

Why big-data analysis of police activity is inherently biased

May 10 2017, by William Isaac And Andi Dixon



Credit: AI-generated image (disclaimer)

In early 2017, Chicago Mayor Rahm Emanuel announced a new initiative in the city's <u>ongoing battle with violent crime</u>. The most common solutions to this sort of problem involve <u>hiring more police</u> <u>officers or working more closely with community members</u>. But Emanuel declared that the Chicago Police Department would expand its



use of software, enabling what is called "predictive policing," particularly in neighborhoods on the city's south side.

The Chicago police will use data and computer analysis to identify neighborhoods that are more likely to experience violent crime, assigning additional police patrols in those areas. In addition, the software will identify individual people who are expected to become – but have yet to be – victims or perpetrators of violent crimes. Officers may even be assigned to visit those people to warn them against committing a violent crime.

Any attempt to curb the alarming rate of <u>homicides in Chicago</u> is laudable. But the city's new effort seems to ignore evidence, including recent research from members of our policing study team at the <u>Human</u> <u>Rights Data Analysis Group</u>, that predictive policing tools reinforce, rather than reimagine, existing police practices. Their expanded use could lead to further targeting of communities or people of color.

Working with available data

At its core, any predictive model or algorithm is a combination of data and a statistical process that seeks to identify patterns in the numbers. This can include looking at police data in hopes of learning about crime trends or recidivism. But a useful outcome depends not only on good mathematical analysis: It also needs good data. That's where predictive policing often falls short.

Machine-learning algorithms learn to make predictions by analyzing patterns in an initial training data set and then look for similar patterns in new data as they come in. If they learn the wrong signals from the data, the subsequent analysis will be lacking.

This happened with a Google initiative called "Flu Trends," which was



launched in 2008 in hopes of using information about people's online searches to spot disease outbreaks. Google's systems would monitor users' searches and identify locations where many people were researching various flu symptoms. In those places, <u>the program would</u> <u>alert public health authorities</u> that more people were about to come down with the flu.

But the project failed to account for the potential for periodic changes in Google's own search algorithm. In an early 2012 update, Google modified its search tool to suggest a diagnosis when users searched for terms like "cough" or "fever." On its own, this change <u>increased the number of searches for flu-related terms</u>. But Google Flu Trends interpreted the data as predicting a flu outbreak twice as big as federal <u>public health</u> officials expected and far <u>larger than what actually happened</u>.

Criminal justice data are biased

The failure of the Google Flu Trends system was a result of one kind of flawed data – information biased by factors other than what was being measured. It's much harder to identify bias in criminal justice prediction models. In part, this is because <u>police data aren't collected uniformly</u>, and in part it's because what data police track reflect longstanding institutional biases along <u>income</u>, <u>race</u> and gender lines.

While police data often are described as representing "crime," that's not quite accurate. Crime itself is a largely hidden social phenomenon that happens anywhere a person violates a law. What are called "crime data" usually tabulate specific events that aren't necessarily lawbreaking – like a 911 call – or that are influenced by existing police priorities, like arrests of people suspected of particular types of crime, or reports of incidents seen when <u>patrolling a particular neighborhood</u>.



Neighborhoods with lots of police calls aren't necessarily the same places the most crime is happening. They are, rather, where the most police attention is – though where that attention focuses can often be <u>biased by</u> <u>gender</u> and <u>racial factors</u>.

It's not possible to remove the bias

Some researchers have argued that machine learning algorithms can address systemic biases by designing "neutral" models that don't take into account sensitive variables like <u>race or gender</u>. But while it may seem possible in hypothetical situations, it doesn't appear to be the case in real life.

Our recent study, by Human Rights Data Analysis Group's Kristian Lum and William Isaac, found that <u>predictive policing vendor PredPol's</u> <u>purportedly race-neutral algorithm</u> targeted black neighborhoods at roughly twice the rate of white neighborhoods when trained on historical drug crime data from Oakland, California. We found similar results when analyzing the data by income group, with low-income communities targeted at disproportionately higher rates compared to high-income neighborhoods.

But estimates – created from <u>public health surveys</u> and <u>population</u> <u>models</u> – suggest illicit drug use in Oakland is roughly equal across racial and income groups. If the algorithm were truly race-neutral, it would spread drug-fighting police attention evenly across the city.

Similar evidence of racial bias was found by ProPublica's investigative reporters <u>when they looked at COMPAS</u>, an algorithm predicting a <u>person's risk of committing a crime</u>, used in bail and sentencing decisions in Broward County, Florida, and elsewhere around the country. These systems learn only what they are presented with; if those data are biased, their learning can't help but be biased too.



Fixing this problem is not a matter of just doing more advanced mathematical or statistical calculations. Rather, it will require rethinking how police agencies collect and analyze data, and how they train their staff to use data on the job.

Understanding the biases to improve the data

Using predictive analytics in the real world is challenging, particularly when trying to craft government policies to minimize harm to vulnerable populations. We do not believe that police departments should stop using analytics or data-driven approaches to reducing crime. Rather, police should work to understand the biases and limitations inherent in their data.

In our view, police departments – and all agencies that use <u>predictive</u> <u>algorithms</u> – should make their systems transparent to public scrutiny. This should start with community members and <u>police departments</u> discussing policing priorities and measures of police performance. That way any software the police use can be programmed to reflect the community's values and concerns.

Ensuring transparency

It is not enough to claim or assume an algorithm is unbiased just because it is computerized and uses data: A lack of bias must be proven by evaluating the algorithm's performance itself. Police agencies should get independent experts or human rights groups to perform regular audits of the algorithms and the data they process. Much like the annual financial reviews large companies do, these examinations can ensure the input data are valid and are analyzed properly to avoid discrimination. If a company wants to claim its algorithm is proprietary and should be kept secret, it should still be required to offer robust testing environments so



outside experts can examine its performance.

Further, <u>police</u> departments that use algorithms to make predictions about individuals, like <u>Chicago's Strategic Subject List does</u>, should have policies similar to a <u>new European Union regulation</u> requiring humanunderstandable explanations of computer algorithms' decisions. And no agency or company should be allowed to discriminate against people who have been identified by predictive policing.

Used correctly, predictive policing can be used to address the complex factors underlying <u>crime</u> trends. For example, rather than stepping up patrols, <u>Toronto and other cities in Canada</u> are using predictive modeling to connect residents to local social services. By improving the quality of <u>data</u> cities collect, and analyzing the information with more transparent and inclusive processes, cities can build safer communities, rather than cracking down harder on areas that are already struggling.

This article was originally published on <u>The Conversation</u>. Read the <u>original article</u>.

Provided by The Conversation

Citation: Why big-data analysis of police activity is inherently biased (2017, May 10) retrieved 27 April 2024 from https://phys.org/news/2017-05-big-data-analysis-police-inherently-biased.html

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.