

Algorithm improves fairness of search results

August 18 2020, by Melanie Lefkowitz



Credit: CC0 Public Domain

When you search for something on the internet, do you scroll through page after page of suggestions—or pick from the first few choices?

Because most people choose from the tops of these lists, they rarely see



the vast majority of the options, creating a potential for bias in everything from hiring to <u>media exposure</u> to e-commerce.

In a new paper, Cornell researchers introduce a tool they've developed to improve the fairness of online rankings without sacrificing their usefulness or relevance.

"If you could examine all your choices equally and then decide what to pick, that may be considered ideal. But since we can't do that, rankings become a crucial interface to navigate these choices," said <u>computer</u> science doctoral student Ashudeep Singh, co-first author of "Controlling Fairness and Bias in Dynamic Learning-to-Rank," which won the Best Paper Award at the Association for Computing Machinery SIGIR Conference on Research and Development in Information Retrieval, held virtually July 25-30.

"For example, many YouTubers will post videos of the same recipe, but some of them get seen way more than others, even though they might be very similar," Singh said. "And this happens because of the way search results are presented to us. We generally go down the ranking linearly and our attention drops off fast."

The researchers' method, called FairCo, gives roughly equal exposure to equally relevant choices and avoids preferential treatment for items that are already high on the list. This can correct the unfairness inherent in existing algorithms, which can exacerbate inequality and political polarization, and curtail personal choice.

"What ranking systems do is they allocate exposure. So how do we make sure that everybody receives their fair share of exposure?" said Thorsten Joachims, professor of computer science and information science, and the paper's senior author. "What constitutes fairness is probably very different in, say, an e-commerce system and a system that ranks resumes



for a job opening. We came up with computational tools that let you specify fairness criteria, as well as the algorithm that will provably enforce them."

Online <u>ranking systems</u> were originally based on library science from the 1960s and '70s, which sought to make it easier for users to find the books they wanted. But this approach can be unfair in two-sided markets, in which one entity wants to find something and another wants to be found.

"Much of machine learning work in optimizing rankings is still very much focused on maximizing utility to the users," Joachims said. "What we've done over the last few years is come up with notions of how to maximize utility while still being fair to the items that are being searched."

Algorithms that prioritize more popular items can be unfair because the higher a choice appears in the list, the more likely users are to click on and react to it. This creates a "rich get richer" phenomenon where one choice becomes increasingly popular, and other choices go unseen.

Algorithms also seek the most relevant items to searchers, but because the vast majority of people choose one of the first few items in a list, small differences in relevance can lead to huge discrepancies in exposure. For example, if 51% of the readers of a news publication prefer opinion pieces that skew conservative, and 49% prefer essays that are more liberal, all of the top stories highlighted on the home page could conceivably lean conservative, according to the paper.

"When small differences in relevance lead to one side being amplified, that often causes polarization, where some people tend to dominate the conversation and other opinions get dropped without their fair share of attention," Joachims said. "You might want to use it in an <u>e-commerce</u>



system to make sure that if you're producing a product that 30% of people like, you're getting a certain amount of exposure based on that. Or if you have a resume database, you could formulate safeguards to make sure it's not discriminating by race or gender."

More information: Marco Morik et al. Controlling Fairness and Bias in Dynamic Learning-to-Rank, *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval* (2020). DOI: 10.1145/3397271.3401100

Provided by Cornell University

Citation: Algorithm improves fairness of search results (2020, August 18) retrieved 26 April 2024 from <u>https://techxplore.com/news/2020-08-algorithm-fairness-results.html</u>

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.