

AI use may mask racial disparities in credit, lending, study suggests

January 30 2019, by Melanie Lefkowitz



Credit: Cornell University

By law, credit and loan decisions cannot discriminate on the basis of race or lead to outcomes that differ substantially by race. But to ensure that they don't discriminate, banks and other lenders aren't allowed to ask about race on most applications. This makes it challenging for auditors to make sure credit decisions are fair.

To evaluate racial disparities in lending decisions, lenders or auditors have to infer applicants' races, generally using a system – known as a proxy – that guesses applicants' races based on what they do know, such as their neighborhoods and surnames.

But these proxies – including a method used by the Consumer Financial Protection Bureau to audit lenders – can yield very different results depending on tiny changes in how they guess applicants' races, according to a new Cornell-led study.

"It's worrisome that these models are being used to determine whether financial institutions comply with the law," said Madeleine Udell, the Richard and Sybil Smith Sesquicentennial Fellow and assistant professor in the School of Operations Research and Information Engineering. "They're clearly not assessing what they're supposed to."

Their paper, "Fairness Under Unawareness: Assessing Disparity When Protected Class Is Unobserved," will be presented at the ACM Conference on Fairness, Accountability and Transparency, Jan. 29-31 in Atlanta. Cornell Tech doctoral student Xiaojie Mao is the lead author. Co-authors included Udell; Nathan Kallus, assistant professor of operations research and information engineering at Cornell Tech; and financial industry data scientists Jiahao Chen and Geoffry Svacha.

Understanding the risks of discrimination when using artificial intelligence is especially important as [financial institutions](#) increasingly rely on [machine learning](#) for lending decisions. Machine learning models can analyze reams of data to arrive at relatively accurate predictions, but their operations are opaque, making it difficult to ensure fairness.

"How can a computer be racist if you're not inputting race? Well, it can, and one of the biggest challenges we're going to face in the coming years is humans using machine learning with unintentional bad consequences

that might lead us to increased polarization and inequality," Kallus said. "There have been a lot of advances in machine learning and [artificial intelligence](#), and we have to be really responsible in our use of it."

Race is one of several characteristics protected by state and federal law; others include age, gender and disability status.

The researchers used data from mortgages – the one type of consumer loan that includes race on applications – to test the accuracy of the Bayesian Improved Surname Geocoding (BISG) auditing system. They found its results often either underestimated or overestimated racial discrepancies, depending on several factors. Assuming race based on the census tracts where applicants live erases black applicants who live in mostly white neighborhoods and white applicants who live in mostly black neighborhoods.

The BISG model estimates the probability that someone is a certain race, and in performing calculations a user can set a minimum probability – for example, choosing to use any examples in which the probability of a given race is 80 percent or more. But differences in that minimum probability yielded unexpectedly large variations in the results, the researchers found.

"Depending on what threshold you picked, you would get wildly different answers for how fair your credit procedure was," Udell said.

The researchers' findings not only shed light on BISG's accuracy, they could help developers improve the machine learning models that make credit decisions. Better models could help banks make more informed decisions when they approve or reject loans, which may lead them to give credit to qualified but lower-income applicants.

"You can figure out who will actually default or not in ways that are

fair," Kallus said. "What we want to do is make sure we put these constraints on the machine learning systems that we build and train, so we understand what it means to be fair and how we can make sure it's fair from the outset."

More information: Jiahao Chen et al. Fairness Under Unawareness, *Proceedings of the Conference on Fairness, Accountability, and Transparency - FAT* '19* (2019). [DOI: 10.1145/3287560.3287594](https://doi.org/10.1145/3287560.3287594)

Provided by Cornell University

Citation: AI use may mask racial disparities in credit, lending, study suggests (2019, January 30) retrieved 9 April 2024 from

<https://techxplore.com/news/2019-01-ai-mask-racial-disparities-credit.html>

<p>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.</p>
--