are most likely to occur," or even whether certain people will be involved in crime ("Are We at a Tipping Point", 2020; Sandhu, 2020). Police forces can use these predictions to distribute their officers to areas deemed at-risk of crime. However, these

projections are developed with historical crime data, which is problematic due to the huge social biases already ingrained in the data. This leads the software to end up enforcing these biases to a much larger degree. In doing so, a very dangerous feedback loop of unjust policing efforts is created that unequally and negatively impact minority communities, especially African Americans.

If left unchecked, it is possible that predictive policing software operating off of this biased data will lead to even more biased policing. It is therefore important to study the use and effects of predictive policing software to determine if they increase bias in law enforcement. An effort can then be made to minimize or eliminate bias from these systems, or to remove the systems entirely, to allow for the least biased form of policing.

U.S. Policing: A Troubled Past

To understand the effect predictive policing software has on bias in policing, it is important to first discuss the inherent injustices built into modern policing at its founding. Policing in the United States mostly started as an informal community-based "watch," and remained this way until the 1830s, when major cities started implementing centralized municipal police departments (Gargurevich, 2013). These departments, whose creation were guided by the new commercial elites of the time, were made to help curb the "disorder" of the growing cities. However, they were more for social control than crime control, as they gave the elite a way to control the working class citizens. In fact, they were able to mask this social control as crime control by the creation of "dangerous classes." As Gargurevich (2013) states, "[t]he suggestion was that public drunkenness, crime, hooliganism, political protests and worker "riots" were the products of a biologically inferior, morally intemperate, unskilled and uneducated

underclass…This underclass was easily identifiable because it consisted primarily of the poor, foreign immigrants and free blacks" (p. 3).

The foundation of more rural police departments in the South followed a different path: they were born out of the former "Slave Patrols." The purpose of these patrols was to catch runaway slaves, discipline slaves for violating plantation rules, and provide a form of organized terror to deter slave revolts. After the Civil War and the emancipation from slavery, the police departments developed out of these patrols continued to control the newly freed slaves and strongly enforce Jim Crow segregation laws (Gargurevich, 2013).

While society has made much progress since the formations of these departments, police forces still struggle with their roots, especially with discrimination against African Americans. One 2018 briefing by the U.S. Department of Justice showed that "black people were overrepresented among persons arrested for nonfatal violent crimes (33%) and for serious nonfatal violent crimes (36%) relative to their representation in the U.S. population (13%)" (Beck, 2021). For all offenses, these numbers lower to 27%, but this is still more than twice the representation relative to population size ("Arrests by offense", n.d.). In another study from 2010 to 2015 that reviewed 19,000 cases of police use-of-force incidents, black residents were 3.6 times more likely to have force used against them than white residents (273 vs. 76 incidents per 100,000 interactions, respectively) and 2.5 times more than the overall rate of 108 incidents per 100,000 interactions (Williams, 2016).

Sadly, lethal examples of police interactions with African American communities are too prevalent. Eric Garner, Michael Brown, Tamir Rice, Alton Sterling, Philando Castile, Breonna Taylor, and George Floyd, to name just a few prominent examples, were all killed by police in the last eight years ("George Floyd", 2021). Even the shooting of Trayvon Martin by George

Zimmerman, whose acquittal sparked the beginning of the Black Lives Matter movement, occurred just barely over a decade ago ("The Black Lives Matter Movement", n.d.).[1] The recency of these events and data make it clear that discriminatory policing still exists in the modern United States, and that serious work still needs to be done in order to create a just and fair policing system that works to protect everyone equally.

Bias in Software Systems

Bias present in software systems must also be discussed before looking at software's application in policing. It is dangerous but common to assume that algorithms and artificial intelligence are inherently unbiased. After all, it sounds strange that algorithms could have biases when they seem to lack the emotions, conceptions and pretexts about the world that humans gather through life experiences. This thought process overlooks one issue that is summarized well in an article on the topic: while the algorithms are not inherently human, "[p]eople write the algorithms, people choose the data used by algorithms and people decide how to apply the results of the algorithms" (Best & Rao, 2021).

Bias can enter software in a myriad amount of ways and at any time in its life cycle. This can make it very difficult to pinpoint an exact 'location' or 'event' where the software becomes biased. It may be programmed directly into the code by a biased programmer, usually subconsciously. It may stem from the software being used in a way or by a group it was not 'expected' to be used for. One common way bias works its way into a software system is by inputting biased data. This is not done on purpose, but is rather a byproduct of most real-world data being inherently biased. Artificial intelligence learns by processing and forming correlations

[1] It should be noted that George Zimmerman was not a police officer. This was included to highlight how recent the Black Lives Matter movement is.

on large amounts of data, “which can include biased human decisions or reflect historical or social inequities, even if sensitive variables such as gender, race, or sexual orientation are removed” (Manyika, Silberg, & Presten, 2019).

In fact, this removal of sensitive variables is one reason why bias becomes overlooked. It is assumed that bias towards these variables cannot form if they are not included in the calculator, and so as long as it is removed then there won’t be bias present. This assumption is incorrect because it does not account for proxy attributes, which are “seemingly innocuous attributes that correlate with socially-sensitive attributes, serving as proxies for the socially-sensitive attributes themselves” (Johnson, 2020). These attributes may not seem to be related to the sensitive attributes, but an underlying correlation between the two means that including the less sensitive attribute will have a clear effect on the sensitive attribute. This can make it very difficult to identify where certain bias comes from and can make it challenging to remove bias from these systems. Even when it is clear that bias is coming from certain proxy attributes, removing them may ruin the accuracy of the program.

Examples of this bias are unfortunately prevalent: Google’s sentiment analyzer, made to determine if a sentence had a negative or positive connotation, gave sentences with the words “Jew” and “homosexual” a negative score (Thompson, 2017); error rates of three commercially available face-recognition softwares did not fall below 0.8% when analyzing light skinned men, but reached over 20% for one program and 34% for the other two when attempting to recognize darker-skinned women (Hardesty, 2018); and Amazon scrapped an AI recruiting tool that, by analyzing a decade of mostly men’s resumes from a male-dominated tech field, taught itself that men were more desirable candidates than women (Dastin, 2018).

By studying software bias, an effort can be made to understand where bias comes from and how it enters software systems. This can be used to help discover when it is present and take steps towards mitigating it or removing it from the system.

Predictive Policing

Over the last 10 years, police forces in the U.S. have begun to use software to identify crime trends and assist in law enforcement (Lee, 2020). As the combination of policing and software, predictive policing attempts to offer all of the benefits of software analysis while trying to solve some of the problems that plague modern policing.

There are two main categories of predictive policing software: location- and person-based. Person-based algorithms use a person's attributes, "such as their age, gender, marital status, history of substance abuse, and criminal record, to predict who has a high chance of being involved in future criminal activity" (Heaven, 2020). A person's race is not included in these attributes, as doing so would be illegal. However, there are plenty of proxy attributes such as a person's zip code, education and socioeconomic status that are allowed. This type of software can be used by police to monitor those deemed more likely to commit a crime, as well as by judicial systems to predict how likely someone is to recommit criminal behavior after being released (Heaven, 2020).

In contrast, location-based algorithms use information about a location to determine the likelihood of crimes occurring there. This can include information such as time of day, weather, historical crime rates, and large gatherings like schools and sports games (Heaven, 2020). Police forces can use the predictions created by these softwares to patrol areas determined to be at a higher risk of criminal activity, with the goal of more efficient and effective policing. For this

paper, location-based algorithms will be the main focus, as they are primarily used by police forces and not in judicial manners like person-based systems.

Producers of predictive policing software, such as Equivant, PredPol and ShotSpotter, argue that current policing is "too subjective," and that these programs would "help officers objectively determine where and when to police and, therefore, more effectively…prevent crime" ("Are We at a Tipping Point", 2020; Sandhu, 2021). Instead of assigning officers areas to patrol based on human intuition or geography, software can create correlations between location and crime based on historical crime events and send patrols to areas that need them more. The proposed benefits of using this software includes the removal of human bias in patrolling, faster response times to crime, and higher crime deterrence by increasing the police presence in high-risk areas.

Garbage In, Garbage Out

While predictive policing programs may sound like a good idea, current attempts at such software seem to only worsen present bias and offer little benefits to police agencies and even less to those being policed. The problem with these programs arises from how they make their predictions. Because many of these algorithms use historical crime data as part of their dataset, any correlations found are inherently built on the social biases of real-world law enforcement.

Brian Jefferson (2016), a professor at the University of Illinois Urbana–Champaign's Department of Geography and GIScience, states that "reliance on official crime statistics works to further entrench and legitimize the geographic knowledge and practices of racialized policing." In fact, a study observing the Chicago Police Department's use of predictive policing conducted by Mr. Jefferson (2016) concluded that "predictive crime mapping does not incur

more precise applications of police force but rather legitimizes the widespread criminalization of racialized districts". This can be demonstrated as a negative feedback loop in which more local policing leads to more crime observation and enforcement, which in turn leads software to suggest even more local policing.

Another issue with predictive policing software is the possibility of increasing an officer's confirmation bias about crime in a specific area. When a program predicts that one area will have a higher likelihood of crime occuring, it increases the possibility that "algorithmic output triggers possibilistic thinking, as police officers patrolling in the designated area are guided by the expectation that there must be a [criminal] around" (Egbert & Krasmann, 2018). This mindset means that anyone in the area is under much more suspicion, regardless of what they are doing, because of the "speculative connection between the spatiotemporal crime prediction and the risk potential of the people present at that location." The implications of this are very dangerous, as it may compound with other implicit biases to create harsher policing environments even with no increase in crime.

Even without all of these negative effects, predictive policing software seems to offer little if no benefits. One study done in Los Angeles and Kent County in southern England used a model "that estimates the risk associated with both long-term hotspots and short-term models of near-repeat risk" and found a 7.4% decrease in crime as a function of patrol time (Mohler, Short, Malinowski, Johnson, Tita, Bertozzi, & Brantingham, 2015). Although promising, this drop cannot be solely explained by the use of the models in the study. Another study in Shreveport, Louisiana in 2012 tested a location-based algorithm named PILOT (Predictive Intelligence Operational Targeting) to determine areas of higher criminal activity. The conclusion: "there is no considerable evidence that the application of PILOT leads to a reduction in crime rates when

compared to the control districts which used conventional crime mapping" (Meijer & Wessels, 2019).

Because of these issues, a growing number of people and groups are starting to question the morality of predictive policing software. These groups argue that their use is immoral due to unavoidable bias in data, lack of transparency, and increased policing of minority groups, all of which are caused "by the social reality on the basis of which the algorithms take form" (Benbouzid, 2018; Gilbertson, 2020). These groups demand accountability and transparency, and use auditing and lawsuits when necessary to achieve this goal (Benbouzid, 2018). They include activists, researchers, and even those who help make the software (Castelvecchi, 2020; Durán, 2019). As one of the targets of these demands and lawsuits, some large police organizations that use these or similar technologies such as the NYPD and LAPD have phased out the use of predictive policing technologies. However, others have simply made their own to avoid public scrutiny (Gilbertson, 2020; Lau, 2020).

Solutions

In order to create a full picture of the problems surrounding predictive policing, it was important to break predictive policing into its subsections: software and policing. Approaching possible solutions must be done in a similar manner, as these problems tend to stem from issues present in each of the subsections.

Police reform has been a hot button issue over the last decade due to numerous highly public instances of police brutality, including those discussed earlier in the paper. While opinions on what should be done vary, a key issue stopping American police reform is the local level at which police forces operate. There are more than 18,000 police agencies in the United States,

over 15,000 of which are organized at the city or county level (Stoughton, Noble & Alpert, 2020). Because of this, when reform does happen, it tends to be incredibly localized, and usually only occurs after a controversial event. A more universal federal policing policy could allow for sweeping reforms that apply to all agencies without the need for a controversial or tragic event (Armstrong, 2016). These policies should focus on broad actions that could improve policing as a whole, such as strict data collection, mandatory bias and de-escalation training hours, and clear definitions on use of force. To make sure that the policies are actually enforced, the withholding of federal funds could act as an incentive for uncooperative departments.

From the software perspective, the biggest issue is biased data. As discussed in *Bias in Software Systems*, removing bias from software is more difficult than it seems, but this does not mean no attempt should be made to. Andrew Ferguson (2017), a law professor at American University's Washington College of Law, argues that three steps can be taken to help reduce error in predictive policing software. First, by simply acknowledging there is error in the data, something can be done about it. "Acknowledging error," he states, "does not discount the value of predictive technologies, but only qualifies the findings and tempers the unquestioning acceptance of the information" (p. 1151). Predictive policing companies hesitate to do this because it shows their algorithms have problems, which may mean a loss in trust or economic value.

Once the error is acknowledged, it can be caught and corrected (Ferguson, 2017). Auditing mechanisms and data services can be put in place to catch, correct and clean data that may be inaccurate, duplicated or of poor quality. While this is challenging and cannot be done perfectly due to the vast amounts of datasets used in predictive policing, any attempt is better than none at all.

Lastly, training police officers to input data correctly and consistently is needed in order for the system to work properly (Ferguson, 2017). Combining this with any technology that makes inputting data easier for officers can lead to cleaner, clearer and more accurate data for predictive policing software to work off of.

While these solutions as written about are specific to predictive policing software, they can be easily broadened to include other types of software systems. Companies are reluctant to admit there is bias in data they use to train their algorithms because it looks bad and may be expensive to fix. However, acknowledging that bias is bound to exist in data and taking steps such as cleaning data and training those inputting information can help to remove, or at least reduce, any biases that may appear.

Conclusion

Predictive policing covers a wide range of social issues including proper use of technology, police reform, racism and discrimination in the US, and more. It is not an inherently bad thing, but if proper steps are not taken to remove bias from its operation and from police forces around the country, predictive policing algorithms have the potential to further entrench biases already present in modern policing. It is important to study the effects its use has to better understand how similar technologies may impact our society, as there are plenty of other influential software systems with ingrained biases that impact our daily lives without us even realizing. Advancements in bias reduction in predictive policing can be applied to these systems in order to reduce overall bias and improve equality in society as a whole.

References

Are We at a Tipping Point in Police-Community Relations? (2020, June 11). *Predictive Policing Blog.* https://blog.predpol.com/are-we-at-a-tipping-point-in-police-community-relations

Armstrong, K. (2016, July 13). How to Fix American Policing. The Marshall Project. https://www.themarshallproject.org/2016/07/13/how-to-fix-american-policing

Arrests by offense, age, and race. (n.d.) *Statistical Briefing Book; Office of Juvenile Justice and Delinquency Prevention.* https://www.ojjdp.gov/ojstatbb/crime/ucr.asp?table_in=2&selYrs=2018&rdoGroups=3&rdoData=rp

Beck, A. (2021, Jan.). Race and Ethnicity of Violent Crime Offenders and Arrestees, 2018. *Statistical Brief, DOJ.* https://bjs.ojp.gov/content/pub/pdf/revcoa18.pdf

Benbouzid, B. (2018, Feb. 26). Values and Consequences in Predictive Machine Evaluation. A Sociology of Predictive Policing. *Science & Technology Studies, Volume 31.* https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3123315

Best, M., & Rao, A. (2021, Jan. 18). Understanding algorithmic bias and how to build trust in AI. *PwC.* https://www.pwc.com/us/en/tech-effect/ai-analytics/algorithmic-bias-and-trust-in-ai.html

Brenan, M. (2020, Aug. 12). Amid Pandemic, Confidence in Key U.S. Institutions Surges. *Gallup.* https://news.gallup.com/poll/317135/amid-pandemic-confidence-key-institutions-surges.aspx

Castelvecchi, D. (2020, June 19). Mathematicians urge colleagues to boycott police work in wake of killings. *Nature*. https://www.nature.com/articles/d41586-020-01874-9

Dastin, J. (2018, Oct. 10). Amazon scraps secret AI recruiting tool that showed bias against women. *Reuters.* https://www.reuters.com/article/us-amazoncom-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-thatshowed-bias-against-women-idUSKCN1MK08G

Durán, L. (2019, Dec. 11). Why We Need Data For Black Lives. *Forbes.* https://www.forbes.com/sites/ashoka/2019/12/11/why-we-need-data-for-black-lives/?sh=2e0a701b7bd4

Egbert, S. & Krasmann, S. (2018, Dec. 14) Predictive policing: not yet, but soon preemptive? *Policing and Society, Volume 30.*

Ferguson, A.G. (2017). Policing Predictive Policing. *94 WASH. U. L. REV. 1109*. https://openscholarship.wustl.edu/law_lawreview/vol94/iss5/5

Gargurevich, I. A. (2013, July 30). The History of Policing in the United States. *Academia.* https://www.academia.edu/43389392/The_History_of_Policing_in_the_United_States?auto=citations&from=cover_page

Gilbertson, A. (2020, Aug. 20). Data-Informed Predictive Policing Was Heralded As Less Biased. Is It? *The Markup.* https://themarkup.org/ask-the-markup/2020/08/20/does-predictive-police-technology-contribute-to-bias

George Floyd: Timeline of black deaths and protests. (2021, April 22). *BBC News.* https://www.bbc.com/news/world-us-canada-52905408

Hardesty, L. (2018, Feb. 11). Study finds gender and skin-type bias in commercial artificial-intelligence systems. *MIT News.* https://news.mit.edu/2018/study-finds-gender-skin-type-bias-artificial-intelligence-systems-0212

Heaven, W. D. (2020, July 17). Predictive policing algorithms are racist. They need to be dismantled. *MIT Technology Review.* https://www.technologyreview.com/2020/07/17/1005396/predictive-policing-algorithms-racist-dismantled-machine-learning-bias-criminal-justice/

Jefferson, B. (2016, July 1). Predictable Policing: Predictive Crime Mapping and Geographies of Policing and Race. *Annals of the American Association of Geographers, Volume 108.*

Johnson, G.M. (2020, June 20). Algorithmic bias: on the implicit biases of social technology. *Synthese* 198, 9941–9961. https://doi.org/10.1007/s11229-020-02696-y

Lau, T. (2020, April 1). Predictive Policing Explained. *Brennan Center for Justice.* https://www.brennancenter.org/our-work/research-reports/predictive-policing-explained

Lee, J. (2020, Nov. 17). How AI technology is helping solving crime. *Police1*. https://www.police1.com/police-products/police-technology/police-software/articles/how-ai-technology-is-helping-solving-crime-7vb577RVrWliW57H/

Manyika, J., Silberg, J., & Presten, B. (2019, October 25). What Do We Do About the Biases in AI? *Harvard Business Review.* https://hbr.org/2019/10/what-do-we-do-about-the-biases-in-ai

Meijer, A. & Wessels, M. (2019, Feb. 12) Predictive Policing: Review of Benefits and Drawbacks. *International Journal of Public Administration*.

Mohler, G. O., Short, M. B., Malinowski, S., Johnson, M., Tita, G.E., Bertozzi, A.L., & Brantingham, P. J. (2015) Randomized Controlled Field Trials of Predictive Policing, *Journal of the American Statistical Association*, 110:512, 1399-1411.

Ortiz, A. (2020, Aug. 12). Confidence in Police Is at Record Low, Gallup Survey Finds. *The New York Times.* https://www.nytimes.com/2020/08/12/us/gallup-poll-police.html

Sandhu, A. (2021, Aug. 4). The 'uberization of policing'? How police negotiate and operationalise predictive policing technology. *Policing and Society*.

Stoughton, S. W., Noble, J. J., & Alpert, G. P. (2020, June 3). How to Actually Fix America's Police. *The Atlantic.* https://www.theatlantic.com/ideas/archive/2020/06/how-actually-fix-americas-police/612520/

The Black Lives Matter Movement - A Brief History of Civil Rights in the United States. (n.d.). *HUSL Library at Howard University School of Law.* https://library.law.howard.edu/civilrightshistory/BLM

Thompson, A. (2017, Oct. 25). Google's Sentiment Analyzer Thinks Being Gay Is Bad. *VICE.* https://www.vice.com/en/article/j5jmj8/google-artificial-intelligence-bias

Williams, T. (2016, July 7). Study Supports Suspicion That Police Are More Likely to Use Force on Blacks. *The New York Times.* https://www.nytimes.com/2016/07/08/us/study-supports-suspicion-that-police-use-of-force-is-more-likely-for-blacks.html