



HAL
open science

On the Need and Applicability of Causality for Fair Machine Learning

Rūta Binkytė, Ljupcho Grozdanovski, Sami Zhioua

► **To cite this version:**

Rūta Binkytė, Ljupcho Grozdanovski, Sami Zhioua. On the Need and Applicability of Causality for Fair Machine Learning. 2023. hal-04329115v1

HAL Id: hal-04329115

<https://inria.hal.science/hal-04329115v1>

Preprint submitted on 7 Dec 2023 (v1), last revised 21 Jan 2024 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial - NoDerivatives 4.0 International License

On the Need and Applicability of Causality for Fair Machine Learning

Ruta Binkyte
Ljupcho Grozdanovski
Sami Zhioua

RUTA.BINKYTE@INRIA.FR
LGROZDANOVSKI@ULIEGE.BE
SAMI.ZHIOUA@INRIA.FR

Editors: Francesco Locatello and Vanessa Didelez

Abstract

Besides its common use cases in epidemiology, political, and social sciences, causality turns out to be crucial in evaluating the fairness of automated decisions, both in a legal and everyday sense. We provide arguments and examples, of why causality is particularly important for fairness evaluation. In particular, we point out the social impact of non-causal predictions and the legal anti-discrimination process that relies on causal claims. We conclude with a discussion about the challenges and limitations of applying causality in practical scenarios as well as possible solutions.

Keywords: Fairness, Discrimination, Causality, Anti-discrimination law

1. Introduction

Causal reasoning has an indispensable role in how humans make sense of the world and come to decisions in everyday life. It is true across epochs and continents although the particularities and points of emphasis may differ [Balogun \(2010\)](#). Although, in response to the critique from Karl Pearson, Ernst Mach, and Bertrand Russel, 20th century science reserved from making causal claims as too strong and not achievable, the 21th century is marked by the return of causality encouraged by the mathematization of causal notions and introduction of the non-deterministic concept of cause [Illari et al. \(2011\)](#).

Measurement of causal relationships can be made using experimental data (experimental causality) or observational data (statistical causality). The first approach is based on the random assignment experiment, which is widely recognized as the gold standard for proving causal relationships [Pearl \(2009\)](#); [Fisher \(1936\)](#); [Bloom \(2008\)](#). However, random assignment is often impractical or impossible. For example, it is not ethical to assign random participants in the experiment to smoke or engage in other hazardous activities. In a fairness-related scenario, it is impossible to change someone's race or gender, to measure its impact on hiring decisions or income. This is why causality is often determined based on observable outcomes (the second approach). However, it requires specific instruments to distinguish causality from statistical correlations.

The recent increase in the application of AI in domains of high social impact, such as healthcare, insurance, or hiring, has shown the risk of amplifying historical discrimination through automated decision systems. The best-known examples are racial discrimination in the Compas algorithm used to predict recidivism [Brackey \(2019\)](#) and sex discrimination in the selection of Amazon job candidates [Dastin \(2018\)](#). Accurate measurement of discrimination is important for evaluating the data or algorithm and for advising methods for achieving fairness. Recently, the domain has seen an increase in the use of statistical causality methods to evaluate and mitigate discrimination in data and algorithmic decisions.

This article consolidates the statistical and legal arguments for using causality in fair AI as well as practical challenges. We argue that causality is needed to appropriately address the problem of fairness in ML based automated decision systems. We summarize the benefits of using causality in three arguments, namely, (1) reliably measuring discrimination, (2) mediation analysis, and (3) establishing causal evidence in legal practice. Compared to existing work, the latter argument can be seen as the first attempt to connect causality in fair AI with the European AI legislation.

Tackling the problem of fairness from a causal perspective is plagued with practical obstacles that hinder its use in real scenarios. This includes the existence of several constraining assumptions that need to be satisfied and the availability of the causal graph. The last part of the paper describes the different assumptions and discusses their implications in the specific context of ML fairness.

2. Reliably measuring discrimination

Measuring discrimination without taking into consideration the causal structure underlying the relationships between variables may lead to misleading conclusions. That is, a biased estimation of discrimination. In extreme cases, such as Simpson’s paradox, the bias may lead to reversing the conclusions (e.g. the biased estimation indicates a positive discrimination, while the unbiased estimation is actually a negative discrimination). Figures 1-3 show the three basic causal structures that can lead to statistical anomalies, and consequently makes common statistical metrics of fairness unreliable.




Figure 1: Confounder structure. Figure 2: Mediator structure. Figure 3: Collider structure.

2.1. Confounder structure

The first situation where ignoring the causal structure of the data may lead to an unreliable estimation of discrimination is due to a failure to consider a confounder variable. Consider the hypothetical example in Figure 1 of an automated system to select candidates for job positions. Assume that the system takes as input two features, namely, the socio-economic status (SES) denoted as Z and the political belief of the candidate A . The outcome Y is whether the candidate is selected for the next stage of hiring (or the probability the candidate is selected). The outcome Y is influenced by the SES (A better SES makes it possible for candidates to attend more reputable academic institutions and to be enrolled in costly trainings). Both variables can be either binary (Z might be either rich or poor while A might be either liberal or conservative) or continuous (how rich/poor a candidate is for Z and the degree of conservativeness of the candidate for A). The political belief A of a candidate can be influenced by several variables, but in this example, assume that it is only influenced by the SES of the candidate. Finally, assume that the automated decision system is suspected to be biased by the political belief of candidates. That is, it is claimed that the system will more likely select candidates with a particular political belief.

A simple approach to check the fairness of the automated selection Y with respect to the sensitive attribute A is to contrast the conditional probabilities: $\mathbb{P}(Y = 1 \mid A = 0)$ and $\mathbb{P}(Y = 1 \mid A = 1)$, corresponding to statistical disparity, which quantifies the disparity in the selection rates between both types of candidates (conservatives and liberals). However such estimation of discrimination is biased due to the confounding path through Z . As Z variable causes both the sensitive variable A and the outcome Y , it creates a correlation between A and Y which is not causal. In other words, high SES (rich) candidates tend to have a more conservative political belief and at the same time more chances to be selected for the job (better academic institutions and training) which creates the following correlation in the data: employers will have more candidates with conservative political beliefs, and hence less candidates with liberal political beliefs. Such correlation is due to the confounder Z and should not count as discrimination. Most of statistical notions of fairness (equal opportunity, predictive parity, etc.) are not suitable to measure discrimination in presence of such statistical anomaly.

2.2. Mediator structure

The second situation where not accounting for the causal structure behind the data may lead to unreliable estimation of discrimination involves the presence of one or several mediator variables. The problem emerges from whether to consider the discrimination through a mediator variable as justifiable/acceptable or not. Similarly to confounding structure, a mediator variable may lead to Simpson's paradox. A famous example of Simpson's paradox caused by a mediator structure is the gender bias in 1973 Berkley admission [Bickel et al. \(1975\)](#); [Loftus et al. \(2018\)](#). Figure 2 shows the causal graph underlying the data where the sensitive variable (A) is gender, the outcome (Y) is the admission for Berkley graduate studies, and a single mediator variable (M) representing the department for which a candidate applied. In that year 1973, 44% of male applicants were admitted against only 34% of female applicants. While this looks like a bias against female candidates, when the same data has been analyzed by department, acceptance rates were approximately the same. In a simple mediator structure, there are two possible paths from A to Y : a direct path $A \rightarrow Y$ and an indirect path $A \rightarrow M \rightarrow Y$. Comparing the global admission rates of male and female candidates correspond to considering both paths when measuring discrimination. Whereas, comparing the admission rates per department correspond to considering only the direct path $A \rightarrow Y$. Hence, whether or not to consider mediator paths when measuring discrimination may lead to contradictory conclusions such as in Simpson's paradox.

2.3. Collider structure

A biased estimation of discrimination may be due to the presence of common effect (collider) variable and a data generation process implicitly conditioning on that variable. Using the same hypothetical example of job selection, consider the causal graph in Figure 3. A and Y are the same as in the previous example. Assume that data for training the automated decision system is collected from different sources, but mainly from labor union records. Assume also that variable W representing the labor union activism of the candidate is caused by both A and Y . On one hand, the political belief A influences whether a candidate is an active member of labor union (individuals with liberal political beliefs are more likely to enroll in labor unions). On the other hand, if a candidate is selected/hired, then there are higher chances that she becomes a member of labor union and conse-

quently that her case is recorded in the labor union records. Consistent with previous work, a box around a variable (W) indicates that data is generated by implicitly conditioning on that variable.

Again the simple approach of contrasting the selection rates between both types of candidates (conservatives and liberals) leads to a biased estimation of discrimination due to the colliding path through W . Intuitively, an individual has a record in the collected data either because she has liberal political beliefs or because she is selected for the job. Individuals who happen to have liberal political beliefs and at the same time selected for the job are still present in the data, however conditioning on labor union activism creates a correlation between A and Y which is not causal: data coming from labor union records includes fewer liberal candidates which are selected for the job than conservative candidates. Again, this is a discrimination against candidates with liberal political beliefs. Such correlation is due to the collider structure and should not count as discrimination.

3. Mediation Analysis




Figure 4: Causal graph with two mediated paths.

In presence of one or several mediator variables, it is useful to know how much of discrimination is direct and how much is mediated. More precisely, how much of discrimination is conveyed through each mediator variable. Mediation analysis is about distinguishing the different paths through which discrimination is going through and the portion of discrimination conveyed through each path. Consider another variant of the job hiring example in Figure 4 with two mediator variables, address (T) and community service (W). There are in total four different paths from the sensitive variable A (political belief) to the outcome variable Y (job hiring):

- $A \leftarrow Z \rightarrow Y$: confounding path
- $A \rightarrow Y$: direct path
- $A \rightarrow T \rightarrow Y$: indirect path through T
- $A \rightarrow W \rightarrow Y$: indirect path through W .

The first confounding path is non-causal and hence any effect slipping away through it should not be considered when estimating discrimination. As described in Section 2.1, that spurious effect is

due to how data is generated/collected and consequently should not count as actual discrimination. The direct path is present whenever there is an edge between A and Y . The effect through $A \rightarrow Y$ is always discriminatory, that is, it can never be justified and considered as acceptable discrimination. The two remaining paths are indirect paths going through mediator variables. Discrimination through an indirect path can or cannot be justified depending on the nature of the mediator variable. For instance, in the job hiring example of Figure 4, T (home address) is a mediator variable because, on one hand, having a certain political inclination may indicate where a candidate is living, and on the other hand, the job hiring decision may depend on the home address of a candidate. W (community service) is another mediator variable because, on one hand, political views of a candidate can influence how much involved she can be in community service, and on the other hand, community service record is a good indicator on how suitable she will be for a given position. Discrimination through the path $A \rightarrow W \rightarrow Y$ can be acceptable as an employer can justify a disparity between candidates with different political beliefs by their records of community service. However, discrimination through the path $A \rightarrow T \rightarrow Y$ is typically not acceptable because an employer cannot justify discrimination on the basis of the addresses of candidates. T is called a proxy variable, whereas W is called an explaining variable¹.

Causality, through the concepts of intervention and counterfactual, provides the tools required to distinguish between discrimination conveyed through different paths. Intervening on A , blocks all paths starting with an incoming edge to A which include all confounding paths between A and Y . Discrimination through all causal paths is captured by the average causal effect (ACE):

$$ACE(Y, A) = \mathbb{P}(Y = y^+ | do(A = 1)) - \mathbb{P}(Y = y^+ | do(A = 0)) \quad (3.1)$$

Where $Y = y^+$ is a positive decision and $A = 1, A = 0$ are the values of the sensitive attribute. In Figure 4, ACE expression captures discrimination through all paths, except $A \leftarrow Z \rightarrow Y$. For notation simplicity, we represent $Y = y^+$ simply as y^+ , $A = 1$ (resp. $A = 0$) as a_1 (resp. a_0), and the $do()$ operator with subscription. Hence, the right-hand side of Equation 3.1 becomes simply $\mathbb{P}(y_{a_1}^+) - \mathbb{P}(y_{a_0}^+)$.

To distinguish the direct discrimination from indirect discrimination, two expressions can be used, namely natural direct effect (NDE) and natural indirect effect (NIE) Pearl (2001):

$$NDE(Y, A) = \mathbb{P}(y_{a_1, \mathbf{Z}_{a_0}}^+) - \mathbb{P}(y_{a_0}^+) \quad (3.2)$$

where \mathbf{Z} is the set of all mediator variables and $\mathbb{P}(y_{a_1, \mathbf{Z}_{a_0}}^+)$ is the probability of a counterfactual situation where $Y = y^+$ had A been 1 and had \mathbf{Z} been the value it would naturally take if $A = 0$. Intuitively, $\mathbb{P}(y_{a_1, \mathbf{Z}_{a_0}}^+)$ is considered counterfactual because it corresponds to a candidate who is conservative ($A = 1$) across the direct path $A \rightarrow Y$ but liberal $A = 0$ across all indirect paths. NIE has a similar form but A values are reversed in the counterfactual expression:

$$NIE(Y, A) = \mathbb{P}(y_{a_0, \mathbf{Z}_{a_1}}^+) - \mathbb{P}(y_{a_0}^+) \quad (3.3)$$

Finally, distinguishing the discrimination conveyed through specific indirect paths is possible through path-specific effect (PSE) Pearl (2001); Chiappa (2019):

$$PSE(Y, A, \pi) = \mathbb{P}(y_{a_1 | \pi, a_0 | \bar{\pi}}^+) - \mathbb{P}(y_{a_0}^+) \quad (3.4)$$

1. In presence of a single path with a sequence of two or more mediators, the existence of at least one explaining variable among the mediators makes discrimination through that path justifiable and hence acceptable.

Where π is set of the variables on the path of interest, $\bar{\pi}$ is the set of variables not in π and $\mathbb{P}(y_{a_1}^+ | \pi, a_0 | \bar{\pi})$ is the counterfactual probability of $Y = y^+$ had A been 1 on the paths π and 0 on the remaining paths $\bar{\pi}$.

4. Uncovering causality through legal evidence: the regulatory approach in the European Union

The method of mediator structure in uncovering causation, as discussed in the previous section, is certainly a useful model for the proof of causality in judicial instances dealing with algorithmic discrimination. The question is, however, if procedural law, namely in the European Union (EU), is designed to support such an analysis. As a preliminary observation, we should stress that in law, the expression ‘causal fairness’ generally refers to the procedural conditions under which instances of fairness (or, unfairness for that matter) are causally represented. With this in mind, in this section, we will focus on two important and interrelated issues: evidence and procedural fairness². In the eye of the law, causality is a question of fact, calling for legally established discovery procedures - and corresponding reasoning models - meant to yield accurate causal representations i.e. allow for causality proper to be singled out from a myriad of correlations (positive associations between candidate-causes and a harm suffered) Haack (2014). However, in adjudicatory contexts, causality is proven for the purpose of fairness, typically compensation as a ‘fair’ outcome to the suffering of harm. Indeed, legal systems committed to the rule of law³ share a commitment to procedural fairness, the normative creed being that only fairly designed procedures can be conducive to fair outcomes. In contemporary systems of evidence and judicial remedies, including those in EU law, the fair procedures/outcomes parallelism is epitomized in the fair trial safeguards - procedural entitlements meant to uphold a level of basic equality (or procedural parity) and effectiveness in the ways in which litigants participate in a dispute resolution Solum (2004). This equality not only applies to the litigants’ ability to access judicial remedies, but also to their ability to access and give evidence, the idea being that one party should in no way be advantaged or disadvantaged over the other, in terms of their access to the facts needed to make their views known (usually, before a court). In short, ‘casual fairness’ in law calls for accurate - or at least, plausible Rescher (1980) - evidence of causality, presented in conditions of procedural fairness.

Proof of causality in connection to algorithmic discrimination has profoundly upset these long-standing legal postulates. From a procedural fairness perspective, a major thorny issue has been that AI’s relative or total opacity makes AI systems’ decisional processes inscrutable, obstructing the victims’ ability to properly establish and argue causation. One of the topical examples in this regard is Cook vs. HSBC North America⁴, a credit scoring case where the system used as a relevant variable the applicants’ places of residence, ultimately favoring ‘white’ areas and discriminating

2. In the EU, the Independent High Level Expert Group on AI, set up by the European Commission, defined procedural fairness as “entails the ability to contest and seek effective redress against decisions made by AI systems and by the humans operating them.” See HLEG, Ethics Guidelines for trustworthy AI, available at <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>, at 13.

3. In the EU, the concept of rule of law is understood to include the following principles: legality, legal certainty, prohibition of arbitrariness of the executive powers, independent and impartial courts, effective judicial review, including respect for fundamental rights and equality before the law. See Communication from the Commission to the European Parliament and the Council ‘A new EU Framework to strengthen the Rule of Law,’ COM(2014) 158 final, at 4.

4. US District Court for the Northern District of Illinois, 21 March 2014, County of Cook v. HSBC North America Holdings Inc et al., 1:2014cv02031.

against members of ethnic minorities. Those ‘subtly discriminatory’ variable associations (such as zip code/ethnic background) combined with the practical difficulties of accessing relevant information on how an AI system associated different variables, meant that the right to access evidence and courts (as a fair trial safeguard) were under serious threat. To remedy this, regulators across the world and in the EU sought to answer two main questions: 1. which evidence do litigants need to have access to in order to effectively prove causation?; 2. once that evidence has been identified, how should legal procedures be (re)designed to ‘open’ the victims’ access to it?

4.1. Using causal tools to establish causal evidence

Regarding the first question, the emerging - but not yet consolidated - global AI liability caselaw reveals an interesting trend. Though many judicial instances can be cited as examples, for the purpose of this article, we shall highlight three cases we view as illustrative of the ‘new approach’ to proving causation in AI-related disputes. These cases are Pickett ⁵ (dealing with a DNA matching system - TrueAllele - used by police authorities to track down harm-doers), Loomis ⁶ (dealing with COMPAS, a recidivism-predicting system used by courts) and Ewert ⁷ (also dealing with the use of recidivism-predicting systems by Canadian correctional services). In all three cases, the plaintiffs argued that the automated decisions were inaccurate because they were unfair that is, contained unfair biases: gender in Pickett and Loomis, ethnic background in Ewert. To uncover the bias-conducive variable association (i.e the causal link), the plaintiffs requested that the systems be reverse engineered. This was hardly possible. For example, in Pickett, independent experts confirmed that reverse engineering would take up to 8,5 years to be completed ⁸. In the face of the practical unfeasibility of reverse-engineering, the court in Pickett (and in Loomis) turned to general expertise, as a *faute de mieux* solution: the lack of direct evidence (reverse-engineering) able to reveal the presence of an unfair bias, was ‘compensated’ by the recourse to already existing expertise assessing a system’s functionalities in general. If the majority of experts agreed that a system, like TrueAllele in Pickett or COMPAS in Loomis, was generally well-performing (i.e. was unbiased and therefore accurate), the courts would be inclined to accept that, in the disputes they were called to resolve, it could be presumed that the systems concerned had made unbiased decisions.

Hence, the role of experts is to assess the strength of the causal link between sensitive variables and the decision (A zero causal effect indicates absence of discrimination) in presence of different causal structures (Section 2) which may lead to different types of bias. A possible approach would be to identify the causal graph in order to unveil the causal relations between variables and then to use causal notions of fairness (Section 3) to assess discrimination. A suggested procedure to identify the causal graph is to first use a causal discovery algorithm (e.g. PC [Spirtes et al. \(1999\)](#)). Then, seek the input of experts in the domain of application to adjust the discovered graph (e.g. adding/removing causal links, enforcing assumptions, etc.). The input of experts can be useful also to clarify the role of each variable, in particular, classifying mediator variables into explaining

5. Superior Court of New Jersey (Appellate Division), 2 February 2021, *State of New Jersey v. Corey Pickett*, Docket N° A-4207-19T4.

6. Supreme Court of Wisconsin, 13 July 2016 (decided), *State of Wisconsin v. Eric L. Loomis*, 881 N.W. 2d 749 (2016) 2016 WI 68.

7. *Ewert vs. Canada*, 2018 SCC 30, File n° 37233, 13 June 2018.

8. See Superior Court of New Jersey (Appellate Division), 2 February 2021, *State of New Jersey v. Corey Pickett*, Docket N° A-4207-19T4, at 17.

(leading to justifiable discrimination) and proxy (leading to unjustifiable discrimination) variables. This is essential to select the suitable causal fairness metric (Section 3) to use.

4.2. But-for test using counterfactuals

From the perspectives of procedural fairness and the mediator structure model, this trend is of course open to criticism. First, general expert opinions on a system's accuracy are not as probative as direct evidence (reverse engineering) able to provide highly reliable information on the mediator association having led to a discriminatory outcome. Second, the inability to prove causation through reliable evidence seems to have given way to a peculiar application of the so-called but-for test. In principle, this test translates to the deployment of counterfactual reasoning seeking to determine if a harm would have been suffered, had an alleged cause not occurred. In the Cook vs HSBC case (credit scoring) e.g., a standard application of said test would translate to determining if the same loan applicants would have been approved, if the system had not taken their places of residence as a relevant variable. However, the cases cited in this section (in particular Pickett and Loomis), reveal a slight shift in the application of the but-for test. In 'ordinary' disputes (non-AI related) cases, this test seeks to answer a question of factive causal association: would an outcome be the same (or different) without certain facts (address, gender, age, etc) in the causal structure? In AI-related disputes, the but-for test answers a question of (human) reliance on AI output, the relevant (causal) issue being if a human decision based on AI would have been the same or different, had the AI not been used at all. In this case, statistical causality tools can be applied to detect discrimination in data reflecting previous hiring or loan-granting practices in the company concerned. If the association between the sensitive attribute and the outcome is detected, then one can conclude that the decision would be the same without the algorithmic assistance. This brings the focus of attention from the AI system (and its architects) to the general practice in the company. Here, again, causality can help to distinguish between spurious association, explainable disparity, or discrimination. On the other hand, if the data with ingrained discrimination is the same as that used to train the algorithm, compliance with AI designing guidelines can be further scrutinized. Finally causality tools provide mathematical expressions to capture the intangible concept of counterfactual [Shpitser and Pearl \(2007, 2008\)](#) very useful to check directly the but-for test.

This allows us to raise the second issue mentioned above: should systems of evidence include a right to access/to request disclosure of evidence?

4.3. Disclosing causal evidence to victims of discrimination

From a procedural fairness perspective, such a right seems paramount for a victim of algorithmic discrimination to at least have a shot at requesting the 'lifting of the opacity veil' that might cover a causal chain [Grozdanovski \(2021\)](#). In the EU, recent regulatory developments seemed - on the surface at least - to move toward the recognition of such a right. First came the AI Act ⁹ - a horizontal, across-the-board legislation which makes two important contributions. On the one hand, it includes a four-level taxonomy of risks-of-harm related to AI systems: non-high, limited, high and unacceptable. On the other hand, and against the backdrop of said risk-taxonomy, the AI Act includes a set of technical standards (transparency, data governance, risk-mitigation strategy

9. Proposal for a Regulation of the European Parliament and of the Council laying down harmonized rules on Artificial Intelligence (AI Act) and amending certain Union legislative acts, COM(2021) 206 final.

etc) targeting high-risk AI systems, used in mainly eight market sectors ¹⁰. To complement the AI Act and to afford procedures designed for the compensation of harm associated with high-risk systems, the AI Liability Directive (AILD) ¹¹ came next. This instrument establishes a system of evidence which grants victims the right to request disclosure of evidence. By virtue of the AILD, if the defendant (a programmer or user) refused to disclose the evidence requested by the victim or if, upon disclosure, a national or EU court found that the evidence was probative and plausible, the defendant would be presumed responsible for the harm (e.g. discrimination) suffered by the claimant. It should however be stressed that the evidence a victim can ask disclosure of under the AILD does not include the evidence flagged as ‘necessary’ (i.e. expertise) in the cases cited earlier. The AILD allows the disclosure of evidence so long as that evidence pertains to the defendant’s compliance with the technical standards listed in the AI Act. In other words, the defendant would not be asked to provide information (e.g. access to the code, reverse-engineering, when feasible) able to support a proper causal analysis. They would be asked to - merely - provide information confirming that they complied with, say, their duty for human control and oversight. The reason for this is, no doubt, that the AILD relies on the assumption that if harm (like discrimination) does occur, it is because the AI Act had not been fully observed. In doing so, the AILD narrows down the scope of the evidentiary debate in the sense that the parties in future AI discrimination cases, will not seek to be called to uncover the actual casual structure underlying discriminatory AI output, but to identify the human agent who had failed to meet a legally prescribed duty of care.

5. Practical Considerations for Using Causality for Fairness

In the previous sections, we illustrate the situations where the causality approach is relevant for evaluating fairness and how it can be attained using causal fairness notions and approximation techniques. Despite the apparent advantages, the applicability of the causal framework is limited because of its reliance on prior knowledge and often untestable assumptions. Many causal requirements can be achieved by applying a specific experiment design (ideally, random assignment). However, in fairness scenarios, it is often not a plausible option. Therefore, discrimination is usually evaluated from observational data. Here, we will list some requirements for applying causal inference that are most relevant for fairness applications.

5.1. Possibility for Intervention

In the fairness estimation, the sensitive attribute is considered to be the exposure or treatment attribute. The goal is to measure its impact on the outcome. Most definitions of causal effect are based on a notion of intervention or manipulation of a cause variable (exposure) [VanderWeele and Hernán \(2012\)](#); [Holland \(1986\)](#). This makes it hard to justify causal claims related to nonmanipulable quantities, such as sensitive attributes, for example, race or gender. Some approaches in the literature suggest shifting attention from an actual manipulation to changes in perception [Rahmatlabi and Xiang \(2022\)](#). For example, instead of changing the gender of candidates to estimate the

10. The ‘high-risk’ sectors are listed in Annex III of the AI act and include Employment, education, healthcare, transport, energy, public sector (including asylum, migration, border controls, judiciary and social security services), defence and security, finance, banking, and insurance.

11. Proposal for a Directive of the European Parliament and of the Council on adapting non-contractual civil liability rules to Artificial Intelligence (AI Liability Directive) COM(2022)496 final.

effect on a hiring decision, the researcher could manipulate the perception of gender by the employer. It could be easily done by submitting two analogous resumes but varying the name or title of the candidate. This approach corresponds to the methodology applied in social experiments on the impact of race or gender of an applicant on hiring decisions [Neumark \(2018\)](#). [VanderWeele and Hernán \(2012\)](#) further differentiate immutable sensitive attributes into those that are randomized at birth (biological sex) and those that are not (race, social gender). This distinction is important when estimating the causal effect of the sensitive attribute. If the sensitive attribute is randomized, then its causal effect on an outcome can be estimated just by comparing exposure levels. For example, it is possible to estimate the total causal effect of biological sex by taking the observed differences in the outcome between men and women [VanderWeele and Hernán \(2012\)](#). On the contrary, race is not random but depends on many ancestral factors. For this reason, even at the biological level estimating the effect of race is more complicated and requires information about the causal structure of the covariates. However, these types of estimation are relevant in medical scenarios, where the independence between the sensitive attribute and the outcome (for example, the probability of a disease) cannot be reasonably assumed. In the possible discrimination scenarios [VanderWeele and Hernán \(2012\)](#), similarly, to [Rahmattalabi and Xiang \(2022\)](#) shift attention to the direct effect of the *perceived* gender or race on the decision.

5.2. Causal assumptions

The SUTVA [Rubin \(1986\)](#) (Stable Unit Treatment Value Assumption) entails the requirements of no interference and consistency. No interference assumption requires that the interaction between individuals does not influence the effect of the sensitive attribute on the outcome. The likelihood of interaction and feedback loops is high in social sciences research in general and calls for a clear discussion and restricted interpretations of causal estimation [Morgan and Winship \(2015\)](#). Fairness is usually measured in a social context. Therefore, the possibility of interaction should be carefully evaluated. Using the hiring example, the violation of the SUTVA requirement would occur in a situation where hiring more participants of one political spectrum increases the likelihood of privileging the same political spectrum in future hiring decisions. Such a scenario is plausible because current employees may favor those who have political beliefs similar to their own. The assumption of consistency requires that each treatment level leads to the same potential outcomes [Keele \(2015\)](#). In fairness evaluation, treatment is replaced by the sensitive attribute, which is often a social construct such as race or gender. Identifying the causal effect of gender on hiring can be problematic if gender itself does not have a consistent effect on hiring. For example, only women with a certain level of "femininity" are discriminated against. This scenario cannot be excluded and should be considered if a fine-grained causal analysis is a goal of a study. In summary, SUTVA assumptions are likely to be violated in fairness scenarios, however, causal approaches can still be applied if the results are interpreted with caution. Some methods for identifying the causal effect under the violations of SUTVA are discussed here [Laffers and Mellace \(2020\)](#).

Ignorability [Rubin \(1986\)](#) assumption requires that the sensitive attribute and the outcome are independent given the observable variables. In other words, no unobserved variables create a significant link between the sensitive attribute and the outcome. In fairness evaluation, the presence of such a link could mean that the portion of discrimination is, in fact, a spurious effect induced by the confounder. For example, if the education confounder is not present in the data, the confounding effect cannot be controlled. As a result, it is not possible to estimate the causal effect of political

belief on the hiring decision that is separate from the effect of education. Unobserved confounders are not likely for immutable sensitive attributes such as sex or race. These sensitive attributes are unlikely to have a temporally prior cause. However, the noise terms can still be not independent between the sensitive attribute and the outcome. [Fawkes et al. \(2022\)](#) point out, the implications of assuming ignorability, when using causal counterfactuals. Following the reasoning by [Fawkes et al. \(2022\)](#), in the case of college admission (Figure 2), an average male who applied to the technical profession could be counterfactually exchanged with an average woman who applied to the same profession. However, given the social expectations tied to gender roles, a woman applying to a technical profession is likely to be more motivated and hard-working than an average male with the same professional goals.

Positivity [Rubin \(1986\)](#) is violated if some of the combinations of a sensitive attribute and a covariate have zero probability. Violations of positivity can be deterministic or random [Westreich and Cole \(2010\)](#). For example, positivity would be violated if a certain level of education always corresponds to liberal political beliefs. In this scenario, the positivity violation would most likely be random. It is unlikely that certain education would have a deterministic relationship on political beliefs. In the random case, statistical methods are available for analysis under violation of positivity [Westreich and Cole \(2010\)](#). However, consider a case where having a Harvard degree is considered an explanatory mediator between ethnicity and hiring. Certain ethnicities may have zero probability of having obtained a Harvard degree due to long-term discrimination and poverty. In this case, it should be reconsidered if a specific Harvard degree is essential for the job considered despite the potential exclusion and disparate impact.

The identifiability of path-specific effects in the presence of multiple mediators requires the absence of causal links between the mediators [VanderWeele and Vansteelandt \(2014\)](#). Evaluating path-specific effects is particularly important to understand the mechanism of the effect of the sensitive attribute on the outcome. As outlined earlier (Section 3), the effect can be deemed justifiable or discriminatory depending on the mediating variables on the path. However, the link between two or more mediators is likely in fairness scenarios. For example, consider the case where race and hiring decisions are mediated by social status (redlining) and education (explaining variable). It is very likely that the level of education is influenced by social status. In this case, the indirect effect through social status and education separately is not identifiable. Work by [VanderWeele and Vansteelandt \(2014\)](#) proposed a method based on the treatment of multiple mediators together. In some cases, this method can help identify individual indirect effects in the presence of causal links between the mediators.

5.3. Availability of Causal Graph

One of the most significant restrictions for using causality is knowledge of the relationship between variables in the form of a directed acyclic graph (DAG) ¹² [Pearl \(2009\)](#). The research by [Binkytė-Sadauskienė et al. \(2022\)](#) shows a significant disagreement between estimations of causal fairness notions due to slight differences in the causal structure. The availability of DAG is particularly important in the presence of collider structures, because including a collider in a conditioning set induces bias in measuring causal effect [Pearl \(1988\)](#). DAG is also important for the evaluation

12. The DAG is subject to further assumptions of causal Markov condition, causal faithfulness, and causal sufficiency. Causal Markov condition, causal faithfulness, and Causal sufficiency together encode the same requirements as the SUTVA and Ignorability in the potential outcome framework, therefore, will not be discussed separately.

of path-specific effects, important for distinguishing redlining and explaining variables in fairness scenarios.

The causal structure (or causal graph) can be obtained by consulting domain experts or learning from observational data. Both approaches have their own limitations. Domain experts can disagree or have biased assumptions. Learning from the data often requires additional assumptions on the data distribution, functional relationships, the relations of exogenous unobserved variables, and the informed choice of the learning algorithm. The research by [Binkytė-Sadauskienė et al. \(2022\)](#) show how different algorithms to recover the causal structure yield different results when applied to the same data set. Learning causal relationships from observational data alone may not be realistic [Guyon et al. \(2011\)](#). However, the combination of causal discovery and expert knowledge could give more reliable results.

6. Conclusions

We outline the circumstances in which causality is needed to evaluate the fairness of AI decisions. We show how using purely statistical tools discrimination can be exaggerated, diminished, or even reversed. We link the reliable evidence of algorithmic discrimination to court practice and the European legal framework. More precisely, we discuss how statistical causality relates to inculcation in court and what evidence can be used to prove discrimination. We distinguish direct evidence in the form of reverse engineering the algorithm and indirect evidence. As direct evidence is often not feasible in practice, we discuss options for using causal analysis on the indirect evidence. This can be done with two approaches to "but for" counterfactual reasoning. The first approach refers to the sensitive variable or redlining mediator in the algorithmic decision. In this case, the test data can be used to explore the connection between the sensitive attribute and the labels produced by the algorithm. The second approach explores whether the decision would be taken if the AI system was not used. Here we focus on the ground truth decision in the training data or the data that reflect the historical practices in the company. Finally, we discuss the causal assumptions in the context of fairness.

Acknowledgments

We thank a bunch of people and funding agency.

References

- Abiodun Balogun. Causality in western science and traditional african thought: a study in comparative philosophy. *Saarbrücken: Lambert Academic Publishers*, 2010.
- Peter J Bickel, Eugene A Hammel, and J William O’Connell. Sex bias in graduate admissions: Data from berkeley. *Science*, 187(4175):398–404, 1975.
- Rūta Binkytė-Sadauskienė, Karima Makhlof, Carlos Pinzón, Sami Zhioua, and Catuscia Palamidessi. Causal discovery for fairness. *arXiv preprint arXiv:2206.06685*, 2022.
- Howard S Bloom. The core analytics of randomized experiments for social research. *The SAGE handbook of social research methods*, pages 115–133, 2008.

- Adrienne Brackey. *Analysis of Racial Bias in Northpointe’s COMPAS Algorithm*. PhD thesis, Tulane University School of Science and Engineering, 2019.
- Silvia Chiappa. Path-specific counterfactual fairness. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 7801–7808, 2019.
- Jeffrey Dastin. Amazon scraps secret ai recruiting tool that showed bias against women. In *Ethics of Data and Analytics*, pages 296–299. Auerbach Publications, 2018.
- Jake Fawkes, Robin Evans, and Dino Sejdinovic. Selection, ignorability and challenges with causal fairness. In *Conference on Causal Learning and Reasoning*, pages 275–289. PMLR, 2022.
- Ronald Aylmer Fisher. Design of experiments. *British Medical Journal*, 1(3923):554, 1936.
- Ljupcho Grozdanovski. In search of effectiveness and fairness in proving algorithmic discrimination in eu law. *Common Market Law Review*, 58(1), 2021.
- Isabelle Guyon, Constantin Aliferis, Gregory Cooper, André Elisseeff, Jean Philippe Pellet, Peter Spirtes, and Alexander Statnikov. Causality workbench. In *Causality in the sciences*. Oxford University Press, 2011.
- Susan Haack. *Evidence matters: Science, proof, and truth in the law*. Cambridge University Press, 2014.
- Paul W Holland. Statistics and causal inference. *Journal of the American statistical Association*, 81(396):945–960, 1986.
- PM Illari, Federica Russo, and Jon Williamson. Why look at causality in the sciences? a manifesto. *Causality in the Sciences*, pages 3–22, 2011.
- Luke Keele. The statistics of causal inference: A view from political methodology. *Political Analysis*, 23(3):313–335, 2015.
- Lukas Laffers and Giovanni Mellace. Identification of the average treatment effect when sutva is violated. *Discussion Papers on Business and Economics, University of Southern Denmark*, 3, 2020.
- Joshua R Loftus, Chris Russell, Matt J Kusner, and Ricardo Silva. Causal reasoning for algorithmic fairness. *arXiv preprint arXiv:1805.05859*, 2018.
- Stephen L Morgan and Christopher Winship. *Counterfactuals and causal inference*. Cambridge University Press, 2015.
- David Neumark. Experimental research on labor market discrimination. *Journal of Economic Literature*, 56(3):799–866, 2018.
- Judea Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan kaufmann, 1988.
- Judea Pearl. Direct and indirect effects. In *Proceedings of the Seventeenth conference on Uncertainty in artificial intelligence*, pages 411–420, 2001.

- Judea Pearl. *Causality*. Cambridge university press, 2009.
- Aida Rahmattalabi and Alice Xiang. Promises and challenges of causality for ethical machine learning. *arXiv preprint arXiv:2201.10683*, 2022.
- Nicholas Rescher. Plausible reasoning: An introduction to the theory and practice of plausibilistic inference. *Philosophy and Rhetoric*, 13(3), 1980.
- Donald B Rubin. Statistics and causal inference: Comment: Which ifs have causal answers. *Journal of the American Statistical Association*, 81(396):961–962, 1986.
- Ilya Shpitser and Judea Pearl. What counterfactuals can be tested. In *23rd Conference on Uncertainty in Artificial Intelligence, UAI 2007*, pages 352–359, 2007.
- Ilya Shpitser and Judea Pearl. Complete identification methods for the causal hierarchy. *Journal of Machine Learning Research*, 9(Sep):1941–1979, 2008.
- Lawrence B Solum. Procedural justice. *S. Cal. l. rev.*, 78:181, 2004.
- Peter Spirtes, Christopher Meek, and Thomas Richardson. An algorithm for causal inference in the presence of latent variables and selection bias. *Computation, causation, and discovery*, 21: 211–252, 1999.
- Tyler VanderWeele and Stijn Vansteelandt. Mediation analysis with multiple mediators. *Epidemiologic methods*, 2(1):95–115, 2014.
- Tyler J VanderWeele and Miguel A Hernán. Causal effects and natural laws: towards a conceptualization of causal counterfactuals for nonmanipulable exposures, with application to the effects of race and sex. *Causality: statistical perspectives and applications*, pages 101–113, 2012.
- Daniel Westreich and Stephen R Cole. Invited commentary: positivity in practice. *American journal of epidemiology*, 171(6):674–677, 2010.