Predictive Policing in Pittsburgh: A Primer

Contents

1 What is predictive policing?

Predictive policing is the practice of using computational models to predict where crime will occur before it happens. While police departments and crime analysts have been trying for years to predict crime trends (in order to better distribute policing resources), predictive policing refers to the use of large-scale artificial intelligence and machine learning[1] models to do this task. These models can process large amounts of data, and provide predictions faster and on a larger scale than previous methods. Generally, predictive policing models are owned by for-profit companies, which can limit public access to information about the details and functioning of such models.

There are two main types of predictive policing models: location-based models and individual-based models. Location-based models predict *where* crime is likely going to happen, and individual-based models predict *who* is likely to be a perpetrator or victim of crime.

Predictive policing models use past data (such as 911 calls, minor arrests, major arrests, criminal histories, and miscellaneous reports) to train an algorithm. This algorithm attempts to learn patterns of past criminal activity in order to predict future crime locations and involved parties. Police then target these locations or people with interventions to prevent crime before it happens–for instance by driving through an area, walking through that area looking for criminal activity, or visiting a person's home.

1.1 Location-based predictive models

Location-based policing models provide geographic information on clusters of crime. These models use broad information about the city (e.g. 911 calls, police reports, prior incidents, etc.) to predict "hot spot" locations where previous crimes are likely to happen again. These models may predict chronic hot spots (specific areas where crime frequently occurs) or temporary hot spots (that change over time or based on events in the city). **Pittsburgh uses a temporary hot spot prediction model**.

[1] Machine learning is the use of computational and statistical models to extract insights from large amounts of data.

Hot spot prediction. In the words of the creators of Pittsburgh's model, "This project...use[s] predictive modeling and crime mapping to predict and target crime flare-ups in micro areas on the scale of a few to several blocks and a few days to weeks before they occur" [1].

Prior literature in criminology emphasizes that crimes do not occur uniformly over time and space. They generally occur in specific locations that present likely opportunities for successful crimes. Hot spot policing is based on the notion that these locations can be predicted in advance using prior data from these locations.

The Pittsburgh model specifically uses "leading indicator" variables which may be predictive of the outcome. To the best of our knowledge, these variables, which are explicitly mentioned in the researchers' preliminary report [1], are:

1. Criminal Homicide
2. Forcible Rape
3. Robbery
4. Aggravated Assault
5. Burglary
6. Larceny/Theft
7. Vehicle Theft
8. Geotagged 911 calls
9. "Other Crimes."

We do not have information on what "other crimes" are included in the model. From one of the authors' prior work in Chicago's CityScan hot spot policing project and from the researchers' grant proposal for Pittsburgh's hot spot policing [2], we speculate that these "other crimes" are likely to include (but are not limited to):

1. "Shots fired" reports
2. "Gang Loitering" reports
3. Simple battery
4. Vandalism
5. Data on mental health and substance abuse
6. Data on high-risk youths
7. Data on homeless individuals
8. Data on persons living in public housing.

A machine learning model trained on these inputs would then learn to map "leading indicator" variables and their date of occurrence to specific times and city blocks that are likely to see a resurgence in the reported indicator variables. The authors focus on predicting Part I Violent crimes (namely criminal homicide, forcible rape, robbery, and aggravated assault) in their model.

Potential drawbacks of hot-spot-based predictive policing. Hot-spot-based predictive policing methods are likely to generate feedback loops that amplify existing biases. This is because predictive policing models use (biased) historical data to make predictions of where crime will occur. However, since the data is biased (e.g. racially biased), the model's predictions will be biased. These (biased) predictions are then used to inform where police should patrol, causing more police to be sent to already over-policed areas, and creating even more arrests in these areas. This (biased) arrest data is then fed back into the predictive policing model to help make future predictions, leading to even more biased predictions, leading to even more biased patrolling and arrests, and so forth. In general, predictive policing models have been empirically shown to be highly susceptible to feedback loops [3].[2]

In the preliminary report on Pittsburgh's model [1], the authors attempt to address the issue of feedback loops by "select[ing] hot spots based on predictions of [Part I Violent] crimes...[as they] believe that reporting of these extremely serious and violent offenses is less likely than other crime types to depend on the presence or absence of police in an area." They do not go into further detail about why this is a sound assumption, nor do cite sources which may confirm this assumption.

Additionally, while the Pittsburgh model only predicts hot spots for Part 1 Violent crimes, they make these predictions using all of the information detailed above, which includes lower-level crimes. Arrest rates for lower-level crimes are known to be disproportionately higher for black and Hispanic populations; basing prediction of Part 1 Violent crimes off these biased inputs is potentially discriminatory [4].

1.2 Individual-based predictive models

Individual-based predictive policing attempts to predict which specific people will be involved in future crimes. These models use data about people, such as criminal history, demographic information (e.g. age), and recent interactions with the police, in order to predict who is likely to be involved in a crime. It is often unclear how people get on and off of the lists of people that the algorithm searches through in order to make predictions.

[2]It should be noted that referenced study trained a model on historical crime incident data, but evaluated the feedback effects on a synthetic population of likely drug users based on public health data.

In the grant proposal about Pittsburgh's model, the researchers mention that their intention to also create individual-level crime prediction models [2], although they have distanced themselves from that idea in more recent publications about their project. More information about such programs in other US cities is described in Section 4.

Potential drawbacks of individual-based policing. Keeping lists of "probable" offenders/victims has the potential to exacerbate racial disparities in policing. This is because the people who are likely to make it onto such lists (based on higher arrest count or interactions with police) have a larger chance of being black or Hispanic, due to racial bias in policing in America. Once someone is on that list, their likelihood of being targeted by the police increases. Based on information about Chicago's individual-based policing program, this is exactly what happened: while the list was claimed to be used as an investigative tool to recommend people for "police visits" and "referral to social services," the list was mostly used to make arrests. Notably this list consisted of 56% of all black men aged 20 to 29 and 23% of all Hispanic men in Chicago aged 20 to 29, while only 6% of all white men in the same age group were included [5]. More information about the repercussions and use of predictive policing models in other cities is described in section 4 of this primer.

2 Background on predictive policing in Pittsburgh

As detailed above, the Pittsburgh Bureau of Police has been using a predictive policing model since 2016, which was developed under Chief McClay by researchers at CMU. It was initially piloted in August 2016 in Homewood. However, at the request of Mayor Peduto, a renewed city-wide experiment launched in May 2017, funded by CMU Metro21. More information about Pittsburgh's model can be found at the following sources [1, 2].

To the best of our knowledge, Pittsburgh's model is trained on data gathered from the Western Pennsylvania Regional Data Center (WPRDC), which lists the age, gender, race, offense, arrest time, and incident location, among other variables for every arrest and incident that occurs in Western Pennsylvania. It then predicts crime hot spots in order to proactively police communities through patrols.

This model is currently being used as part of a *randomized control trial* where Pittsburgh is separated into control areas (where the predictive policing model is not used) and treatment areas (where new patrols are allocated based on the predictive policing model's predictions). The patrols in treatment areas, which may result in arrests or new incident reports, are then included in the aggregate WPRDC data. The predictive policing model is then retrained with this new data "on a rolling basis throughout the evaluation period using two years of training data" [1].

The authors of the program make efforts to curb potential bias and the downstream effects of overpolicing by predicting temporary hot spots (which change with time and based on events in the city). The researchers have found that temporary hot spots are 5 times more equitable than chronic hot sports with regard to the distribution of patrol efforts. They account in their analysis for crime spillover effects, i.e. the possibility that patrols are just displacing crime rather than decreasing it, and determine that this form of hot spot policing does not displace crime and does in fact lower it.

3 Our concerns about predictive policing in Pittsburgh

(Lack of) transparency. As alluded to in the above sections, it is unclear what exact variables are used to train Pittsburgh's model. While the preliminary report on hot spot policing in Pittsburgh (authored by the creators of Pittsburgh's model) mentions some variables, describes the model's current use, and evaluates the model's effectiveness, there is no way to verify the claims that the authors make, as the model itself is not public.

Carnegie Mellon University's "Tech4Society" group attempted to initiate a Right-To-Know request in order to receive information on the exact variables and locations of control and treatment areas over the 3 year period. However, these requests were not be fulfilled, with the rationale that "Hot Spot locations would jeopardize the safety of officers and the effectiveness of crime patrols. These records are exempt under RTKL. 65 P.S. §§67.708(b)(1)(ii); (b)(2)" [2]. A similar response was received for specific model information, with the rationale that "trade secrets and confidential proprietary information are exempt under RTKL. 65 P.S. §§67.708 (b)(11)" [2].

Without information about treatment/control locations or the specific variables used, it is impossible for anyone outside of the creators of the model or the Pittsburgh Bureau of Police to verify or reproduce the results of the study. Restricting this information hinders any attempt at a third party audit of the model and its use in Pittsburgh.

Lack of community engagement. While this program has been ongoing as a pilot project for over 3 years, only partners in the city government have been consulted about its implementation.[3] Other stakeholders, such as Pittsburgh community groups or groups who could conduct a third-party audit of the model, did not (as far as we are aware) play a meaningful role in the development and deployment of the model.

While the researchers said they had made connections with other groups, the strength of these connections may have been overstated. In particular, Carnegie Mellon University's "Tech4Society" group reached out to the researchers regarding the broader concerns with their predictive policing model and the concerns about the lack of community involvement. The authors of the project stated via email that "we have had positive and productive discussions with the Center for Policing Equity, who have both stated their appreciation for the project's emphasis on equity issues, as well as providing us with constructive and valuable advice that we will work to incorporate as the project moves forward." After Tech4Society reached out to the Center for Policing Equity about these claims, the center responded with, "While we have had a few conversations with this research team, their letter suggests we have had more substantive contact with them than we have ... We have communicated to the research team that we would not take part in their research as we have concerns regarding how it was both conceptualized and rolled out to the community, and none of our research staff engaged with them substantively."

How police use these models. The preliminary report on hot spot policing in Pittsburgh [1] and the creators' grant proposal [2] state that the intended intervention when a hotspot is predicted is for police to patrol those hotspots in order deter future crimes. Specifically, the authors state, "The target treatment dosage for each 500 ft. by 500 ft. hot spot cell was three 15-minute foot patrols or nine 5-minute patrols per hot spot per day" [1]. Officers additionally have a "hot spot dashboard" which provides hot spot locations to officers while they are on patrol.

One of the notable issues with the method outlined is that while the authors intend for police to act solely as a deterrent, the researchers do not have the capability of enforcing this behavior. Throughout their study, data is not collected about whether officers are creating new arrests, tickets, reports, etc. when patrolling a predicted hot spot, and such information is not used to evaluate the effectiveness of the intervention. Without transparent information on the field experiment, unexpected effects of the patrols cannot be tracked.

Potential bias. An additional concern involves Pittsburgh's propensity toward having higher arrest rates for people of color: Pittsburgh arrests black adults at a rate 1.5 times higher than the national average, and black youths at 2.5 the national average [6]. Since this historical arrest data is racially biased, and predictive policing models use this (biased) arrest data to make predictions, the predictions these models make will likely also be biased.

Predictive policing models are also commonly motivated by the theory of "broken windows policing," which is the concept that urban disorder (such as abandoned and run-down buildings) are indicators of future crime. For instance, as the creators of Pittsburgh's model cite in their project proposal, "it is well established that signals of urban disorder (eg., 'broken windows') can lead to or attract criminal behavior that hardens over time" [2]. However, broken windows policing has been long debated. There remains a significant amount of disagreement as to whether or not there is an actual long-term link between disorder and serious crime [7, 8]. In addition, researchers have found that policies that emphasize neighborhood disorder as an indicator of crime negatively impact minority groups in the US [9].

Feedback loops The Pittsburgh model does attempt to minimize the effects of feedback loops and overpolicing by focusing on temporary hot spots. The intuition behind these temporary hot spots is that they spread out resources so that if an incident occurs while an officer is on patrol, it will not influence future models to send officers back to the same area.

Despite these measures, we are concerned that feedback loops can take different forms than just those addressed by temporary hot spots. These loops can be learned through common features of offenders, such as age, race, gender, etc. As mentioned above, even if officers are sent to different neighborhoods, if incidents occur more often with specific demographics, the model will learn to associate these features such as age, race, gender, etc with locations of crime, which in turn causes the model to send officers to these areas more frequently.

Links to prior legal issues related to predictive policing

- Harvard Civil Rights - Civil Liberties Law Review Minority Report: Why We Should Question Predictive Policing

[3]Based on the Right to Know request, these partners were limited to the Pittsburgh Department of Innovation and Performance, the Allegheny County Department of Human Services, the Pittsburgh Bureau of Police, and the Pittsburgh Department of Public Works.

- The Undue Influence of Surveillance Technology Companies on Policing
- Statement of Concern About Predictive Policing By ACLU and 16 Civil Rights Privacy, Racial Justice, and Technology Organizations

4 How has predictive policing affected other cities?

4.1 Los Angeles

Types of predictive policing models. Person-based and location-based. Due to concerns from the public, the person-based program was disbanded, and the location-based program is on hold.

(Lack of) transparency. There was little transparency about how individuals were put on the likely offenders list for the person-based model – probably partially due to the fact that there was no consistent protocol.

For the location-based model, police did publish micro-areas that they visited each day; however, this was only published towards the end of the program. The creators of the algorithm claim that the only three attributes used to predict location of future crime were crime type, location, and time of past crimes.

Community engagement. From what we can find, the LAPD had a symposium on predictive policing in 2010, and the program was rolled out in 2011. There is also an article from the National Institute of Justice Journal [10] reporting on the roll out of the program. We are unsure about the degree to which the community was informed of the program (town halls, press release, etc.) prior to the roll out. The LAPD has a civilian oversight panel that responded to concerns raised by activists about the use of predictive policing. There was a hearing about LAPD's use of predictive policing, and the oversight panel declined to let the LAPD move forward with their use of it, instead requesting more information on its efficacy.

Potential bias. In a city where black residents make up 9% of the population, 33% of individuals on the list of likely offenders were black. Black and Hispanic people combined made up 83% of the list.
Based on reports from WIRED about community feeling about the use of predictive policing, community members often felt targeted and noticed the increased police presence in their neighborhoods.
The Stop LAPD Spying Coalition conducted a survey that found the predictive policing program exacerbated racial bias by leading to more minority interactions with police in areas that were already over-policed.

Feedback loops. There was not enough data collected and released to verify whether feedback loops were happening with the use of LAPD's predictive policing model ("PredPol").

How it is used. According to the LA Times:

- **Location-based:** "Each day the PredPol program generates 10 "boxes," about 500 feet by 500 feet each, designated as zones for possible property crimes such as burglaries and car thefts. Officers are expected to patrol the areas when time allows during their shifts; the GPS system in patrol cars allows their movements to be tracked." "Specialized units are also sent to the [PredPol] areas." The officers look for suspicious activity and cars without plates, get out and engage in policing if they deem it is necessary.
- **Individual-based:** Technically, individuals are put on the list of watched people based on large number of points (e.g. gang affiliation is five points, being on probation is five points, and every contact with police has some point number). However, in practice, LAPD did not use consistent standards to calculate points and track the individuals–in fact, detectives could put people with zero points on the list, arbitrarily. It was reported that 44% of those labeled "chronic offenders" had either zero or one arrest for violent crimes.

Effectiveness at reducing crime. Based on review by Inspector General Mark Smith, there was not enough data collected to make a good case for the effectiveness of the model. Crime in the Hill District, where the data was used, went down by 13% during the use of PredPol, but crime also went down in areas unaffected by PredPol.

Sources

1. https://www.latimes.com/local/lanow/la-me-ln-lapd-data-20190312-story.html
2. https://www.govtech.com/public-safety/LAPD-Predictive-Policing-Tool-Raises-Racial-Bias-Concerns.html
3. https://www.latimes.com/local/lanow/la-me-lapd-precision-policing-data-20190703-story.html
4. https://www.wired.com/story/los-angeles-police-department-predictive-policing/
5. https://www.vice.com/en_us/article/ezp8zp/minority-retort-why-oakland-police-turned-down-predictive-policing

4.2 New York

Types of predictive policing models. Person-based, location-based.

(Lack of) transparency. It took two years of litigation by the Brennan Center for Justice to get information about how the NYPD tested, purchased, and implemented their predictive policing software. The Brennan Center commented that the long legal battle, as well as the "piecemeal" information that they eventually received, showed the NYPD's reluctance to release information about predictive policing.

Additionally, **by design**, NYPD's predictive policing model does not save the predictions it generates, nor does it keep an audit log of who creates or accesses them. This lack of saved data makes it extremely difficult to audit the model.

Community engagement. We could not find any information about community engagement in NY.

Potential bias. According to the documents received in the FOIL request by the Brennan Center, the algorithms used by the NYPD requested very rich data, many of which are strongly correlated with race and ethnicity, for example:

- Median income in the past 12 months for a given area
- How many people in the area receive food stamps or other public assistance income
- School enrollment and educational attainment of the population for 25 years and older
- Household language.

We note that it is unclear if this data was actually handed over, but it seems that accoring to an article from *The Daily Beast* the NYPD "is very willing to hand over and integrate" such demographic data with its preditive policing strategies.

Additionally, the model used stop-and-frisk data without even checking if these stop and frisk stops lead to arrests or convictions. Notably, New York was ruled to have been using stop-and-frisk in an unconstitutional manner in 2014.

Feedback loops. There was not enough data collected and released to verify whether feedback loops were happening with the use of PredPol in the NY. However, if the extremely sensitive and racially correlated demographic data was used as a part of crime prediction, it is quite conceivable that some feedback loops may have been created.

How it is used. It's not exactly clear, but there were several neighborhoods in Brooklyn and Queens that have been patrolled by PredPol since 2013. The algorithm decides how police resources like patrol cars are spread around the city, as well as when and were to use stop and frisk.

Effectiveness at reducing crime. While no studies we know of have been done to test the effectiveness of PredPol at reducing crime in NY specifically, during a 30-day test of the accuracy of PredPol at predicting crime in NY, it was found to be 29% accurate at predicting where crimes were going to occur. It was the most effective at predicting gun violence, but its accuracy for minor crimes was as low as 10%. Since it was predicting such low-level crimes with pretty low accuracy, it's easy to imagine how the model may send police to areas with high "risk" when they are not necessary, and end up over-policing minority areas.

Since the model explicitly does not save the results of its predictions, it is impossible to test the efficacy or potential biases inherent in this model.

Sources

1. https://www.thedailybeast.com/red-flags-as-new-documents-point-to-blind-spots-of-nypd-predictive-policing?ref=scroll
2. https://www.brennancenter.org/sites/default/files/Keystats%20Desired%20Data%20Elements.pdf
3. https://www.brennancenter.org/our-work/research-reports/nypd-predictive-policing-documents

4.3 Chicago

Types of predictive policing models. Person-based: the strategic subject list (SSL).

(Lack of) transparency. The Chicago Police Dept. denied records requests about the SSL from the Sun-Times, citing the probability that criminals may thwart the police with knowledge of the list. Chicago Magazine notes that this is strange since the Chicago police advertise the system as targeting individuals for social services, and not targeting individuals for arrests.

Community engagement. No information for community engagement found for Chicago.

Potential bias. The list of "strategic subjects" consisted of 56% of all black men aged 20 to 29 and 23% of all Hispanic men aged 20 to 29 in Chicago, while only 6% of all white men in the same age group were included [5].

Feedback loops. There was not enough data collected and released to verify whether feedback loops were happening with the use of predictive policing in Chicago.

How it is used. According to a RAND report, cited below, there was no consistent use strategy of the list. The police department proclaimed that the list was a list of both perpetrators and victims of gun violence, and would be used for interventions to decrease arrests (such as introduction to social services), as opposed to being used for arrests. This was not the case in practice: the vast majority of interactions between police and individuals on the SSL list were indeed arrests. According to Chiago Magazine, Chicago has been expanding its use of predictive policing capabilities to implement a "total overhaul" of mission assignments, which may include deciding where police resources are sent more globally across the city, not just with respect to those on the SSL list.

Effectiveness at reducing crime. The only study of the effectiveness of predictive policing was done by the RAND corporation [11] on the use of predictive policing in Chicago. They found that the effects of using predictive policing are minimal at best, and are essentially non-existent: individuals on the SSL are "not more or less likely" to become a victim of a homicide or shooting than the control group, and police use of the list also had no effect on citywide violence levels in Chicago.

The police replied by saying that the RAND report was on an outdated form of their system. However, the writers of the RAND report stated that they believed this didn't matter, since there was no consistent use strategy of the strategic subject list; no matter how accurate the list was, it would still not affect the police department's ability to successfully prevent crime.

Despite its ineffectiveness, the program is rapidly expanding in Chicago, with a recent $10 million dollar grant to the system.

Sources

1. https://www.chicagomag.com/city-life/August-2017/Chicago-Police-Strategic-Subject-List/
2. https://www.chicagomag.com/city-life/August-2016/Chicago-Police-Data/
3. https://abc7chicago.com/chicago-given-/10m-to-expand-predictive-policing-officer-training/3327651/

4.4 New Orleans

Types of predictive policing models. Person-based.

(Lack of) transparency. While the New Orleans government did not explicitly keep the program a secret, (contracts with Palantir, the company who provided the predictive policing service "Gotham", were on publicly available budgets), the use of predictive policing in New Orleans went under the radar for six years before there was any article published about it. The first public remark from the police department on the predictive policing partnership was in 2018, when Mayor Mitch Landrieu's press office noted that the partnership would not be continued. It started in 2012.

However, it was recently announced that New Orleans will start another person-based predictive policing program, much like the program in Chicago—it purports to be based around increasing access to social services, but no further information about the funding of these social programs, or how they will be provided to the community, is available.

Community engagement. Based on information we can find, there was no effort to engage the community on the decision to begin or end the program.

Potential bias. The NOLA program used data from field interview cards (FICs) that allow police officers to collect information about individuals after any sort of encounter, not just one that leads to arrest. This data is heavily racially biased, and likely lead to bias in the person-based predictive policing program.

Feedback loops. There was not enough data collected and released to verify whether feedback loops were happening with the use of predictive policing in New Orleans.

How it is used. According to a recently released powerpoint from the New Orleans police department, the Gotham software was an "integral part" of the investigative process. The software was used to predict who was involved in gangs in the city, and was involved in the arrest of 83 suspected gang members. Kentrell Hickerson, one individual who was convicted of gang violence and racketeering, attempted to protest his convictions and requested a re-trial, citing unfair treatment due to the lack of transparency surrounding the Gotham algorithm. His request was denied by the judge.

Effectiveness at reducing crime. According to The Lens NOLA, there is not sufficient enough data to support the effectiveness at the violence reduction programs, including Gotham, at this time.

After several "call-ins" of gang members, where gang members were pressured to provide information or receive much harsher policing and sentences for their next crime, there was a temporary downtick in crime. However, according to *The Lens NOLA*, this downtick did not last, and crime returned to normal levels.

Sources

1. https://www.theverge.com/2018/2/27/17054740/palantir-predictive-policing-tool-new-orleans-nopd
2. https://thelensnola.org/2018/10/24/months-after-end-of-predictive-policing-contract-cantrell-administration-works-on-new-tool-to-id-high-risk-residents/
3. https://www.nola.com/news/courts/article_03d1e5d6-ae5c-11e9-ad5d-63ea4d1bb62e.html
4. https://www.documentcloud.org/documents/5018779-Asher-PowerPoint.html#document/p15/a462822
5. https://www.aclu.org/blog/privacy-technology/new-orleans-program-offers-lessons-pitfalls-predictive-policing

5 What have other movements against predictive policing done?

1. ACLU issued a joint statement of concern with 16 other organizations (including the NAACP).
2. The Brennan Center has filed an Article 78 action against the New York Police Department for failing to respond adequately to a Freedom of Information Law request for records regarding the acquisition and implementation of predictive policing technology.
3. NYC passed a law to look into how information on agencies' automated decision systems may be shared with the public, and how agencies may address instances where people are harmed by these automated decision systems.

4. Newer predictive policing companies such as Civicscape have committed to algorithmic transparency by publishing a version of their source code on the online code repository Github. They have also pledged not to use their tools to predict drug crimes because of concerns that the bias present in crime data is too difficult to model out of their predictions.

5. The state of Washington has passed landmark algorithmic accountability laws.

6. Taking Algorithms to Court: When polite requests for transparency don't succeed, lawyers and activists have successfully argued that trade secrets and NDAs violate individual rights to due process.

7. ACLU-WA was instrumental in helping pass Seattle's 2017 surveillance ordinance, one of many transparency rules that city legislatures have passed in recent years.

References

[1] D. Fitzpatrick, W. Gorr, and D. B. Neill, "Hot-spot-based predictive policing in pittsburgh: A controlled field experiment," tech. rep., 2018.

[2] "Right to know 70-02-2019." `https://drive.google.com/open?id=1Vjsu91QYkD6NJi9cID0DbZSD32obaoFu`.

[3] K. Lum and W. Isaac, "To predict and serve?," *Significance*, vol. 13, no. 5, pp. 14–19, 2016.

[4] S. Project, "Report of the sentencing project to the united nations human rights committee: Regarding racial disparities in the united states criminal justice system," 2013.

[5] Y. Kunichoff and P. Sier, "The contradictions of chicago police's secretive list," *Chicago Magazine*, 2017.

[6] L. Davis and R. Bangs, "Pittsburgh's racial demographics 2015: Differences and disparities." `http://www.crsp.pitt.edu/sites/default/files/Final%20version%20for%20publishing.pdf`.

[7] B. E. Harcourt, *Illusion of order: The false promise of broken windows policing*. Harvard University Press, 2009.

[8] M. Cerdá, M. Tracy, S. F. Messner, D. Vlahov, K. Tardiff, and S. Galea, "Misdemeanor policing, physical disorder, and gun-related homicide: a spatial analytic test of" broken-windows" theory," *Epidemiology*, pp. 533–541, 2009.

[9] S. Childress, "The problem with broken windows policing." `https://www.pbs.org/wgbh/frontline/article/the-problem-with-broken-windows-policing/`, 2016.

[10] B. Pearsall, "Predictive policing: The future of law enforcement," *National Institute of Justice Journal*, vol. 266, no. 1, pp. 16–19, 2010.

[11] J. Saunders, P. Hunt, and J. S. Hollywood, "Predictions put into practice: a quasi-experimental evaluation of chicago's predictive policing pilot," *Journal of Experimental Criminology*, vol. 12, no. 3, pp. 347–371, 2016.