

Ava Weisheit

Younai Soe

ILZ-Z 410

April 29, 2026

Coded Prejudice: How Predictive Policing Turns Bias Into Law

Imagine this: two men walk out of the same Walmart having stolen the same item, at the same time, with identical criminal records. One is white. One is Black. They are arrested together, tried together, and stand before the same judge. But when sentencing comes, the Black man receives nearly twice the prison time - not because of anything he did differently, but because an algorithm decided he was a higher risk to society. No explanation. No appeal. Just a number generated by a machine, built on data that was never neutral to begin with.

How does that make you feel? Would you accept that outcome if it were you?

America promised us that data-driven policing would make the justice system fairer and smarter. Instead, it has handed racist outcomes a scientific disguise - and the communities bearing the cost are the same ones that have always been over-policed, over-surveilled, and under-protected. Predictive policing does not predict crime. It predicts who society has already decided to criminalize, and until that truth is confronted, no algorithm will make our justice system just.

Predictive policing uses algorithms to analyze existing crime data to forecast what crimes will occur, where they will occur, and who is most likely to commit them (Browning & Arrigo, 2021). One of its most widely used tools, COMPAS, assigns individuals a numerical risk score between one and ten to determine how likely they are to reoffend - a number that judges across the country have cited directly in sentencing decisions (Heaven, 2020). The promise was simple: because these tools are built on math rather than human judgment, they would be free from the prejudices that have long plagued law enforcement (Browning & Arrigo, 2021). But this promise rested on a deeply flawed assumption - that the data itself was clean.

It was not. The data these algorithms learn from is not a neutral record of criminal behavior. It is a record of who has been policed, arrested, and processed by a system with a centuries-long history of racial discrimination (Browning & Arrigo, 2021). As Browning and Arrigo (2021) argue, predictive policing is "nothing more than a digital analogue to traditional stop and frisk practices" - a policy already ruled unconstitutional for its discriminatory targeting of Black and Hispanic communities. We banned the human version. We are now paying private companies to run the algorithmic one. Hung and Yen (2023) make the idea clear: when the data that police provide already contain contextual priorities, filtering, and decisions, the results will also reflect those same assumptions. Biased data does not produce unbiased predictions. It produces bias at scale, dressed in the language of an AI algorithm.

The real-world consequences are not theoretical. Angwin et al. (2016) found in their landmark ProPublica investigation that Black defendants were nearly twice as likely as white defendants to be falsely flagged as future offenders by the COMPAS algorithm, even after controlling for criminal history, age, and gender. In Chicago, the city's predictive "heat list" grew to over 400,000 people - with 56 percent of Black men between the ages of 20 and 29 receiving a risk score (Hung & Yen, 2023). In Plainfield, New Jersey, the predictive software Geolitica was found to have a success rate of less than half a percent across more than 23,000 predictions (Thomas, 2024). These tools are not just biased. They do not even work.

Supporters argue that algorithms are still more consistent than biased human judges, and that some studies show modest reductions in property crime in targeted areas. Friend (n.d.), writing for the FBI Law Enforcement Bulletin, points to experiments in Santa Cruz and Los Angeles where predictive algorithms produced twice the accuracy of traditional policing methods, with property crime in one LAPD division declining by 12 percent during the trial, concluding that "the computer eliminates the bias that people have." This is the strongest version of the counterargument, and it deserves a direct response: consistency is not justice, and crime reduction in one neighborhood tells us nothing about who paid the price for it elsewhere. An algorithm that is consistently wrong about Black defendants in the same direction is not an improvement over human bias - it is human bias with better public relations. As Hung and Yen (2023) demonstrate, it is mathematically impossible to satisfy all definitions of algorithmic fairness simultaneously. The companies selling these tools simply choose the metric that makes

their product look unbiased, while the measure that reveals racial harm is quietly set aside. Meanwhile, Heaven (2020) notes that these algorithms do not even use race directly - they use zip code, education level, and employment status as approximations, variables so correlated with race in America that the distinction barely matters.

Beyond racial injustice, predictive policing threatens the constitutional rights of every citizen. The Fourth Amendment protects against unreasonable searches and seizures, and the Fourteenth guarantees equal protection under the law. Browning and Arrigo (2021) argue that by directing officers toward minority communities based on biased data, and subjecting individuals to surveillance without probable cause, predictive policing violates both. Research also shows that living under constant algorithmic surveillance causes measurable psychological harm, eroding the sense of safety and belonging essential to participation in a democratic society (Browning & Arrigo, 2021).

Our justice system was already broken before predictive policing arrived. Adding an algorithm does not fix a broken foundation - it just automates the cracks. These tools have over-promised and under-delivered, and as Thomas (2024) notes, communities of color have carried the brunt of their failures. The solution is not better data or more oversight. The solution is to stop pretending that a just outcome can be computed from an unjust history. Until we reckon with the racism embedded in every data point these systems feed on, predictive policing will remain exactly what it has always been - discrimination with a digital disguise.

References

- Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2016, May 23). *Machine bias: There's software used across the country to predict future criminals. And it's biased against blacks*. ProPublica. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- Browning, M., & Arrigo, B. (2021). Stop and risk: Policing, data, and the digital age of discrimination. *American Journal of Criminal Justice*, 46(2), 298–316. <https://doi.org/10.1007/s12103-020-09557-x>
- Friend, Z. (n.d.). Predictive policing: Using technology to reduce crime. *FBI Law Enforcement Bulletin*. <https://leb.fbi.gov/articles/featured-articles/predictive-policing-using-technology-to-reduce-crime>
- 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/>
- Hung, T., & Yen, C. (2023). Predictive policing and algorithmic fairness. *Synthese*, 201, 206. <https://doi.org/10.1007/s11229-023-04189-0>
- Thomas, G. (2024, October 15). Politicians move to limit predictive policing after years of controversial failures. *TechPolicy Press*. <https://www.techpolicy.press/politicians-move-to-limit-predictive-policing-after-years-of-controversial-failures/>