Tulane researcher gets NSF grant to teach algorithms to be fair

November 1st, 2021

Barri Bronston bbronst@tulane.edu

View PDF



People who shop online may not always get the fairest of recommendations, but Tulane researcher Nick Mattei hopes a new study on fairness in algorithms leads to more equitable internet experiences.

Anyone who's spent much time on social media knows that the algorithms are always watching. These machine learning systems know what we like, what we consume online and what we may want to buy — sometimes even before we do based on the digital footprints we all leave online. But these highly curated and personalized recommender systems aren't always fair and can sometimes inherit prejudices and biases from the data, search design or programmers of the systems, said Tulane artificial intelligence expert <u>Nick Mattei</u>.

"These systems are central to many of our experiences on the Internet," said Mattei, an assistant professor of computer science in the Tulane School of Science and Engineering. "Imagine a system that always recommended the most expensive products in a category – or only recommended news from a certain set of outlets. These systems can unfairly make options invisible to you on the internet."

Mattei is part of a new study funded by the National Science Foundation to design more equitable algorithm recommender systems that can be applied generally to many organizations no matter what types of products or services they are recommending to users.

"The big problem is that fairness can and is defined in many ways by many different stakeholders. The question is how to balance all these competing concerns in these complex settings." *Nick Mattei, assistant professor of computer science at Tulane*

He is teaming up on the \$930,000 study with Robin Burke and Amy Voida, professors of information science at the University of Colorado, and Kiva.org, a nonprofit lending institution geared to underserved communities around the world. Kiva.org will serve as a co-principal investigator, providing data and the ability to test the new algorithms in a live setting.

"A key piece of the research is understanding how fairness concepts are understood and used in real-world settings," Mattei said. "Kiva.org is a micro-lending site that has an overall goal of promoting development through its platform."

Fairness in recommendation systems has been the subject of recent research and media attention, and for good reason, Mattei said. For example, job recommendation systems may leverage data that excludes minorities or other historically underrepresented groups from certain opportunities. Individuals whose preferences are very different from the average user may not receive high-quality recommendations. Coming up with solutions has been difficult because fairness has too often been conceived in simple, narrow ways, and has remained largely divorced from realworld organizational practices.

"The big problem is that fairness can and is defined in many ways by many different stakeholders: the operators of the platform, the people receiving the loans and the users of the website. The question is how to balance all these competing concerns in these complex settings," Mattei said.

Kiva.org's challenge, for example, is in its wide range of borrowers from different sectors of the economy and the world.

"Sometimes being fair to one group may mean being very unfair to another group," Mattei said. "If for some reason all loans to one country are only women, then there will be no men in the search results, so optimizing for one type of fairness may mean moving away from fairness across another dimension."

To help counter this, Mattei will use mathematical tools from computational social choice, fair allocation and algorithmic game theory to create new systems for deploying and understanding multi-stakeholder fairness in the context of recommender systems.

"My piece of the research is using concepts from multi-agent systems including computational social choice and fair allocation to create new systems for deploying multi-stakeholder fairness in recommender systems," Mattei said.

Researchers will conduct a detailed analysis of fairness within Kiva.org, ensuring that the fairness concepts that are implemented in the system are grounded in real organizational needs. Working with Kiva.org, researchers will conduct interviews and focus groups with diverse stakeholders, building models of the different ways that fairness is used within the organizational context, and generalizing these techniques to apply to other organizations.