


# Algorithms can be more fair than humans

August 18 2016, by H V Jagadish

---



Credit: AI-generated image ([disclaimer](#))

Amazon recently began to offer [same-day delivery](#) in selected metropolitan areas. This may be good for many customers, but the rollout shows how computerized decision-making can also deliver a strong dose of discrimination.

Sensibly, the company began its service in areas where delivery costs would be lowest, by identifying ZIP codes of densely populated places

home to many existing Amazon customers with income levels high enough to make frequent purchases of products available for same-day delivery. The company provided a web page letting customers enter their ZIP code to see if same-day delivery served them. Investigative journalists at Bloomberg News used that page to [create maps of Amazon's service area for same-day delivery](#).

The Bloomberg analysis revealed that many poor urban areas were excluded from the service area, while more affluent neighboring areas were included. Many of these excluded poor areas were predominantly inhabited by minorities. For example, all of Boston was covered except for Roxbury; New York City coverage included almost all of four boroughs but completely excluded the Bronx; Chicago coverage left out the impoverished South Side, while extending substantially to affluent northern and western suburbs.

While it is tempting to believe data-driven decisions are unbiased, [research and scholarly discussion](#) are beginning to demonstrate that unfairness and discrimination remain. In my [online course on data ethics](#), students learn that [algorithms can discriminate](#). But there may be a bit of a silver lining: As the Bloomberg research suggests, basing decisions on data may also make it easier to detect when biases arise.

## **Bias can be unintentional**

Unfairness like that in Amazon's delivery policy can arise for many reasons, including hidden biases – such as assumptions that populations are distributed uniformly. Algorithm designers likely don't intend to discriminate, and may not even realize a problem has crept in.

Amazon told Bloomberg it had no discriminatory intent, and there is every reason to believe that claim. In response to the Bloomberg report, [city officials](#) and [other politicians](#) called on Amazon to fix this problem.

The company [moved quickly to add](#) the originally excluded poor urban ZIP codes to its service area.

A similar question has been [asked of Uber](#), which seems to provide better service to areas inhabited by higher proportions of white people. It is likely there will be more retail and service industry examples of unintentional algorithmic discrimination discovered in the future.

## **Asking too much of algorithms?**

We should pause a moment to consider whether we are unduly demanding of algorithmic decisions. Companies operating brick-and-mortar stores make location decisions all the time, taking into account criteria not that different from Amazon's. Stores attempt to have locations that are convenient for a large pool of potential customers with money to spend.

In consequence, few stores choose to locate in poor inner-city neighborhoods. Particularly in the context of grocery stores, this phenomenon has been studied extensively, and the term "[food desert](#)" has been used to describe urban areas whose residents have no convenient access to fresh food. This [location bias](#) is less studied for retail stores overall.

As an indicative example, I looked at the 55 Michigan locations of Target, a large comprehensive retail chain. When I sorted every Michigan ZIP code based on whether its average income was in the top half or bottom half statewide, I found that only 16 of the Target stores (29 percent) were in ZIP codes from the lower income group. More than twice as many, 39 stores, were sited in ZIP codes from the more affluent half.

## Identifying discrimination

Moreover, there are no Target stores in the city of Detroit, though there are several in its (wealthier) suburbs. Yet there has been no popular outcry alleging Target unfairly discriminates against poor people in its store location decisions. There are two main reasons the concerns about Amazon are justified: rigidity and dominance.

Rigidity has to do with both the online retailer's decision-making processes and with the result. Amazon decides which ZIP codes are in its service area. If a customer lives just across the street from the boundary set by Amazon, she is outside the service area and can do little about it. By contrast, someone who lives in a ZIP code without a Target store can still shop at Target – though it may take longer to get there.

It also matters how dominant a retailer is in consumers' minds. Whereas Target is only one of many physical store chains, Amazon enjoys [market dominance as a web retailer](#), and hence attracts more attention. Such dominance is a characteristic of today's [winner-takes-all](#) web businesses.

While their rigidity and dominance may cause us greater concern about online businesses, we also are better able to detect their discrimination than we are for brick-and-mortar shops. For a traditional chain store, we need to guess how far consumers are willing to travel. We may also need to be cognizant of time: Five miles to the next freeway exit is not the same thing as five miles via congested streets to the other side of town. Furthermore, travel time itself can vary widely depending on the time of day. After identifying the likely areas a store serves, they may not map neatly into geographic units for which we have statistics about race or income. In short, the analysis is messy and requires much effort.

In contrast, it would have taken journalists at Bloomberg only a few hours to develop a map of Amazon's service area and correlate it with

income or race. If Amazon had done this internally, they could have performed the same analysis in just minutes – and perhaps noticed the problems and fixed them before same-day service even began.

## How do humans compare?

Let us look at a very different example to see how the same points apply broadly. Recently, ProPublica published [an excellent analysis of racial discrimination](#) by an algorithm that predicts a criminal's likelihood of offending again. The algorithm considers dozens of factors and calculates a probability estimate. ProPublica's analysis found significant systematic racial bias, even though race was not among the specific factors considered.

Without the algorithm, a human judge would make a similar estimate, as part of a sentencing or parole decision. The human decision might consider a richer set of factors, such as the criminal's courtroom demeanor. But we know, from [studies in psychology](#), that [human decision-making is replete with bias](#), even when we try our best to be fair.

But any errors that result from bias in human judges' decisions are likely to be different among judges, and even for different decisions made by the same judge. In the aggregate, there may be racial discrimination due to [subconscious bias](#), but establishing this conclusively is tricky. A U.S. Justice Department study found strong evidence of [disparities in sentencing white and black convicts](#), but could not clearly determine whether race itself was a factor in those decisions.

In contrast, the exact same algorithm ProPublica looked at is used in thousands of cases across many states. Its rigidity, and the large volume, ease the job of determining if it discriminates – and can offer ways to efficiently rectify the problem.

The use of information technology seems to make lines brighter, differences starker and data about all of this much more easily available. What could be brushed under the rug yesterday now clamors for attention. As we find more and more uses for data-driven algorithms, it is not yet common to analyze their fairness, particularly before the roll out of a new data-based service. Making it so will go a long way to measuring, and improving, the fairness of these increasingly important computerized calculations.

*This article was originally published on [The Conversation](#). Read the [original article](#).*

Provided by The Conversation

Citation: Algorithms can be more fair than humans (2016, August 18) retrieved 24 April 2024 from <https://techxplore.com/news/2016-08-algorithms-fair-humans.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.