


Study: When allocating scarce resources with AI, randomization can improve fairness

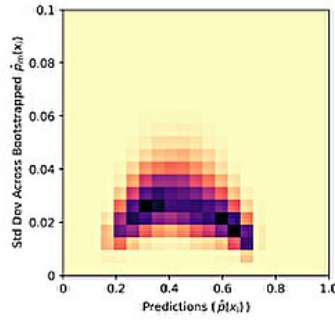
July 24 2024, by Adam Zewe

(a) - (c) Swiss Unemployment Data – Predictions x Std Dev Across Bootstrapped Predictions


(a) Logistic Regression



(b) Random Forest




(c) Decision Tree




(d) - (f) Swiss Unemployment Data – Predictions x Conformal P-Values for Outlier Detection


(d) Logistic Regression



(e) Random Forest

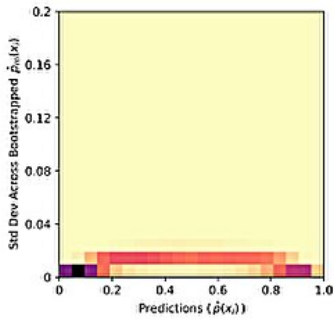


(f) Decision Tree




(g) - (i) Census Income Data – Predictions x Std Dev Across Bootstrapped Predictions


(g) Logistic Regression



(h) Random Forest




(i) Decision Tree




(j) - (l) Census Income Data – Predictions x Conformal P-Values for Outlier Detection


(j) Logistic Regression



(k) Random Forest



(l) Decision Tree



Density of Predictions by Point Estimate $\hat{p}(x_i)$ & Uncertainty Metric (Darker Colors = Higher Density). Credit: *arXiv* (2024). DOI: 10.48550/arxiv.2404.08592

Organizations are increasingly utilizing machine-learning models to allocate scarce resources or opportunities. For instance, such models can help companies screen resumes to choose job interview candidates or aid hospitals in ranking kidney transplant patients based on their likelihood of survival.

When deploying a model, users typically strive to ensure its predictions are fair by reducing bias. This often involves techniques like adjusting the features a model uses to make decisions or calibrating the scores it generates.

However, researchers from MIT and Northeastern University argue that these fairness methods are not sufficient to address structural injustices and inherent uncertainties. In a [new paper](#) posted to the *arXiv* preprint server, they show how randomizing a model's decisions in a structured way can improve fairness in certain situations.

For example, if multiple companies use the same machine-learning model to rank job interview candidates deterministically—without any randomization—then one deserving individual could be the bottom-ranked candidate for every job, perhaps due to how the model weighs answers provided in an online form. Introducing randomization into a model's decisions could prevent one worthy person or group from always being denied a scarce resource, like a job interview.

Through their analysis, the researchers found that randomization can be especially beneficial when a model's decisions involve uncertainty or when the same group consistently receives negative decisions.

They present a framework one could use to introduce a specific amount of randomization into a model's decisions by allocating resources through a weighted lottery. This method, which an individual can tailor to fit their situation, can improve fairness without hurting the efficiency or accuracy of a model.

"Even if you could make fair predictions, should you be deciding these social allocations of scarce resources or opportunities strictly off scores or rankings? As things scale, and we see more and more opportunities being decided by these algorithms, the inherent uncertainties in these scores can be amplified. We show that fairness may require some sort of randomization," says Shomik Jain, a graduate student in the Institute for Data, Systems, and Society (IDSS) and lead author of the paper.

Jain is joined on the paper by Kathleen Creel, assistant professor of philosophy and computer science at Northeastern University; and senior author Ashia Wilson, the Lister Brothers Career Development Professor in the Department of Electrical Engineering and Computer Science and a principal investigator in the Laboratory for Information and Decision Systems (LIDS). The research will be presented at the International Conference on Machine Learning ([ICML 2024](#)), held in Vienna, Austria, July 21–27.

Considering claims

This work builds off a [previous paper](#) in which the researchers explored harms that can occur when one uses deterministic systems at scale. They found that using a machine-learning model to deterministically allocate resources can amplify inequalities that exist in training data, which can

reinforce bias and systemic inequality.

"Randomization is a very useful concept in statistics, and to our delight, satisfies the fairness demands coming from both a systemic and individual point of view," Wilson says.

In this paper, they explored the question of when randomization can improve fairness. They framed their analysis around the ideas of philosopher John Broome, who wrote about the value of using lotteries to award scarce resources in a way that honors all claims of individuals.

A person's claim to a scarce resource, like a kidney transplant, can stem from merit, deservingness, or need. For instance, everyone has a right to life, and their claims on a [kidney transplant](#) may stem from that right, Wilson explains.

"When you acknowledge that people have different claims to these scarce resources, fairness is going to require that we respect all claims of individuals. If we always give someone with a stronger claim the resource, is that fair?" Jain says.

That sort of deterministic allocation could cause systemic exclusion or exacerbate patterned inequality, which occurs when receiving one allocation increases an individual's likelihood of receiving future allocations. In addition, [machine-learning models](#) can make mistakes, and a deterministic approach could cause the same mistake to be repeated.

Randomization can overcome these problems, but that doesn't mean all decisions a model makes should be randomized equally.

Structured randomization

The researchers used a weighted lottery to adjust the level of randomization based on the amount of uncertainty involved in the model's decision-making. A decision that is less certain should incorporate more randomization.

"In kidney allocation, usually the planning is around projected lifespan, and that is deeply uncertain. If two patients are only five years apart, it becomes a lot harder to measure. We want to leverage that level of uncertainty to tailor the randomization," Wilson says.

The researchers used statistical uncertainty quantification methods to determine how much randomization is needed in different situations. They show that calibrated randomization can lead to fairer outcomes for individuals without significantly affecting the utility, or effectiveness, of the [model](#).

"There is a balance to be had between overall utility and respecting the rights of the individuals who are receiving a scarce resource, but oftentimes the tradeoff is relatively small," says Wilson.

However, the researchers emphasize there are situations where randomizing decisions would not improve fairness and could harm individuals, such as in criminal justice contexts.

But there could be other areas where randomization can improve fairness, such as [college admissions](#), and the researchers plan to study other use cases in future work. They also want to explore how randomization can affect other factors, such as competition or prices, and how it could be used to improve the robustness of machine-learning models.

"We are hoping our paper is a first move toward illustrating that there might be a benefit to randomization. We are offering randomization as a

tool. How much you are going to want to do it is going to be up to all the stakeholders in the allocation to decide. And, of course, how they decide is another research question all together," says Wilson.

More information: Shomik Jain et al, Scarce Resource Allocations That Rely On Machine Learning Should Be Randomized, *arXiv* (2024). DOI: [10.48550/arxiv.2404.08592](https://doi.org/10.48550/arxiv.2404.08592)

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: Study: When allocating scarce resources with AI, randomization can improve fairness (2024, July 24) retrieved 24 July 2024 from <https://techxplore.com/news/2024-07-allocating-scarce-resources-ai-randomization.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.