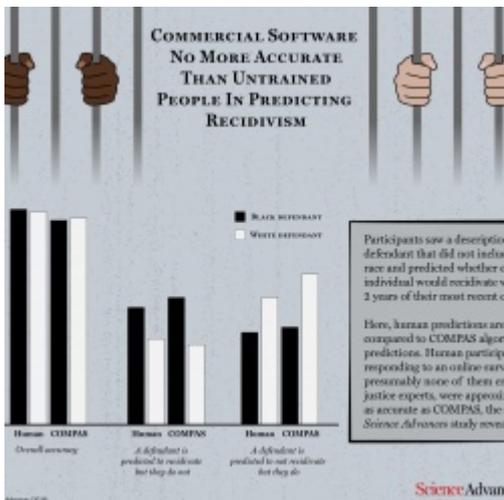


ProPublica Is Wrong In Charging Racial Bias In An Algorithm



By Chuck Dinerstein — January 17, 2018



Credit: Carla Schaffer / AAAS

Predicting the future is not only the provenance of fortune tellers or media pundits. Predictive algorithms, based on extensive datasets and statistics have overtaken wholesale and retail operations as any online shopper knows. And in the last few years algorithms, are used to automate decision making for bank loans, school admissions, hiring and infamously in predicting recidivism – the probability that a defendant will commit another crime in the next two years.

COMPAS, which stands for Correctional Offender Management Profiling for Alternative Sanctions, is such a program and was singled out by [ProPublica](#) [1] earlier this year as being racially biased. COMPAS utilizes 137 variables in its proprietary and unpublished scoring algorithm; race is not one of those variables. ProPublica used a dataset of defendants in Broward County, Florida. The data included demographics, criminal history, a COMPAS score [1] and the criminal actions in the subsequent two years. ProPublica then crosslinked this data with the defendants' race. Their findings are generally accepted by all sides

- COMPAS is moderately accurate in identifying white and black recidivism about 60% of the time.
- COMPAS's errors reflect apparent racial bias. Blacks are more often wrongly identified as recidivist risks (statistically a false positive) and whites more often erroneously identified as not being a risk (a false negative).
- The “penalty” for being misclassified as a higher risk is more likely to be stiffer punishment. Being misclassified as a lower risk is like a “Get out of jail” card.

As you might anticipate, significant controversy followed couched mostly in an academic fight over

which statistical measures or calculations more accurately depicted the bias. A study in Science Advances revisits the discussion and comes to a different conclusion.

The study

In the current study, assessment by humans was compared to that of COMPAS using that Broward County dataset. The humans were found on Amazon’s Mechanical Turk [2], paid a dollar to participate and a \$5 bonus if they were accurate more than 60% of the time.

- Humans were just as accurate as the algorithm (62 vs. 65%)
- The errors by the algorithm and humans were identical, overpredicting (false positives) recidivism for black defendants and underpredicting (false negatives) for white defendants.

	Humans		COMPAS	
		Black %	White %	Black %
Accuracy*	68.2	67.6	64.9	65.7
False Positives	37.1	27.2	40.4	25.4
False Negatives	29.2	40.3	30.9	47.9

* Accuracy is the sum of true positives and true negatives statistically speaking

The human assessors used only seven variables, not the 137 of COMPAS. [3] This suggests that the algorithm was needlessly complex at least in deciding recidivism risk. In fact, the researchers found that just two variables, defendant age, and the number of prior convictions was as accurate as COMPAS’s predictions.

Of more significant interest is the finding that when human assessors were given the additional information regarding defendant race it had no impact. They were just as accurate and demonstrated the same racial disparity in false positives and negatives. Race was an associated confounder, but it was not the cause of the statistical difference. ProPublica’s narrative of racial bias was incorrect.

Algorithms are statistical models involving choices. If you optimize to find all the true positives, your false positives will increase. Lower your false positive rate and the false negatives increase. Do we want to incarcerate more or less? The [MIT Technology Review](#) [2] puts it this way.

“Are we primarily interested in taking as few chances as possible that someone will skip bail or re-offend? What trade-offs should we make to ensure justice and lower the massive social costs of imprisonment?”

COMPAS is meant to serve as a decision aid. The purpose of the 137 variables is to create a

variety of scales depicting substance abuse, environment, criminal opportunity, associates, etc. Its role is to assist the humans of our justice system in determining an appropriate punishment. [4] None of the studies, to my knowledge, looked at the sentences handed down. The current research ends as follows:

“When considering using software such as COMPAS in making decisions that will significantly affect the lives and well-being of criminal defendants, it is valuable to ask whether we would put these decisions in the hands of random people who respond to an online survey because, in the end, the results from these two approaches appear to be indistinguishable.”

The answer is no. COMPAS and similar algorithms are tools, not a replacement FOR human judgment. They facilitate but do not automate. ProPublica is correct when saying algorithmic decisions need to be understood by their human users and require continuous validation and refinement. But ProPublica’s narrative, that evil forces were responsible for a racially biased algorithm are not true.

[1] COMPAS is scored on a 1 to 10 scale, with scores greater than 8 being at high risk for repeat offenders

[2] Amazon's Mechanical Turk is a crowd-sourced marketplace for human assistance. In this case for people to participate in a survey.

[3] The variables were gender, age, criminal charge, misdemeanor or felony, prior adult convictions and prior juvenile felony and misdemeanor charges.

[4] As the COMPAS manual states: “Criminal justice agencies across the nation use COMPAS **to inform decisions** regarding the placement, supervision and case management of offenders.” [Emphasis added]

Source: [The accuracy, fairness and limits of predicting recidivism](#) [3] Science Advances 2018

COPYRIGHT © 1978-2016 BY THE AMERICAN COUNCIL ON SCIENCE AND HEALTH

Source URL: <https://www.acsh.org/news/2018/01/17/propublica-wrong-charging-racial-bias-algorithm-12427>

Links

[1] <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>

[2] <https://www.technologyreview.com/s/607955/inspecting-algorithms-for-bias/>

[3] <http://advances.sciencemag.org/content/4/1/eaao5580>