



**HAL**  
open science

## When causality meets fairness: A survey

Karima Makhoulf, Sami Zhioua, Catuscia Palamidessi

► **To cite this version:**

Karima Makhoulf, Sami Zhioua, Catuscia Palamidessi. When causality meets fairness: A survey. *Journal of Logical and Algebraic Methods in Programming*, 2024, 141, pp.101000. 10.1016/J.JLAMP.2024.101000 . hal-04950308

**HAL Id: hal-04950308**

<https://inria.hal.science/hal-04950308v1>

Submitted on 16 Feb 2025

**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 4.0 International License

# When Causality Meets Fairness: A Survey

Karima Makhlouf<sup>a,b</sup>, Sami Zhioua<sup>1b</sup>, Catuscia Palamidessi<sup>a,b</sup>

<sup>a</sup>*INRIA, Palaiseau, 91120, France*

<sup>b</sup>*École Polytechnique, Palaiseau, 91120, France*

---

## Abstract

Addressing the problem of fairness is crucial to safely using machine learning algorithms to support decisions that have a critical impact on people’s lives, such as job hiring, child maltreatment, disease diagnosis, loan granting, etc. Several notions of fairness have been defined and examined in the past decade, such as statistical parity and equalized odds. However, the most recent notions of fairness are causal-based and reflect the now widely accepted idea that using causality is necessary to appropriately address the problem of fairness. This paper examines an exhaustive list of causal-based fairness notions and studies their applicability in real-world scenarios. As most causal-based fairness notions are defined in terms of non-observable quantities (e.g., interventions and counterfactuals), their deployment in practice requires computing or estimating those quantities using observational data. This paper offers a comprehensive report of the different approaches to infer causal quantities from observational data, including identifiability (Pearl’s SCM framework) and estimation (potential outcome framework). The main contributions of this survey paper are (1) a guideline to help select a suitable causal fairness notion given a specific real-world scenario and (2) a ranking of the fairness notions according to Pearl’s causation ladder, indicating how difficult it is to deploy each notion in practice.

*Keywords:* Fairness, machine learning, causality, causal inference, intervention, counterfactual

---

## 1. Introduction

Machine learning algorithms are increasingly used to inform automated decisions with critical impact on people’s lives including job hiring, loan granting, predicting recidivism during parole, etc. Correcting bias in the decision prediction requires first

---

<sup>1</sup>Corresponding author

to measure it. The most commonly used fairness notions are observational and rely on mere correlation between variables. For example, statistical parity [1] requires that the proportion of positive outcome (e.g., granting loans) is the same for all sub-populations (e.g., male and female groups). Equal opportunity [2] requires that the true positive rate (TPR) is the same for all sub-populations. The main problem of correlation-based fairness notions is that they fail to detect discrimination in the presence of statistical anomalies such as Simpson’s paradox [3]. A famous example of Simpson’s paradox is the gender bias in 1973 Berkeley admission [4, 5]. In that year, 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 was analyzed by the department, acceptance rates were approximately the same. In other words, the statistical conclusions drawn from the sub-populations differ from that of the whole population. Considering the problem of fairness from the legal and philosophical point of view reveals another limitation of statistical fairness notions. In the disparate treatment liability framework [6], discrimination claims require plaintiffs to demonstrate a causal connection between the challenged decision (e.g., hiring, firing, admission) and the sensitive feature (e.g., gender, race). It is then necessary to investigate the causal relationship between the sensitive attribute and the decision rather than the associated relationship. Because of these two limitations, it is now widely accepted that causality is necessary to appropriately address the problem of fairness [5].

Various causal-based fairness notions have been recently proposed to tackle the problem of fairness through causal inference lenses. These include total effect [7], counterfactual fairness [8], counterfactual effects [9], interventional fairness [10], etc. These notions differ from statistical fairness approaches in that they are not totally based on data but consider additional knowledge about the structure of the world, in the form of a causal model. This additional knowledge helps to understand how data is generated in the first place and how changes in variables propagate in a system. Most of these fairness notions are defined in terms of non-observable quantities such as interventions (to simulate random experiments) and counterfactuals (which consider other hypothetical worlds, in addition to the actual world). Such quantities cannot always be uniquely computed from observational data, which hinders significantly the applicability of causal-based notions in practical scenarios. Each one of the two main causal frameworks in the literature, namely, structural causal model (SCM) with causal graphs [7] and potential outcome [11], uses a different approach to compute/estimate the causal quantities using observational data. The SCM framework relies mainly on the identifiability criterion [12] to generate an expression for the causal quantity based only on observable probabilities. If the identifiability criterion

is unsatisfied, the causal quantity can not be computed using the available observable data. In such case, as an alternative, if the complete structure of the causal model is available, it is possible to estimate the distribution of the latent variables  $U$  and consequently generate an estimation of the counterfactual outcomes [8]. The potential outcome framework approximates causal quantities using several estimation techniques (e.g., matching, re-weighting, etc.) [13].

Nineteen causal-based fairness notions are examined in this paper. Given a real-world scenario, selecting which fairness notion to use is a challenging and error-prone task. Using the wrong fairness notion may indicate unfairness in an otherwise fair scenario, or the opposite (failing to detect unfairness in an unfair scenario). This survey paper provides guidelines to help select a suitable fairness notion given a specific real-world scenario. The guidelines are summarized in a decision diagram that can be easily navigated using the characteristics of the real-world scenario at hand. On the other hand, according to Pearl’s SCM framework, computing causal quantities (interventions and counterfactuals) depends on their identifiability. Hence, even if a fairness notion is appropriate in some setups, it might not be applicable because of identifiability issues. Placing the various causal-based fairness notions in Pearl’s causation ladder with the three corresponding rungs (observation, intervention, and counterfactual) [14] allows us to rank these notions and indicates how difficult it is to deploy each one of them in practice.

This survey paper is a comprehensive report on assessing machine learning fairness with causality lenses. It starts with a related work section (Section 2). Section 3 illustrates the need for causality through a hypothetical example of teacher firing (Section 3). Then, it provides essential background on causal inference in sufficient detail for our analysis (Section 4). Section 5 examines a comprehensive list of causal-based fairness notions. Unlike other surveys in the literature, the subtleties of the fairness notions are illustrated using a very simple numerical job hiring example. A survey on the three approaches to computing causal quantities from observable data, namely, identifiability, estimation based on the full causal model, and potential outcome estimation, is provided in Section 6. The main contributions of the survey, which are the suitability and applicability of causal-based fairness notions, are described in Section 7. The practical considerations of causal-based fairness notions are summarized in Section 8. Finally, Section 9 concludes.

## 2. Related work

With the lack of a universal definition of fairness, several metrics has been defined in the literature. Most of these metrics fall into the category of statistical metrics of

fairness. These include statistical parity [15], equal opportunity [2], calibration [16], individual fairness [15], etc. Existing surveys about fairness in machine learning focused mainly on statistical metrics.

Zliobaite’s survey [17] classified fairness notions into four categories, namely, statistical tests, absolute measures, conditional measures, and structural measures. Regarding the applicability of notions, the only criterion considered was the type of variables (e.g. binary, categoric, numeric, etc.). The survey of Berk et al. [18] listed only group fairness notions that are defined using the confusion matrix. The applicability aspect has not been addressed, as the paper focused only on criminal risk assessment use cases. Verma and Rubin [19] described a list of fairness metrics and illustrated how each notion can be computed in real scenarios through a loan granting real use case (German credit dataset [20]). Mehrabi et al. [21] considered a more general scope for their survey: in addition to briefly listing 10 definitions of fairness notions, they surveyed different sources of bias and different types of discrimination, they listed methods to implement fairness categorized into pre-processing, in-processing, and post-processing, and they discussed potential directions for contributions in the field. A more recent survey by Mitchell et al. [22] presents an exhaustive list of fairness notions in both categories (group and individual) and summarizes most of the incompatibility results in the literature. Makhlof et al.’s survey [23] is the most exhaustive in terms of covered statistical metrics while at the same time including the most notable causal-based fairness metrics (e.g., counterfactual fairness [8]). However, what distinguishes Makhlof et al.’s survey [23] is the focus on the applicability of fairness notions and the identification of a set of thirteen criteria to help select the most suitable notion to use given a scenario at hand.

More related to causal fairness, survey papers focused on advocating the need for causality to appropriately address the problem of fairness. Loftus et al. [5] were among the first to argue that causality is critical for fairness. Their main argument is that causal approaches to fairness depart from the previous statistical approaches in that they are not wholly data-driven but require additional knowledge of the structure of the world, in the form of a causal model. In terms of fairness metrics, Loftus et al. focused mainly on counterfactual fairness, which was introduced by the same authors. Similarly, Kilbertus et al. [24] also stressed the need for causal reasoning in fairness by framing the difference with statistical approaches as shifting the attention from “What is the right fairness criterion” to “What do we want to assume about our model of the causal data generating process”. They then focus on two fairness notions that correspond to two natural causal criteria centered around the notion of interventions, namely, no unresolved discrimination and no proxy discrimination. These are mainly qualitative criteria that cannot be used to

quantify the level of discrimination. More recently, Binkyte et al. [25] summarized the reasons why causality is essential to address fairness into three arguments. First, reliably measure discrimination. Second, mediation analysis allows us to distinguish the different paths through which discrimination can flow. Third, to establish causal evidence in legal practice.

With the wide acceptance of the idea that causality is essential to appropriately address fairness, a large number of fairness metrics have been introduced. These fairness metrics fall into two categories. First, existing causal effect measures are used in fields where causality is historically a well-established approach, such as health sciences, social sciences, and economy. These measures, which can also be applied to study fairness and discrimination, include total effect (TE), average treatment effect (ATE), effect of treatment on treated (ETT), direct effect (NDE), path-specific effect (PSE), etc. and are detailed in the numerous contributions of Judea Pearl [7] and Donald Rubin [11]. Second, causal-based metrics were introduced specifically to assess fairness and discrimination. These include counterfactual fairness [8], Equality of Effort [26], counterfactual error rates [9], etc.. This survey deviates from previous contributions and fairness surveys by exclusively focusing on causal-based fairness metrics and compiling practical aspects of the applicability of those metrics on real scenarios (identifiability, estimation, etc.). It is important to mention, however, that across all the sections of the paper, it is assumed that the causal graph is available and is reliable. Hence, we do not consider the causal discovery step [27, 13, 28] in the fairness assessment pipeline.

### 3. The need for causality: an example

Consider the hypothetical example<sup>2</sup> of an automated system for deciding whether to fire a teacher at the end of the academic year. Deployed teacher evaluation systems have been suspected of bias in the past. For example, IMPACT is a teacher evaluation system used in the city of Washington DC and has been found to be unfair against teachers from minority groups [30, 31, 32]. Assume that the system takes as input two features, namely, the location of the school where the teacher is working ( $C$ ) and the initial<sup>3</sup> average level of the students in her class ( $A$ ). The outcome is whether to fire the teacher ( $Y$ ). Assume also that all 3 variables are binary with the following values: if the school is located in a high-income neighborhood,  $C = 1$ , otherwise (the school is located in a low-income neighborhood),  $C = 0$ . If the initial average score

---

<sup>2</sup>Inspired by the prior convictions example in [29].

<sup>3</sup>At the beginning of the academic year.

for the students assigned to the teacher is high,  $A = 1$ , otherwise (initial level is low),  $A = 0$ . Firing a teacher corresponds to  $Y = 1$ , while retaining her corresponds to  $Y = 0$ . The level of students in a given class can be influenced by several variables, but in this example, assume that it is only influenced by the location of the school; students in high-income neighborhoods are more advantaged and typically perform better in school.

Assume now that the automated decision system is suspected to be biased by the initial level of students assigned to the teacher. That is, it is claimed that the system will more likely fire teachers who have been assigned classes with low-level students at the beginning of the academic year, which is clearly unfair. The sensitive attribute, in this case, is the initial level of students assigned to the teacher ( $A$ ). For concreteness, consider the prediction system that yields the following conditional probabilities:

$$\begin{aligned} \mathbb{P}(Y = 1 \mid A = 1, C = 0) &= 0.02 & \mathbb{P}(A = 1 \mid C = 0) &= 0.2 \\ \mathbb{P}(Y = 1 \mid A = 1, C = 1) &= 0.0675 & \mathbb{P}(A = 1 \mid C = 1) &= 0.8 \\ \mathbb{P}(Y = 1 \mid A = 0, C = 0) &= 0.01 & \mathbb{P}(A = 0 \mid C = 0) &= 0.8 \\ \mathbb{P}(Y = 1 \mid A = 0, C = 1) &= 0.25 & \mathbb{P}(A = 0 \mid C = 1) &= 0.2 \end{aligned}$$

and that the dataset is collected from a population where schools are located with equal proportions in high-income and low-income neighborhoods, that is,  $\mathbb{P}(C = 1) = \mathbb{P}(C = 0) = 0.5$ . Assume also that the proportion of classes with a low initial average level of students is the same as the one with high average initial level of students, that is,  $\mathbb{P}(A = 1) = \mathbb{P}(A = 0) = 0.5$ . To keep the scenario simple, assume that the level of students  $A$  does not depend on any other feature except  $C$  and that the firing decision  $Y$  depends only on  $A$  and  $C$ .

A simple approach to check the fairness of the firing decision  $Y$  with respect to the sensitive attribute  $A$  is to contrast the conditional probabilities:  $\mathbb{P}(Y = 1 \mid A = 0)$  and  $\mathbb{P}(Y = 1 \mid A = 1)$  which quantify, respectively, the likelihood of firing a teacher given that she is assigned students with an initial low level versus and the likelihood of firing a teacher given that she is assigned students with an initial high-level class. Such probabilities can be computed as follows:

$$\begin{aligned} \mathbb{P}(Y = 1 \mid A = a) &= \sum_{c \in \{0,1\}} \mathbb{P}(Y = 1 \mid A = a, C = c) \\ &\quad \times \mathbb{P}(C = c \mid A = a) \end{aligned} \tag{1}$$

Hence,

$$\begin{aligned} \mathbb{P}(Y = 1 \mid A = 1) &= 0.02 \times 0.2 + 0.0675 \times 0.8 = 0.058 \\ \mathbb{P}(Y = 1 \mid A = 0) &= 0.01 \times 0.8 + 0.25 \times 0.2 = 0.058 \end{aligned}$$

As the values are equal, the rates of firing between teachers who were assigned low-level students and high-level students appear to be equal, and hence no discrimination is detected<sup>4</sup>. This conclusion is flawed because it doesn't consider the mechanism by which the observed data is generated. In particular, the location of the school in which the teacher is working influences both the initial level of students assigned to her as well as the decision to fire or retain her. The  $\mathbb{P}(A|C)$  distribution indicates that 80% of classes in low-income neighborhoods have students with low initial levels ( $\mathbb{P}(A = 0 | C = 0) = 0.8$ ) while 80% of classes in high-income neighborhoods have students with high initial levels ( $\mathbb{P}(A = 1 | C = 1) = 0.8$ ). The automated decision system is biased in this case because  $\mathbb{P}(Y = 1 | A = 0, C = 1)$ , the probability of firing a teacher in high-income neighborhoods who is assigned a class with a low initial level, is exceptionally high (0.25). Using simple conditional probabilities (Eq. 1) on this collected dataset fails to appropriately account for that bias because very few teachers in high-income neighborhoods are assigned low-level classes in this particular dataset ( $\mathbb{P}(A = 0 | C = 1) = 0.2$ ). In general, any statistical fairness notion that relies solely on correlation between variables will fail to detect such bias.

To avoid such misleading conclusions, the causal relationships between variables should be considered. Figure 1 illustrates the causal relations between the three variables of the above example where the location of the school  $C$  is a confounder. Based on such a causal graph, a firing decision system is fair if it is as likely to fire teachers in the following two hypothetical cases: (1) when *all teachers in the population are assigned students of low level on average*, and (2) when *all teachers in the population are assigned students of high level on average*. This is achieved using intervention ( $do()$  operator)<sup>5</sup> and allows to break the problematic dependence between  $A$  and  $C$ . The probabilities of firing a teacher in these two hypothetical cases are expressed as  $\mathbb{P}(Y_{A=0} = 1) = \mathbb{P}(Y = 1 | do(A = 0))$  and  $\mathbb{P}(Y_{A=1} = 1) = \mathbb{P}(Y = 1 | do(A = 1))$  respectively. In this simple graph, and assuming no other variable is used in the prediction, these probabilities can be computed as follows:

$$\mathbb{P}(Y_{A=a} = 1) = \sum_{c \in \{0,1\}} \mathbb{P}(Y = 1 | A = a, C = c) \times \mathbb{P}(C = c)$$

---

<sup>4</sup>This corresponds to statistical parity.

<sup>5</sup>Intervention and the  $do()$  operator will be explained further in Section 6.1.1.



Hence,

$$\begin{aligned}\mathbb{P}(Y_{A=1} = 1) &= 0.02 \times 0.5 + 0.0675 \times 0.5 = 0.0437 \\ \mathbb{P}(Y_{A=0} = 1) &= 0.01 \times 0.5 + 0.25 \times 0.5 = 0.13\end{aligned}$$

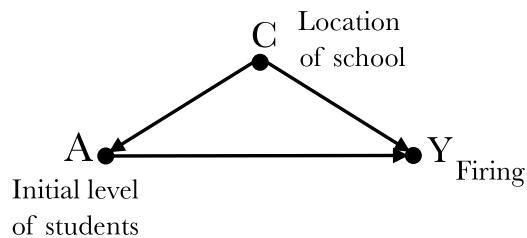


Figure 1: Causal graph of the firing example.

The values confirm the existence of a bias against teachers who are assigned classes with initially low levels.

#### 4. Preliminaries and Notation

Variables are denoted by capital letters. In particular,  $A$  is used for the sensitive variable (e.g., gender, race, age), and  $Y$  is used for the outcome of the automated decision system (e.g., health-care intervention, hiring, admission, releasing on parole). Small letters denote specific values of variables (e.g.,  $A = a'$ ,  $W = w$ ). Bold capital and small letters denote a set of variables and a set of values, respectively.

There are two fundamental frameworks to mathematically represent and characterize causal relations between variables: structural causal model [7] and potential outcome [11]. Formally, the two frameworks are equivalent [33, 34]. However, each of them is more equipped to address different problems in particular situations. For example, accounting for the many causal pathways that may exist in real applications can be more straightforward using SCM. On the other hand, a potential outcome framework is preferred when estimating individual-level causal effects. As most causal-based fairness notions are defined using one of these frameworks, the rest of this section introduces the terminology and the notation for both frameworks.

##### 4.1. Structural Causal Model (SCM) Framework

A structural causal model [7] is a tuple  $M = \langle \mathbf{U}, \mathbf{V}, \mathbf{F}, \mathbb{P}(\mathbf{U}) \rangle$  where:

- $\mathbf{U}$  is a set of exogenous variables that cannot be observed or experimented on but constitute the background knowledge behind the model.
- $\mathbf{V}$  is a set of observable variables that can be experimented on.
- $\mathbf{F}$  is a set of structural functions where each  $f_i$  is mapping  $\mathbf{U} \cup \mathbf{V} \rightarrow V_i$  which represents the process by which the variable  $V_i$  changes in response to other variables in  $\mathbf{U} \cup \mathbf{V}$ .
- $\mathbb{P}(\mathbf{u})$  is a probability distribution over the unobservable (latent) variables  $\mathbf{U}$ .

Various causal assumptions between variables are captured by a causal diagram  $G$ . These include causal Markov condition, causal faithfulness, and causal sufficiency. These assumptions together encode the same requirements as the SUTVA and Ignorability in the potential outcome framework, which is discussed in Section 4.2. A causal diagram  $G$  is a directed acyclic graph (DAG) where vertices represent variables, and directed edges represent functional relationships between the variables. Directed edges can have two interpretations. A probabilistic interpretation where the edge represents a dependency among the variables such that the direction of the edge is irrelevant. A causal interpretation where the edge represents a causal influence between the corresponding variables such that the direction of the edge matters. In the presence of a cause-effect relation between two variables  $A$  and  $Y$ , a confounder is a third variable  $C$ , which affects both the cause  $A$  and the effect  $Y$ . For example, the location of school variable  $C$  in Figure 1 is a confounder. Unobserved variables  $\mathbf{U}$ , which are typically not represented in the causal diagram, can be either mutually independent (Markovian model) or dependent on each other. In case the unobserved variables can be dependent and each  $U_i \in \mathbf{U}$  is used in at most two functions in  $F$ , the model is called semi-Markovian. In causal diagrams of semi-Markovian models, dependent unobservable variables (unobserved confounders) are represented by a dotted bi-directed edge between observable variables. Figure 2 shows causal graphs of Markovian model (Figure 2(a)), semi-Markovian model (Figures 2(b)) and semi-Markovian model after intervening on  $Z$  (Figure 2(c)).

An intervention noted  $do(V = v)$  is a manipulation of the model that consists of fixing the value of a variable (or a set of variables) to a specific value regardless of the corresponding function  $f_v$ . Graphically, it consists of discarding all edges incident to the vertex corresponding to variable  $V$ . Figure 2(c) shows the causal diagram of the manipulated model after intervention  $do(Z = z)$  denoted  $M_{Z=z}$  or  $M_z$  for short. The intervention  $do(V = v)$  induces a different distribution on the other variables. For example, in Figure 2(c),  $do(Z = z)$  results in a different distribution on  $Y$ , namely,  $\mathbb{P}(Y|do(Z = z))$ . Intuitively, while  $\mathbb{P}(Y|Z = z)$  reflects the population

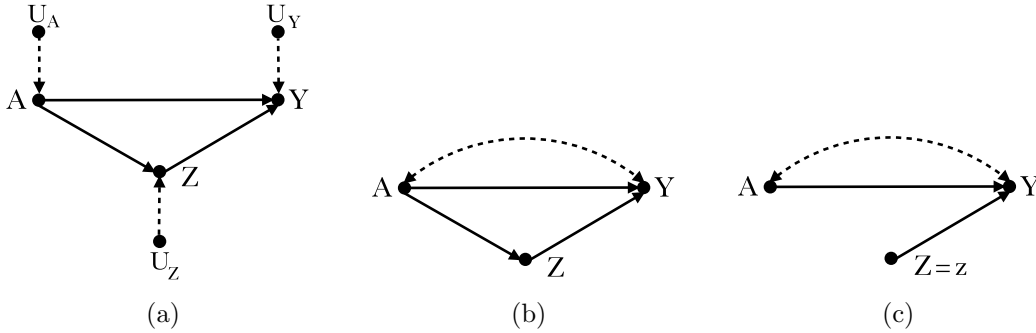


Figure 2: Markovian and semi-Markovian causal models.

distribution of  $Y$  among individuals whose  $Z$  value is  $z$ ,  $\mathbb{P}(Y|do(Z = z))$  reflects the population distribution of  $Y$  if *everyone in the population* had their  $Z$  value fixed at  $z$ . The obtained distribution  $\mathbb{P}(Y|do(Z = z))$  can be considered as a *counterfactual* distribution since the intervention forces  $Z$  to take a value different from the one it would take in the actual world. Such counterfactual variable is noted  $Y_{Z=z}$  or  $Y_z$  for short<sup>6</sup>. The term counterfactual quantity is used for expressions that involve explicitly multiple worlds. In Figure 2(b), consider the expression  $\mathbb{P}(y_{a'}|Y = y, A = a) = \mathbb{P}(y_{a'}|y, a)$ . Such expression involves two worlds: an observed world where  $A = a$  and  $Y = y$  and a counterfactual world where  $Y = y$  and  $A = a'$  and it reads “the probability of  $Y = y$  had  $A$  been  $a'$  given that we observed  $Y = y$  and  $A = a$ ”. In the common example of job hiring<sup>7</sup>, if  $A$  denotes race ( $a$  :white,  $a'$ :non-white) and  $Y$  denotes the hiring decision ( $y$ :hired,  $y'$ :not hired),  $\mathbb{P}(y_{a'}|y, a)$  reads “given that a white applicant has been hired, what is the probability that the same applicant is still being hired had he been non-white”. Nesting counterfactuals can produce complex expressions. For example, in the relatively simple model of Figure 2(b),  $\mathbb{P}(y_{a,z_{a'}}|y'_{a'})$  reads the probability of  $Y = y$  had (1)  $A$  been  $a'$  and (2)  $Z$  been  $z$  when  $A$  is  $a'$ , given that an intervention  $A = a'$  produced  $y'$ . This expression involves three worlds: a world where  $A = a$ , a world where  $Z = z_{a'}$ , and a world where  $A = a'$ . Such complex expressions are used to characterize direct, indirect, and path-specific effects.

Causal-based discrimination discovery aims to determine if the outcome of automated decision-making is fair or discriminative. Several causal-based fairness no-

<sup>6</sup>The notations  $Y_{Z \leftarrow z}$  and  $Y(z)$  are used in the literature as well.  $\mathbb{P}(Y = y|do(Z = z)) = \mathbb{P}(Y_{Z=z} = y) = \mathbb{P}(Y_z = y) = \mathbb{P}(y_z)$  is used to define the causal effect of  $z$  on  $Y$ .

<sup>7</sup>As a typical example, we use a simple job hiring use case where the sensitive attribute is the gender and the outcome is the hiring decision.

tions are defined in the literature (Section 5) and expressed in terms of joint, conditional, interventional, and counterfactual probabilities. Applying a fairness notion requires inputting a dataset  $D$  and a causal graph  $G$ . While joint probabilities (e.g.,  $\mathbb{P}(X = x, Y = y, Z = z)$ ) and conditional probabilities (e.g.,  $\mathbb{P}(Y = y|X = x)$ ) can be trivially estimated from the dataset  $D$ , probabilities involving interventions or counterfactuals cannot always be estimated from  $D$  and  $G$ . When a probability can be estimated from observable data ( $D$ ), it is said to be *identifiable*. Otherwise, it is *unidentifiable*. More formally, let  $M_1$  and  $M_2$  be two causal models sharing the same causal graph (not including the unobservable variable  $\mathbf{U}$ ), and the same set of probability distributions  $\psi$ , a quantity  $Q$  (e.g., intervention or counterfactual) is identifiable using  $\psi$  (noted  $\psi$ -identifiable) if the value of  $Q$  is unique and computable from  $\psi$  in any models  $M_1$  and  $M_2$ . In other words, if there exists two models  $M_1$  and  $M_2$  sharing the same graph structure and the same probability distributions, but yielding different  $Q$  values, then  $Q$  is unidentifiable. Typically, the identifiability of interventional and counterfactual quantities depends on the structure of the graph, in particular, the location of the unobserved confounding variables. Identifiability criteria are summarized in Section 6.1.

#### 4.2. Potential Outcome

Unlike the SCM framework, expressing causal relations in the potential outcome framework starts at the unit level. A unit  $i$  is the atomic research object. Fairness problems typically refer to an individual. For example, every candidate corresponds to a unit  $i$  in the job hiring scenario. Using the same example, the sensitive attribute of the candidate (e.g., gender) corresponds to the treatment in the potential outcome terminology. Given an outcome random variable  $Y$ , applying a treatment  $A = a$  on a unit  $i$  yields a different random variable called the potential outcome  $Y_i(A = a) = Y_i^a$ . For example, if  $A = 0$  refers to male,  $A = 1$  refers to female, and  $Y$  is the hiring decision,  $Y_i^1$  is the potential hiring decision of unit  $i$  when the gender (treatment) is female. Consequently, if the treatment variable  $A$  is binary, there are two potential outcomes  $Y_i^0$  and  $Y_i^1$ . In observational studies (in contrast to experimental studies), only one potential outcome can be observed, which is the factual outcome. The other potential outcome is usually impossible to observe and is called the counterfactual outcome. For example, if a job candidate  $i$  is female ( $A = 1$ ) and is not hired, the potential outcome  $Y_i^1$  is observed and is equal to 0. However, the potential outcome of that candidate  $i$  had she be male  $Y_i^0$  is impossible to observe because this requires going back in time (impossible) and changing the sex of that individual to male (not ethical in the cases where it is possible).

Causal inference in the potential outcome framework typically relies on three

assumptions: SUTVA, ignorability, and positivity [11]. SUTVA (Stable Unit Treatment Value Assumption) has two requirements: the absence of interference among units. In the job hiring example, it means that the hiring decision for a candidate is independent of the hiring decisions of all other candidates. Second, there is only one version of the treatment. This is more relevant in medical scenarios when a treatment (medication) has different versions (e.g., different dosages). For fairness scenarios, this requirement is typically satisfied as the treatment generally corresponds to an individual’s intrinsic attribute (e.g., gender, race, etc.). Ignorability is satisfied when the sensitive attribute  $A$  and the potential outcome variables  $Y^0$  and  $Y^1$  are independent given observable variables  $X$ . That is,  $A \perp Y^0, Y^1 | X$ <sup>8</sup>. This corresponds to the absence of hidden (unobservable) confounders. In the SCM framework, it is equivalent to the graphical Markovian model requirement. Positivity assumption requires that the sensitive attribute is not deterministic with respect to other observable variables. That is,  $\mathbb{P}(A = a | X = x) > 0$ ,  $\forall a$ , and  $x$ . In the job hiring example, any candidate can have any values of variables regardless of the gender  $A$ .

#### 4.3. SCM and graphical models vs. potential outcome

Although both causal frameworks are considered equivalent [33], interesting differences exist between them. Depending on the task at hand, one framework might be more appropriate to use than the other. For example, reasoning about causal effects at the individual (unit) level is more straightforward with the potential outcome framework [34] (Section 3.4). On the other hand, considering the different paths of causal effects (direct, indirect, and spurious) is much easier to achieve using SCMs and causal graphs. More generally, a potential outcome framework is more suitable for causal inference problems where the goal is to narrowly estimate the causal (treatment) effect of a cause variable  $A$  on an outcome variable  $Y$ . There are two justifications for this point. First, developing estimators of causal effects and counterfactuals can be more straightforward using the potential outcome framework [35]. Second, the potential outcome framework can decompose the sources of inconsistency and bias into unaccounted-for baseline differences between individuals and treatment effect bias [34] (Section 3.4). SCMs and causal graphs, however, are more suitable in causal discovery problems where the goal is to learn the causal relations among a set of variables [36]. The potential outcome framework is not well equipped for such problems because the causal effect of variables other than the treatment (sensitive

---

<sup>8</sup>Strong ignorability is a stronger assumption requiring independence between the potential outcomes and any covariate  $X$  ( $X \perp Y^0, Y^1$ ).

attribute) is not defined.

## 5. Causality-based Fairness notions

Assume that the sensitive attribute  $A$  and the outcome  $Y$  are binary variables where  $A = a_0$  denotes the privileged group (e.g., male), typically considered as the reference in characterizing discrimination, and  $A = a_1$  the disadvantaged group (e.g., female).

Whenever needed, the simple job hiring example will be used where  $A$  is the sensitive attribute corresponding to the gender ( $A = 0$  for male and  $A = 1$  for female),  $C$  is a covariate corresponding to the job type ( $C = 0$  for flexible schedule job and  $C = 1$  for non-flexible job schedule), and  $Y$  is the outcome corresponding to the hiring decision ( $Y = 0$  for not-hired and  $Y = 1$  for hired). Table 1 is an example dataset corresponding to this scenario.

Table 1: A job hiring example with 24 applications.  $A$  is the gender (sensitive attribute) where  $A = 1$ : female,  $A = 0$ : male.  $C$  is the job type where  $C = 0$ : flexible time job,  $C = 1$ : non-flexible time job.  $Y$  is the hiring decision (outcome) where  $Y = 0$ : not-hired,  $Y = 1$ : hired.

Female applicants (Treatment group)				Male applicants (Control Group)			
$i$	$A$	$C$	$Y$	$i$	$A$	$C$	$Y$
1:	1	0	1	13:	0	0	1
2:	1	0	1	14:	0	0	0
3:	1	0	0	15:	0	0	0
4:	1	0	0	16:	0	0	0
5:	1	0	0	17:	0	1	1
6:	1	0	0	18:	0	1	1
7:	1	0	0	19:	0	1	1
8:	1	0	0	20:	0	1	1
9:	1	1	1	21:	0	1	0
10:	1	1	1	22:	0	1	0
11:	1	1	1	23:	0	1	0
12:	1	1	0	24:	0	1	0

The most common non-causal fairness notion is total variation (TV), known as statistical parity, demographic parity, or risk difference. The total variation of  $A = a_1$  on the outcome  $Y = y$  with reference  $A = a_0$  is defined using conditional probabilities as follows:

$$\text{TV}_{a_1, a_0}(y) = \mathbb{P}(y | a_1) - \mathbb{P}(y | a_0) \quad (2)$$

Intuitively,  $\text{TV}_{a_1, a_0}(y)$  measures the difference between the conditional distributions of  $Y$  when we (passively) observe  $A$  changing from  $a_0$  to  $a_1$ . In the example of Table 1:

$$\text{TV} = \mathbb{P}(Y = 1 \mid A = 0) - \mathbb{P}(Y = 1 \mid A = 1) = \frac{5}{12} - \frac{5}{12} = 0.$$

So according to TV, the predicted hiring decision is fair. The main limitation of TV is its purely statistical nature, which makes it unable to reflect the causal relationship between  $A$  and  $Y$ ; that is, it is insensitive to the mechanism by which data is generated and collected. Total effect (TE) [7] is the causal version of TV and is defined in terms of experimental probabilities as follows:

$$\text{TE}_{a_1, a_0}(y) = \mathbb{P}(y_{a_1}) - \mathbb{P}(y_{a_0}) \quad (3)$$

TE measures the effect of the change of  $A$  from  $a_1$  to  $a_0$  on  $Y = y$  along all the causal paths from  $A$  to  $Y$ . Intuitively, while TV reflects the difference in proportions of  $Y = y$  in the current cohort, TE reflects the difference in proportions of  $Y = y$  in the entire population. For the binary outcome case, TE is equivalent to the average treatment effect (ATE) [34] in the potential outcome framework, which is defined as follows:

$$\text{ATE}_{a_1, a_0} = \mathbb{E}[Y^{a_1} - Y^{a_0}] \quad (4)$$

$$= \frac{1}{n} \sum_{i=1}^n (Y_i^{a_1} - Y_i^{a_0}) \quad (5)$$

where  $n$  is the number of observed samples. It is important to mention that (5) is an estimation of (4) in the available data. Note also that ATE corresponds exactly to FACE in [37].

Computing exactly ATE requires the knowledge of both potential outcomes: the observed and the counterfactual. As the latter is almost impossible to observe, the exact computation of ATE is typically impossible. However, for the sake of illustration, we assume the counterfactual outcome is available and show how ATE is computed. Later sections will show how ATE and counterfactual outcomes can be estimated from observable data. Table 2, shows the same job hiring dataset, but with counterfactual outcomes.

ATE is computed by considering the average potential outcome if the gender is female  $A = 1$ , that is,  $\frac{1}{n} \sum_{i=1}^n (Y_i^1)$  and the same if the gender is male  $A = 0$ ,  $\frac{1}{n} \sum_{i=1}^n (Y_i^0)$ . The former ( $\sum_{i=1}^n (Y_i^1)$ ) corresponds to the average of the observed outcomes ( $Y$ ) of samples 1 to 12 and counterfactual outcomes ( $Y^{cf}$ ) of samples 13

Table 2: The job hiring example with counterfactual outcomes.  $A^{cf}$  denotes the candidate's gender in the counterfactual world.  $Y^{cf}$  denotes the counterfactual potential outcome.

Female applicants (Treatment group)						Male applicants (Control Group)					
$i$	$A$	$C$	$Y$	$A^{cf}$	$Y^{cf}$	$i$	$A$	$C$	$Y$	$A^{cf}$	$Y^{cf}$
1:	1	0	1	0	1	13:	0	0	1	1	1
2:	1	0	1	0	0	14:	0	0	0	1	1
3:	1	0	0	0	1	15:	0	0	0	1	0
4:	1	0	0	0	0	16:	0	0	0	1	0
5:	1	0	0	0	0	17:	0	1	1	1	1
6:	1	0	0	0	0	18:	0	1	1	1	1
7:	1	0	0	0	0	19:	0	1	1	1	1
8:	1	0	0	0	0	20:	0	1	1	1	1
9:	1	1	1	0	1	21:	0	1	0	1	1
10:	1	1	1	0	1	22:	0	1	0	1	0
11:	1	1	1	0	0	23:	0	1	0	1	0
12:	1	1	0	0	0	24:	0	1	0	1	0

to 24, which gives  $\frac{12}{24} = \frac{1}{2}$ . Similarly, the average potential outcome if gender is male corresponds to the counterfactual outcomes of samples 1 to 12 and the observed outcomes of samples 13 to 24 which gives  $\frac{9}{24} = \frac{3}{8}$ . Hence,  $ATE = \frac{1}{2} - \frac{3}{8} = \frac{1}{8}$  which indicates a positive bias for female.

Computing the causal effect based only on the observed treatment group samples (e.g., female applicants only) corresponds to a variant of TE called the effect of treatment on the treated (ETT) [7] and is defined as:

$$ETT_{a_1, a_0}(y) = \mathbb{P}(y_{a_1} | a_1) - \mathbb{P}(y_{a_0} | a_1) \quad (6)$$

In the binary outcome case, ETT corresponds to the average treatment effect on the treated ATT [34] in the potential outcome framework defined as:

$$ATT_{a_1, a_0} = \mathbb{E}[Y^{a_1} | A = a_1] - \mathbb{E}[Y^{a_0} | A = a_1] \quad (7)$$

$$= \frac{1}{n_1} \sum_{i: A=a_1} (Y_i^{a_1} - Y_i^{a_0}) \quad (8)$$

where  $n_1$  is the number of samples in the treatment group. ATT is also called FACT (fair on average causal effect on the treated) in [37]. In the example of Table 2, ATT corresponds to the difference between the average observable outcome ( $Y$ ) and the



average counterfactual outcome ( $Y^{cf}$ ) in samples 1 to 12, that is,  $ATT = \frac{5}{12} - \frac{4}{12} = \frac{1}{12}$ , which confirms the positive bias for female.

Average treatment effect on the control group (ATC) [34] is the same as ATT but focusing instead on the control group:

$$ATC_{a_1, a_0} = \mathbb{E}[Y^{a_1} | A = a_0] - \mathbb{E}[Y^{a_0} | A = a_0] \quad (9)$$

$$= \frac{1}{n_2} \sum_{i:A=a_0} (Y_i^{a_1} - Y_i^{a_0}) \quad (10)$$

where  $n_2$  is the number of samples in the control group. Using the example of Table 2,  $ATC = \frac{7}{12} - \frac{5}{12} = \frac{1}{6}$ . Conditional average treatment effect (CATE) [34] is defined similarly, but conditioning on some other covariate instead of the sensitive attribute  $A$ :

$$CATE_{a_1, a_0}(X = x) = \mathbb{E}[Y^{a_1} | X = x] - \mathbb{E}[Y^{a_0} | X = x] \quad (11)$$

$$= \frac{1}{n_x} \sum_{i:X=x} (Y_i^{a_1} - Y_i^{a_0}) \quad (12)$$

where  $n_x$  is the number of samples in the subgroup  $X = x$ . Using the covariate  $C = 0$  (flexible schedule jobs) in the hiring example of Tabel 2,  $CATE(C = 0) = \frac{4}{12} - \frac{3}{12} = \frac{1}{12}$ , which is again confirming hiring decisions in favor of female. The choice of the variable  $X$  on which to condition depends on the role of the variable in the causal graph. It is common to condition on a mediator variable which can be used to justify the discrimination. The choice of the conditioning variable can be found in [38].

Unlike the SCM framework, in the potential framework, it is possible to define individual treatment effect ITE [34], which is defined for every unit  $i$  as:

$$ITE_{a_1, a_0}(i) = Y_i^{a_1} - Y_i^{a_0} \quad (13)$$

For instance, in Table 2,  $ITE(i = 3) = 0 - 1 = -1$  which indicates a discrimination against the female applicant  $i = 3$ . ATC, CATE, and ITE are defined and typically used in the potential outcome framework but have no equivalents in the SCM framework. However, although it is possible to defined equivalents of ATC and CATE in the SCM formalism, ITE cannot be formalized in the SCM framework.

The job hiring example of Tables 1 and 2 is interesting because it illustrates a statistical anomaly where some statistical notions such as TV fail to appropriately account for the bias between sub-populations (e.g., female vs male). Notice first that, according to the collected data, both female and male candidates are hired

at the same rate  $\frac{5}{12}$ . Notice also that if the hiring rates are adjusted according to the job type, female candidates are hired at an equal or higher rate for both types of jobs: for flexible schedule jobs ( $C = 0$ ), the hiring rates are the same  $\frac{1}{4}$  and for non-flexible jobs ( $C = 1$ ), the hiring rates are  $\frac{3}{4}$  for female and  $\frac{4}{8} = \frac{1}{2}$  for male. The explanation for such a counter-intuitive result is that most female candidates (8 out of 12) are applying for flexible schedule jobs (for family reasons), in which hiring is more difficult. On the other hand, few male candidates (4 out of 12) are applying for flexible schedule jobs and instead massively applying for the more accessible non-flexible jobs (8 out of 12 applicants). To appropriately assess discrimination in this case, there is a need to adjust on the job type variable  $C$ , that is, assessing discrimination for each job type separately. This simple job hiring scenario is similar to the Berkeley sex discrimination in college admission [4] where data showed a bias for male applicants overall, but when results were analyzed separately for each department, data showed a slight bias in favor of female candidates. The Berkeley scenario is typically used as an example of Simpson’s paradox [3]. In both scenarios, by considering the outcome of the observable samples in the counterfactual setup, the above causal-based fairness notions could appropriately assess gender ( $A$ ) discrimination on the outcome ( $Y$ ). The job hiring example illustrating the statistical anomaly can be easily modified to reflect Simpson’s paradox [3]. Table 3 shows the same example but with 30 observed samples. In such cohort,  $TV = -\frac{1}{15}$  indicates discrimination against female applicants. However, all causal notions ( $TE = ATE = \frac{2}{15}$ ,  $ATT = \frac{2}{15}$ , and  $ATC = \frac{2}{15}$ ,  $CATE(C = 0) = \frac{2}{15}$ , and  $CATE(C = 1) = \frac{2}{15}$ ) are indicating a bias in favor of female.

All the above causal-based fairness notions fall into the framework of disparate impact [6], which aims at ensuring the equality of outcomes among all groups (protected/treatment and unprotected/control). An alternative framework is the disparate treatment [6], which seeks equality of treatment achievable by prohibiting the sensitive attribute’s use in the decision process. The main idea is to split the causal effect between the sensitive attribute  $A$  and the outcome  $Y$  into several causal pathways, each of which is either fair, unfair, or spurious. Common fairness notions from the disparate treatment framework include direct effect, indirect effect, and path-specific effect [39]. An effect can be deemed fair, unfair, or spurious by an expert of the scenario at hand. An unfair effect is called discrimination. Direct discrimination is assessed using causal effect along the direct edge from  $A$  to  $Y$ . Indirect discrimination is measured using the causal effect along causal paths that pass through proxy attributes<sup>9</sup>. A fair or explainable discrimination is measured using causal pathways

---

<sup>9</sup>A proxy is an attribute that cannot be objectively justified if used in the decision-making

Table 3: The job hiring example with a Simpson’s paradox.

Female applicants (Treatment group)						Male applicants (Control Group)					
$i$	$A$	$C$	$Y$	$A^{c_j}$	$Y^{c_j}$	$i$	$A$	$C$	$Y$	$A^{c_j}$	$Y^{c_j}$
1:	1	0	1	0	1	16:	0	0	1	1	1
2:	1	0	1	0	1	17:	0	0	0	1	1
3:	1	0	1	0	0	18:	0	0	0	1	0
4:	1	0	0	0	0	19:	0	0	0	1	0
5:	1	0	0	0	0	20:	0	0	0	1	0
6:	1	0	0	0	0	21:	0	1	1	1	1
7:	1	0	0	0	0	22:	0	1	1	1	1
8:	1	0	0	0	0	23:	0	1	1	1	1
9:	1	0	0	0	0	24:	0	1	1	1	1
10:	1	0	0	0	0	25:	0	1	1	1	1
11:	1	1	1	0	1	26:	0	1	1	1	1
12:	1	1	1	0	1	27:	0	1	1	1	1
13:	1	1	1	0	1	28:	0	1	0	1	1
14:	1	1	1	0	0	29:	0	1	0	1	0
15:	1	1	0	0	0	30:	0	1	0	1	0

passing through explaining variables. The spurious effect corresponds to a pathway starting with an incident edge into the sensitive attribute  $A$ .

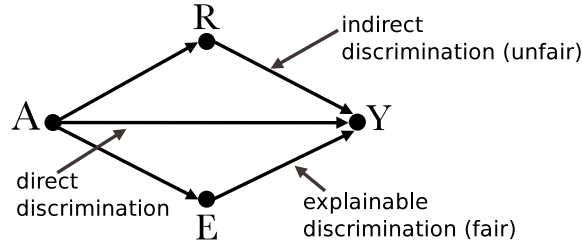


Figure 3: Job hiring scenario where  $A$  is gender,  $Y$ : hiring decision,  $R$ : hobby of a candidate ( $R = 1$  for mechanical hobby,  $R = 0$  for non-mechanical hobby), and  $E$ : education level of the candidate ( $E = 1$  for college degree,  $E = 0$  for no college degree).

Figure 3 presents a causal graph of the job hiring scenario involving an explaining variable  $E$  (e.g., education and academic degrees) and a proxy/redlining variable  $R$

process. It is also known as a redlining attribute.

(e.g., the hobby of the candidate). Hiring discrimination due to education level is legitimate and considered fair. In contrast, discrimination due to the hobby of the candidate is unfair as it is a proxy for gender (the type of hobby generally indicates the candidate’s gender). Direct effect can be computed by simply “blocking” all indirect causal paths. An indirect causal path is a directed path from  $A$  to  $Y$  going through one or several mediator variables. For example, in Figure 3, there are two indirect causal paths  $A \rightarrow R \rightarrow Y$  and  $A \rightarrow E \rightarrow Y$ . To compute the direct causal effect ( $A \rightarrow Y$ ), both indirect causal paths need to be blocked by adjusting on variables  $R$  and  $E$ . As there are no confounders, the direct effect can be computed as:

$$\begin{aligned} \text{DE}_{a_1, a_0}(y) &= \mathbb{P}(y \mid a_1, R, E) - \mathbb{P}(y \mid a_0, R, E) \\ &= \sum_r \sum_e (\mathbb{P}(y \mid a_1, r, e) - \mathbb{P}(y \mid a_0, r, e)) \end{aligned}$$

In the presence of confounders (between  $A$  and  $Y$ , between  $R$  and  $Y$ , etc.), natural direct effect (NDE) [39] is a more general notion that measures the direct causal effect and is defined as:

$$\text{NDE}_{a_1, a_0}(y) = \mathbb{P}(y_{a_1, \mathbf{z}_{a_0}}) - \mathbb{P}(y_{a_0}) \quad (14)$$

Where  $\mathbf{Z}$  is the set of mediator variables and  $\mathbb{P}(y_{a_1, \mathbf{z}_{a_0}})$  is the probability of  $Y = y$  had  $A$  been  $a_1$  and had  $\mathbf{Z}$  been the value it would naturally take if  $A = a_0$ . That is,  $A$  is set to  $a_1$  in the single direct path  $A \rightarrow Y$  and is set to  $a_0$  in all other indirect paths ( $A \rightarrow R \rightarrow Y$  and  $A \rightarrow E \rightarrow Y$ ). To see how NDE is computed, consider the sample dataset in Table 4 corresponding to the causal graph in Figure 3. Similarly to the previous examples, we assume the counterfactual values are available (grayed columns). The cohort consists of 6 female candidates and 6 male candidates.  $Y^{cf}$  is the counterfactual potential outcome (the gender differs from the observed sample).  $E_0$  is the education level had the gender was male.  $R_0$  is the hobby of the candidate had the gender was male.  $Y_{1, E_0, R_0}$  is the hiring decision had (1) the gender was female and (2) the education and hobby were set to the values if the candidate was male. According to Equation 14,  $\text{NDE}_{1,0}(y = 1) = \mathbb{P}(y_{1, E_0, R_0}) - \mathbb{P}(y_0) = \frac{8}{12} - \frac{9}{12} = -\frac{1}{12}$  which indicates a direct discrimination against female candidates. Notice the following:

- For rows 7 to 12, the values of columns  $E_0$ ,  $R_0$ , and  $Y_{1, E_0, R_0}$  are equal to the values in columns  $E$ ,  $R$ , and  $Y^{cf}$ , respectively.
- $\mathbb{P}(y_0)$  is computed based on the values in rows 1 to 6 of column  $Y^{cf}$  and rows 7 to 12 of column  $Y$ .

Table 4: A job hiring scenario corresponding to the causal graph in Figure 3.

$i$	$A$	$E$	$R$	$Y$	$Y^c$	$E_0$	$R_0$	$Y_{1,E_0}$	$E_1$	$R_1$	$Y_{0,E_1}$	$Y_{1,E_0,R_1}$
1:	1	1	0	1	1	1	1	1	1	0	1	1
2:	1	1	0	1	1	1	1	1	1	0	1	1
3:	1	1	1	0	1	0	1	1	1	1	1	0
4:	1	0	0	1	1	1	0	1	0	0	1	1
5:	1	0	0	0	1	0	1	0	0	0	1	1
6:	1	0	0	0	0	0	0	0	0	0	0	0
7:	0	1	1	1	1	1	1	1	1	0	1	1
8:	0	1	0	1	1	1	0	1	1	0	1	1
9:	0	1	1	1	0	1	1	0	1	1	0	0
10:	0	0	1	1	1	0	1	1	1	0	1	1
11:	0	0	1	0	1	0	1	1	0	0	1	0
12:	0	0	1	0	0	0	1	0	0	0	1	1

Natural indirect effect (NIE) [39] measures the indirect effect of  $A$  on  $Y$  and is defined as:

$$\text{NIE}_{a_1, a_0}(y) = \mathbb{P}(y_{a_0, \mathbf{z}_{a_1}}) - \mathbb{P}(y_{a_0}) \quad (15)$$

In the example of Table 4,  $\text{NIE}_{1,0}(y = 1) = \mathbb{P}(y_{0, E_1, R_1}) - \mathbb{P}(y_0) = \frac{9}{12} - \frac{9}{12} = 0$

The problem with NIE is that it does not distinguish between the fair (explainable) and unfair (indirect discrimination) effects. Path-specific effect [7, 40, 41] is a more nuanced measure that characterizes the causal effect in terms of specific paths.

Given a path set  $\pi$ , the  $\pi$ -specific effect is defined as:

$$\text{PSE}_{a_1, a_0}^\pi(y) = \mathbb{P}(y_{a_1 | \pi, a_0 | \bar{\pi}}) - \mathbb{P}(y_{a_0}) \quad (16)$$

where  $\mathbb{P}(y_{a_1 | \pi, a_0 | \bar{\pi}})$  is the probability of  $Y = y$  in the counterfactual situation where the effect of  $A$  on  $Y$  with the intervention ( $a_1$ ) is transmitted along  $\pi$ , while the effect of  $A$  on  $Y$  without the intervention ( $a_0$ ) is transmitted along paths not in  $\pi$  (denoted by:  $\bar{\pi}$ ). Using the job hiring example of Figure 3, Eq. 16 can be used to assess only unfair discrimination which is transmitted through the direct path  $A \rightarrow Y$  and the indirect path  $A \rightarrow R \rightarrow Y$ . The third path  $A \rightarrow E \rightarrow Y$  transmits explainable (fair) discrimination and hence, should not be considered. Given  $\pi = \{A \rightarrow Y, A \rightarrow R \rightarrow Y\}$ ,  $\text{PSE}_{1,0}^\pi = \mathbb{P}(Y_{1, E_0, R_1}) - \mathbb{P}(y_0) = \frac{8}{12} - \frac{9}{12} = -\frac{1}{12}$  which indicates a discrimination against female candidates.

### 5.1. No unresolved discrimination

No unresolved discrimination [24] is a fairness notion that falls into the disparate treatment framework and focuses on the indirect causal effects from  $A$  to  $Y$ . No

unresolved discrimination is satisfied when no directed path from  $A$  to  $Y$  is allowed, except via a resolving (explaining) variable  $E$ . A resolving variable is any variable in a causal graph that is influenced by the sensitive attribute in a manner accepted as non-discriminatory. Figure 4 presents two alternative causal graphs for the job hiring example. The graph at the left exhibits unresolved discrimination along the heavy paths:  $A \rightarrow R \rightarrow Y$  and  $A \rightarrow Y$ . By contrast, the graph at the right does not exhibit any unresolved discrimination as the effect of  $A$  on  $Y$  is justified by the resolved variable  $E$ :  $A \rightarrow E \rightarrow Y$ .

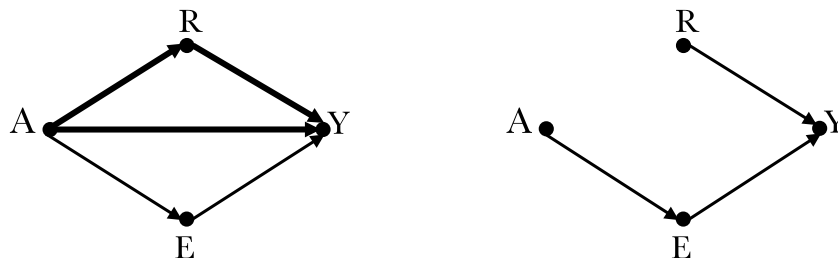


Figure 4:  $Y$  exhibits unresolved discrimination in the left graph (along the heavy paths) but not the right one.

The use of no unresolved discrimination in real scenarios is limited by the assumption of valid causal graph availability. [24] provide formal proof that even with prior knowledge of resolving variables, it is not always possible to tell, based on observational data only, if a predictor  $Y$  satisfies no unresolved discrimination.

### 5.2. No proxy discrimination

Similarly to no unresolved discrimination, no proxy discrimination [24] focuses on indirect discrimination. A causal graph exhibits potential proxy discrimination if a path exists from the protected attribute  $A$  to the outcome  $Y$  that is blocked by a proxy/redlining variable  $R$ . It is called proxy because it is used to decide about the outcome  $Y$  while it is a descendent of  $A$ , which is significantly correlated with it in such a way that using the proxy in the decision has almost the same impact as using  $A$  directly. An outcome variable  $Y$  exhibits no proxy discrimination if the equality:

$$\mathbb{P}(Y \mid do(R = r)) = \mathbb{P}(Y \mid do(R = r')) \quad \forall r, r' \in dom(R) \quad (17)$$

holds for any potential proxy  $R$ .

Figure 5 shows two similar causal graphs for the same job hiring example. The causal graph at the left presents potential proxy discrimination via the path:  $A \rightarrow$

$R \rightarrow Y$ . However, the graph at the right is free of proxy discrimination as the edge between  $A$  and its proxy  $R$  has been removed due to the intervention on  $R$  ( $R = r$ ).

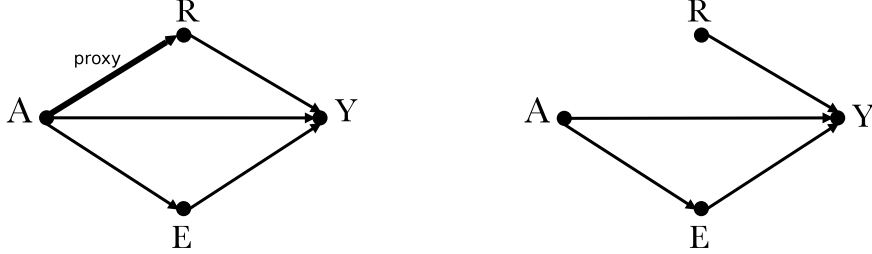


Figure 5: The graph at the left exhibits potential proxy discrimination (along the heavy edge between  $A$  and  $R$ ) but not in the right one.

Similarly to no unresolved discrimination, no proxy discrimination requires a valid causal graph. Hence, both fairness notions depend on the correct output of the causal discovery task.

### 5.3. Counterfactual fairness

Counterfactual fairness [8] is a very strong fairness notion requiring equality between the observed and counterfactual outcomes for every individual. That is, an outcome  $Y$  is counterfactually fair if under any assignment of values  $\mathbf{X} = \mathbf{x}$  and any individual in  $U$ ,

$$\mathbb{P}(y_{a_1}(U) \mid \mathbf{X} = \mathbf{x}, A = a_0) = \mathbb{P}(y_{a_0}(U) \mid \mathbf{X} = \mathbf{x}, A = a_0) \quad (18)$$

where  $\mathbf{X} = \mathbf{V} \setminus \{A, Y\}$  is the set of all remaining variables. As the latent variable  $U$  appears in Equation 18, counterfactual fairness is an individual fairness notion. It is satisfied if the probability distribution of the outcome  $Y$  is the same for every possible individual in the actual and counterfactual worlds. Counterfactual fairness typically coincides with ITE (Eq. 13).

[8] could test counterfactual fairness by making a very strong assumption. That is, they assumed the full structure of the causal model is available, including the latent variables  $U$ . They could then estimate the  $\mathbb{P}(U)$  distribution using Markov chain Monte Carlo methods and the observed data. Thanks to the estimated distribution of  $\mathbb{P}(U)$ , they could compute counterfactuals using Pearl’s three-step process: abduction, action, and prediction [7]. Hence, another sample with counterfactual sensitive value is generated for every individual in the population. Counterfactual fairness is finally assessed by comparing the density functions of the actual and counterfactual samples. Section 6.2 details the process of testing counterfactual fairness.

#### 5.4. Counterfactual Effects

By conditioning on the sensitive attribute  $A = a$ , Zhang and Bareinboim [9] defined two variants of NDE (Eq. 14) and NIE (Eq. 15) which focus on the direct and indirect effect for a specific group. In addition, they characterize a third type of effect, spurious, which considers the back-door paths between  $A$  and  $Y$ , that is, paths with an arrow into  $A$ .

The three effects are defined as follows:

$$\text{DE}_{a_1, a_0}(y|a) = \mathbb{P}(y_{a_1, \mathbf{z}_{a_0}}|a) - \mathbb{P}(y_{a_0}|a) \quad (19)$$

$$\text{IE}_{a_1, a_0}(y|a) = \mathbb{P}(y_{a_0, \mathbf{z}_{a_1}}|a) - \mathbb{P}(y_{a_0}|a) \quad (20)$$

$$\text{SE}_{a_1, a_0}(y) = \mathbb{P}(y_{a_0}|a_1) - \mathbb{P}(y|a_0) \quad (21)$$

where in Eq. 19 and 20,  $a$  can be  $a_0$  or  $a_1$ . Considering the simple job hiring example and focusing on the female group ( $A = 1$ ),  $\text{DE}_{1,0}(y|1)$  measures the change in the probability of  $Y$  (e.g., hiring) had  $A$  been 1 (female), while mediators  $E$  and  $R$  are kept at the level they would take had  $A$  been 0 (male). Using the values in Table 4,  $\text{DE}_{1,0} = \mathbb{P}(y_{1, E_0, R_0}|1) - \mathbb{P}(y_0|1) = \frac{4}{6} - \frac{5}{6} = -\frac{1}{6}$  which indicates a direct counterfactual discrimination against female. Similarly,  $\text{IE}_{1,0} = \mathbb{P}(y_{0, E_1, R_1}|1) - \mathbb{P}(y_0|1) = \frac{5}{6} - \frac{5}{6} = 0$  which indicates the absence of counterfactual indirect discrimination.  $\text{SE}_{1,0}(y)$  reads the change in the probability of hiring  $Y$  had  $A$  been 0 (male) for the female candidates with respect to the probability of hiring male candidates. Using Table 4,  $\text{SE}_{1,0}(y) = \mathbb{P}(y_0|1) - \mathbb{P}(y|0) = \frac{5}{6} - \frac{4}{6} = \frac{1}{6}$  which indicates a spurious effect in favor of female. Compared to NDE and NIE, counterfactual effects focus only on individuals of a specific group (e.g., only female candidates) and characterize the causal effect through spurious (back-door) paths. This spurious effect is what makes causal relations different from mere statistical correlations. However, counterfactual indirect effect IE still does not distinguish between fair and unfair direct effects.

#### 5.5. Counterfactual Error Rates

Equalized odds [2] is an important statistical fairness notion that requires equality of error rates (TPR and FPR) across sub-populations, that is,

$$\text{ER}_{a_1, a_0}(\hat{y}|y) = \mathbb{P}(\hat{y} | a_1, y) - \mathbb{P}(\hat{y} | a_0, y) \quad (22)$$

where  $\hat{y}$  denotes the prediction while  $y$  denotes the true outcome. The problem with this statistical notion is the difficulty in identifying the causes behind the discrimination, if any. [42] decomposes equalized odds (Eq. 22) using three counterfactual measures corresponding to the direct, indirect, and spurious effects of  $A$  on  $\hat{Y}$ . The three measures are counterfactual direct error rate, counterfactual indirect error rate,



and counterfactual spurious error rate. Let  $\hat{y} = f(\hat{\mathbf{p}}\mathbf{a})$  be a classifier where  $\hat{\mathbf{P}}\mathbf{A}$  is the set of input features (parent variables of  $\hat{Y}$ ) for the classifier. The counterfactual error rates for a sub-population  $a, y$  (with prediction  $\hat{y} \neq y$ ) are defined as:

$$\text{ER}_{a_1, a_0}^d(\hat{y} | a, y) = \mathbb{P}(\hat{y}_{a_1, y, (\hat{\mathbf{P}}\mathbf{A} \setminus A)_{a_0, y}} | a, y) - \mathbb{P}(\hat{y}_{a_0, y} | a, y) \quad (23)$$

$$\text{ER}_{a_1, a_0}^i(\hat{y} | a, y) = \mathbb{P}(\hat{y}_{a_0, y, (\hat{\mathbf{P}}\mathbf{A} \setminus A)_{a_1, y}} | a, y) - \mathbb{P}(\hat{y}_{a_0, y} | a, y) \quad (24)$$

$$\text{ER}_{a_1, a_0}^s(\hat{y} | y) = \mathbb{P}(\hat{y}_{a_0, y} | a_1, y) - \mathbb{P}(\hat{y}_{a_0, y} | a_0, y) \quad (25)$$

For example, the counterfactual direct error rate (Eq. 23) measures the error rate (disparity between the true and the predicted outcome) in terms of the direct effects of the sensitive attribute  $A$  on the prediction  $\hat{Y}$ . In the job hiring example, considering the female sub-population that *should* be hired ( $A = 1$  and  $Y = 1$ ), it reads: for a female candidate that should be hired, how would the prediction  $\hat{Y}$  change had the candidate been a female ( $A$  been 1), while keeping all the other features  $\hat{\mathbf{P}}\mathbf{A} \setminus A$  at the level that they would attain had “she was male”, compared to the prediction  $\hat{Y}$  she would receive had “she was male” and should have been hired?

Table 5: A job hiring scenario for counterfactual direct error rate  $\text{ER}^d$  computation.  $E_{0,1}$  is a short version of  $E_{A=0, Y=1}$ .  $R_{0,1}$  means  $R_{A=0, Y=1}$ .  $\hat{Y}_{1,1, E_{0,1}, R_{0,1}}$  means  $\hat{Y}_{A=1, Y=1, E_{0,1}, R_{0,1}}$ .  $Y_{0,1}$  means  $Y_{A=0, Y=1}$ .

$i$	$A$	$E$	$R$	$\hat{Y}$	$Y$	$E_{0,1}$	$R_{0,1}$	$\hat{Y}_{0,1}$
1:	1	1	0	1	1	1	0	1
2:	1	1	0	1	1	1	0	1
3:	1	1	1	0	1	0	1	1
4:	1	0	0	1	1	1	0	0
5:	1	0	0	0	1	0	0	0
6:	1	0	0	0	0	0	0	0
7:	0	1	1	1	1	1	1	1
8:	0	1	0	1	1	1	0	1
9:	0	1	1	1	0	0	1	1
10:	0	0	1	1	1	0	1	0
11:	0	0	1	0	1	0	1	0
12:	0	0	1	0	0	0	1	0

Table 5 shows the values (observed and counterfactual) needed to compute counterfactual direct error rate  $\text{ER}^d$  for the female candidates that should be hired ( $A = 1$

and  $Y = 1$ ).

$$\begin{aligned} \text{ER}^d(\hat{Y} = 1|A = 1, Y = 1) &= \mathbb{P}(\hat{Y}_{A=1, Y=1, E_{0,1}, R_{0,1}}|A = 1, Y = 1) \\ &\quad - \mathbb{P}(\hat{Y}_{A=0, Y=1}|A = 1, Y = 1) \end{aligned}$$

where  $E_{0,1}$  is a short version of  $E_{A=0, Y=1}$  which refers to the education level of the candidate had “she” been male and hired.  $R_{0,1}$  means  $R_{A=0, Y=1}$  and indicates the hobby of the candidate had “she” been male and hired.  $\hat{Y}_{A=1, Y=1, E_{0,1}, R_{0,1}}$  reads the hiring decision had the candidate was female, hired, with education  $E_{0,1}$ , and hobby  $R_{0,1}$ .  $Y_{A=0, Y=1}$  reads the hiring decision had the candidate was male and hired. Using the values in Table 5 (rows 1 to 5 in the last two columns),  $\text{ER}^d(\hat{Y} = 1|A = 1, Y = 1) = \frac{4}{5} - \frac{3}{5} = \frac{1}{5}$  which indicates a higher direct error rate for the female group.

Interestingly, the statistical equalized odd error rate (Eq. 22) can be decomposed in terms of the three above causal-based error rates:

$$\text{ER}_{a_1, a_0}(\hat{y} | y) = \text{ER}_{a_1, a_0}^d(\hat{y} | a_0, y) - \text{ER}_{a_0, a_1}^i(\hat{y} | a_0, y) - \text{ER}_{a_0, a_1}^s(\hat{y} | y) \quad (26)$$

### 5.6. Individual direct discrimination

Individual direct discrimination [43] aims to discover direct discrimination at the individual level. It is based on situation testing [44], a legally grounded technique for analyzing discrimination at an individual level. It compares the individual with similar individuals from both groups (protected and unprotected). That is, for an individual  $i$  in question, find the  $k$  other individuals who are the most similar to  $i$  in the group  $A = a_0$  and  $k$  similar individuals from the group  $A = a_1$ . The first set is denoted as  $\mathbf{S}^+$  while the second is as  $\mathbf{S}^-$ . The target individual is considered as discriminated if the difference observed between the rate of positive decisions in  $\mathbf{S}^-$  and  $\mathbf{S}^+$  is higher than a predefined threshold  $\tau$  (typically 5%).

Causal inference is used to define the distance function  $d(i, i')$  required to select the elements of  $\mathbf{S}^-$  and  $\mathbf{S}^+$ . First, only attributes that are direct causes of the outcome should be considered in the computation of the distance. That is, based on the causal graph,  $\mathbf{Q} = Pa(Y) \setminus \{A\}$  denotes the set of variables that should be used in the distance function. Second, the function definition should consider the causal effect of each of the selected attributes ( $Q_k \in \mathbf{Q}$ ) on the outcome. In particular, for each variable  $Q_k$ ,  $CE(q_k, q'_k)$  measures the causal effect on the outcome when the value of  $Q_k$  changes from  $q_k$  to  $q'_k$  and is defined as:

$$\text{CE}(q_k, q'_k) = \mathbb{P}(y_{\mathbf{q}}) - \mathbb{P}(y_{q'_k, \mathbf{q} \setminus \{q_k\}}) \quad (27)$$

where  $(\mathbb{P}(y_{\mathbf{q}}))$  is the effect of the intervention that forces the set  $\mathbf{Q}$  to take the set of values  $\mathbf{q}$ , and  $(\mathbb{P}(y_{q'_k, \mathbf{q} \setminus \{q_k\}}))$  is the effect of the intervention that forces  $Q_k$  to take value  $q'_k$  and other attributes in  $\mathbf{Q}$  to take the same values as  $\mathbf{q}$ .

The two individual fairness notions mentioned above, namely, ITE (Equation 13) and counterfactual fairness (Section 5.3) rely on the counterfactual outcome to assess fairness for every individual. Individual direct discrimination drops this requirement and use instead the sets  $\mathbf{S}^-$  and  $\mathbf{S}^+$  composed of similar individuals in both groups. Hence, it can be considered as an estimation technique to circumvent the need for counterfactuals. However, the distance function between two individuals  $d(i, i')$  is unnecessarily complex; it is defined in terms of the causal effects of every covariate  $X$  on the outcome  $Y$ . These causal effects are re-computed each time the distance between two individuals is needed. Matching techniques in the potential outcome framework use much simpler distance metrics. Matching techniques are discussed in Section 6.3.2.

### 5.7. Non-Discrimination Criterion

Non-discrimination criterion [45] is a group fairness notion that aims to discover and quantify direct discrimination through the direct causal effect of  $A$  on  $Y$ . Recall that, given a causal graph  $G$ , a direct effect of  $A$  on  $Y$  is the causal effect through the edge  $A \rightarrow Y$ . The idea is to consider a modified graph  $G'$  where the edge in question ( $A \rightarrow Y$ ) is discarded. A *block set*  $\mathbf{Q}$  is a set of variables that blocks all causal effects from  $A$  to  $Y$  in the modified graph  $G'$ . Hence,  $A$  and  $Y$  are independent conditioning on  $\mathbf{Q}$  in  $G'$ , that is,  $(A \perp Y | \mathbf{Q})_{G'}$ . Hence, conditioning on the same variables  $\mathbf{Q}$ , any dependence between  $A$  and  $Y$  in  $G$  is due to the direct effect of  $A$  on  $Y$ , which indicates direct discrimination. This discrimination can be assessed using simply the risk difference [46]:

$$| \Delta P |_{\mathbf{q}} = | \mathbb{P}(y | a_1, \mathbf{q}) - \mathbb{P}(y | a_0, \mathbf{q}) | \quad (28)$$

where  $\mathbf{q}$  is a value assignment for the block set  $\mathbf{Q}$  and the absolute value to consider both positive and negative discriminations. No direct discrimination can be concluded if the risk difference is less than a threshold  $\tau$  for all combinations of values of all block sets, that is, Eq. 28 holds for each value assignment  $q$  of each block set  $\mathbf{Q}$ .

NDE (Equation 14) and counterfactual direct effect DE (Section 5.4) also focus on assessing the direct discrimination, but they both rely on nested counterfactual quantities that are not observable from data. Non-discrimination criterion circumvents this difficulty by using block sets and considering all combinations of values of these block sets. Similarly to individual direct discrimination, it can be considered an estimation technique to avoid dealing with counterfactual quantities. This approach, however, does not work in semi-markovian models as  $A$  and  $Y$  will never be independent in  $G'$  ( $(A \perp Y | \mathbf{Q})_{G'}$ ) because of hidden confounders.

### 5.8. Equality of Effort

Equality of effort [26] fairness notion identifies discrimination by assessing how much effort is needed by the disadvantaged individual/group to reach a certain level of outcome. A treatment variable  $T$  is selected and used to address the question: “to what extent should the treatment variable  $T$  change to make the individual (or a group of individuals) achieve a certain outcome level?”. Hence, this notion focuses on whether the effort to reach a certain outcome level is the same for the protected and unprotected groups. Considering the simple job hiring example, the education level  $E$  is a good choice for the treatment variable. Two equality of effort notions are defined based on the potential outcome framework, individual  $\gamma$ -Equal effort and system  $\gamma$ -Equal effort. Let  $Y_i^{(t)}$  be the potential outcome for individual  $i$  had  $T$  been  $t$  and  $\mathbb{E}[Y_i^{(t)}]$  be the expected outcome for individual  $i$ . Situation testing [44] is used to estimate the counterfactual potential outcome in a similar way as individual direct discrimination (Section 5.6). Let  $\mathbf{S}^+$  and  $\mathbf{S}^-$  be the two sets of similar individuals with  $A = a_0$  and  $A = a_1$ , respectively, and  $\mathbf{E}[Y_{\mathbf{S}^+}^{(t)}]$  be the expected outcome under treatment  $t$  for the subgroup of individuals  $\mathbf{S}^+$ . The minimal effort needed to achieve  $\gamma$ -level of outcome variable within the subgroup  $\mathbf{S}^+$  is defined as:

$$\Psi_{\mathbf{S}^+}(\gamma) = \operatorname{argmin}_{t \in T} \{\mathbb{E}[Y_{\mathbf{S}^+}^{(t)}] \geq \gamma\} \quad (29)$$

Individual  $\gamma$ -Equal effort is satisfied for individual  $i$  if:

$$\Psi_{\mathbf{S}^+}(\gamma) = \Psi_{\mathbf{S}^-}(\gamma) \quad (30)$$

System  $\gamma$ -Equal effort is satisfied for a sub-population (e.g.,  $A = a_1$ ) if:

$$\Psi_{\mathbf{D}^+}(\gamma) = \Psi_{\mathbf{D}^-}(\gamma) \quad (31)$$

where  $\mathbf{D}^+$  and  $\mathbf{D}^-$  are the subsets of the entire dataset with sensitive attributes  $a_0$  and  $a_1$ , respectively. Both criteria can be used to measure the effort discrepancy between protected and unprotected groups by considering the difference  $\Psi_{\mathbf{S}^+}(\gamma) - \Psi_{\mathbf{S}^-}(\gamma)$ . Unlike most causal-based fairness notions that intervene (*do* operator) on the sensitive attribute  $A$  ( $y_a, Y_i^a$ , etc.), equality of effort intervenes instead on a treatment variable  $T$  ( $Y_i^{(t)}$ ). The main limitation of the equality of effort notion is that a single treatment variable typically does not appropriately reflect the discrepancy between protected and unprotected groups.

### 5.9. Interventional and justifiable fairness

Interventional fairness [10] is a group-level fairness that can be seen as a strong version of total effect (Eq. 3). Instead of intervening only on the sensitive attribute

$A$ , interventional fairness intervenes on all remaining variables. Let  $\mathbf{K}$  be a subset of  $\mathbf{V}$  excluding  $A$  and  $Y$ , that is,  $\mathbf{K} \subseteq \mathbf{V} - \{A, Y\}$ . A predicting algorithm is  $\mathbf{K}$ -fair if for any assignment of values  $\mathbf{K} = \mathbf{k}$  and outcome  $Y = y$ :

$$\mathbb{P}(y_{a_1, \mathbf{k}}) = \mathbb{P}(y_{a_0, \mathbf{k}}) \quad (32)$$

A predicting algorithm is interventionally fair if it is  $\mathbf{K}$ -fair for every set of variables  $\mathbf{K}$ . Using the job hiring example of Figure 3, interventional fairness holds between male and female groups if  $\mathbb{P}(y_{1, E_u, R_v}) = \mathbb{P}(y_{0, E_u, R_v})$ ,  $\forall u, v \in 0, 1$ . The interventional fairness formula (Eq. 32) is similar to the non-discrimination criterion formula (Eq. 28). However, while Eq. 28 uses simple conditioning on  $A$  and covariates, Eq. 32 makes an intervention on  $A$  and all other covariates and hence works on Markovian as well as semi-Markovian models.

Justifiable fairness is a relaxation of interventional fairness achieved by classifying the variables as admissible (denoted as  $\mathbf{E}$ ) or inadmissible (denoted as  $\mathbf{R}$ ), which correspond, respectively, to explainable and proxy/redlining variables as defined previously. A predicting algorithm is justifiably fair if it is  $\mathbf{K}$ -fair with respect to only supersets of  $\mathbf{E}$ , that is,  $\mathbf{K} \supseteq \mathbf{E}$ . Hence, instead of intervening on all variables, it is enough to intervene on only admissible variables (or any superset of them). Graphically, suppose all directed paths from the sensitive attribute  $A$  to the outcome  $Y$  go through an admissible attribute in  $\mathbf{E}$ . In that case, the algorithm is justifiably fair, which typically coincides with no-unresolved discrimination (Section 5.1). Using the job hiring example (Figure 3), justifiable fairness holds if  $\mathbb{P}(y_{1, E_u}) = \mathbb{P}(y_{0, E_u})$ ,  $\forall u \in 0, 1$ . Notice that in case  $\mathbf{E} = \emptyset$ , justifiable fairness coincides with interventional fairness. Interestingly, being based solely on interventions, interventional and justifiable notions of fairness do not require the presence of the underlying causal model. The only assumption is the ability to distinguish admissible and inadmissible variables.

#### 5.10. Individual equalized counterfactual odds

Individual equalized counterfactual odds [47] is a stronger version of counterfactual fairness (Section 5.3) requiring, in addition, that the factual-counterfactual pair share the same value of the outcome  $Y$ . The aim is to have a counterfactual version of equalized odds [2]. This is achieved by conditioning both sides of Eq. 18 on the same outcome  $Y = y$ . A predictor satisfies individual equalized counterfactual odds if:

$$\mathbb{P}(\hat{y}_{a_1} \mid \mathbf{X} = \mathbf{x}, y_{a_1}, A = a_0) = \mathbb{P}(\hat{y}_{a_0} \mid \mathbf{X} = \mathbf{x}, y_{a_0}, A = a_0) \quad (33)$$

The only difference with Eq. 18 is the additional conditioning  $Y = y_{a_1}$  in the LHS and  $Y = y_{a_0}$  in the RHS. Counterfactual error rates are the only other causal-based

fairness notions considering the outcome  $Y$  (Section 5.4). However, unlike counterfactual error rates, individual equalized counterfactual odds requires intervention on  $Y$ . This is the only fairness criterion that requires intervention on the prediction  $\hat{Y}$  and on the actual outcome  $Y$ .

## 6. Computing causal quantities from observable data

Using causal-based fairness notions is challenging for two reasons. First, only the factual outcome can be observed among the two possible outcomes. The counterfactual outcome is usually impossible to observe (e.g., if the gender of a candidate is female (factual), it is impossible to observe the counterfactual outcome when the same candidate would have been male). Second, sensitive attributes (e.g., male and female) are typically not randomly assigned in observational data. Hence, the main difficulty in applying causal-based fairness notions is to compute and/or estimate the causal quantities (counterfactual outcomes, causal effects, counterfactual effects, etc.) using observational data. This includes all grayed columns in the simple toy datasets used in Section 5 as well as all fairness notions such as *ATE*, *ETT*, counterfactual fairness, etc. Each causal framework, namely, SCM with causal graphs and potential outcomes, uses a different approach to compute/estimate the causal quantities using observational data. The SCM framework relies mainly on the identifiability criterion to generate an expression for the causal quantity based only on observable probabilities. If the identifiability criterion is unsatisfied, the causal quantity can not be computed using the available observable data. In such case, as an alternative, if the complete structure of the causal model is available, it is possible to estimate the distribution of the latent variables  $U$  and consequently generate an estimation of the counterfactual outcomes. The potential outcome framework approximates causal quantities using several estimation techniques (e.g., matching, re-weighting, etc.). The following subsections illustrate the above three approaches: identifiability, estimation based on full causal model, and potential outcome estimation techniques.

### 6.1. Identifiability

The identifiability of causal quantities has been extensively studied in the literature: causal effect (intervention) identifiability [48, 49, 50, 51, 12, 52, 53, 7], counterfactual identifiability [54, 53, 55, 56], direct/indirect effects [39] and path-specific effect identifiability [57, 55, 58, 59, 60]. This section summarizes the main identifiability conditions as they relate to the specific problem of discrimination discovery.

### 6.1.1. Identifiability of causal effect (intervention)

The causal effect of a cause variable  $X$  on an effect variable  $Y$  is computed using  $\mathbb{P}(Y_x) = \mathbb{P}(Y|do(X = x))$ , the distribution of  $Y$  after the intervention  $X = x$ . In a discrimination setup, the cause is typically the sensitive attribute  $A$ . A basic case where identifiability can be avoided altogether is when it is possible to perform experiments by intervening on the sensitive attribute  $A$ . When this is possible, randomized controlled trial (RCT) [61] can be used to estimate the causal effect. RCT consists in randomly assigning subjects (e.g., individuals) to treatments (e.g., gender), then comparing the outcome  $Y$  of all treatment groups. However, in the context of machine learning fairness, RCT is often not an option as experiments can be too costly to implement, physically impossible, or ethically not acceptable (e.g., changing the gender of a job applicant).

In Markovian models (no unobserved confounding), the causal effect is always identifiable (Corollary 3.2.6 in [7]). The simplest case is when there is no confounding between  $A$  and  $Y$  (Figure 6(a)). In that case, the causal effect matches the conditional probability regardless of any mediator:

$$\mathbb{P}(y_a) = \mathbb{P}(y|do(a)) = \mathbb{P}(y|a) \quad (34)$$

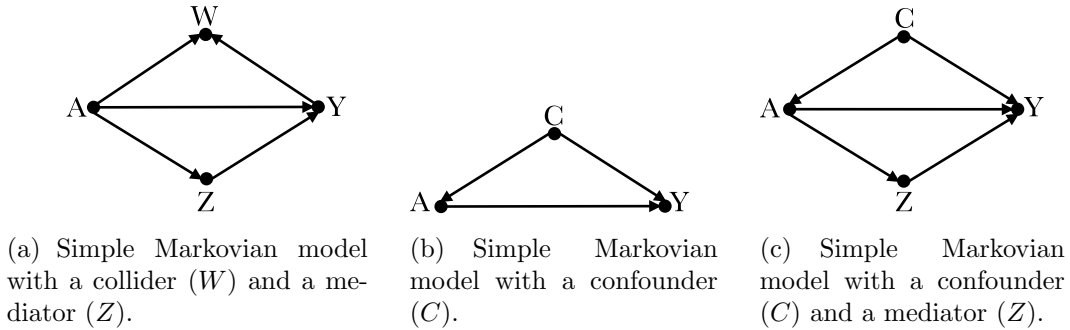


Figure 6: Simple causal graphs

In presence of an observable confounder (Figure 6(b)),  $\mathbb{P}(y_a)$  is identifiable by adjusting on the confounder:

$$\mathbb{P}(y_a) = \sum_C \mathbb{P}(y|a, c) \mathbb{P}(c) \quad (35)$$

where the summation is on values  $c$  in the domain (sample space) of  $C$  denoted as

$dom(C)$ . Eq. 35 is called the back-door formula<sup>10</sup>. For comparison, notice that the joint probability is

$$\mathbb{P}(y, a, c) = \mathbb{P}(y|a, c) \mathbb{P}(a|c) \mathbb{P}(c)$$

while the conditional probability is

$$\mathbb{P}(y|a) = \sum_C \mathbb{P}(y|a, c) \mathbb{P}(c|a)$$

For semi-Markovian models, identifiability of  $\mathbb{P}(y_a)$  is not guaranteed. In case it is identifiable, Pearl [7] proposes a *do*-calculus composed of three rules allowing to expression of interventional probabilities in terms of observational ones:

1.  $\mathbb{P}(y_a|\mathbf{z}, w) = \mathbb{P}(y_a|\mathbf{z})$  provided that the set of variables  $\mathbf{Z}$  blocks all backdoor paths from  $W$  to  $Y$  after all arrows leading to  $A$  have been deleted.
2.  $\mathbb{P}(y_a|\mathbf{z}) = \mathbb{P}(y|a, \mathbf{z})$  provided that the set of variables  $\mathbf{Z}$  blocks all backdoor paths from  $A$  to  $Y$ .
3.  $\mathbb{P}(y_a) = \mathbb{P}(y)$  provided that there are no causal paths between  $A$  and  $Y$ .

*do*-calculus has been proven to be sound and complete in identifying interventional distributions [52]. For example,  $\mathbb{P}(y_a)$  is identifiable in Figure 7(d). By applying the chain rule following the topological order:  $W_2 < A < W_1 < W_3 < Y$ , we get:

$$\begin{aligned} \mathbb{P}(y_a) &= \sum_{w_1 w_2 w_3} \mathbb{P}(y|do(a), w_1, w_2, w_3) \mathbb{P}(w_1|do(a), w_2) \mathbb{P}(w_2) \\ &\quad \times \mathbb{P}(w_3|w_2, w_1, do(a)) \end{aligned} \tag{36}$$

$$= \sum_{w_1 w_2} \mathbb{P}(y|do(a), w_1, w_2) \mathbb{P}(w_1|do(a), w_2) \mathbb{P}(w_2) \tag{37}$$

$$= \sum_{w_1 w_2} \mathbb{P}(y|do(a), w_2) \mathbb{P}(w_1|a, w_2) \mathbb{P}(w_2) \tag{38}$$

$$\begin{aligned} &= \sum_{w_1 w_2} \sum_{a'} \mathbb{P}(y|a', w_2, do(w_1)) \mathbb{P}(a'|do(w_1), w_2) \\ &\quad \times \mathbb{P}(w_1|a, w_2) \mathbb{P}(w_2) \end{aligned} \tag{39}$$

$$= \sum_{w'_1} \sum_{w'_2} \sum_{a'} \mathbb{P}(y|w'_1, w'_2, a') \mathbb{P}(a'|w'_2) \mathbb{P}(w'_1|w'_2, a) \mathbb{P}(w'_2) \tag{40}$$

Note that  $w_3$  is omitted from (37) since it is considered latent [62]. Applying Rule 2 followed by Rule 3 to the first term in (37) yields  $\mathbb{P}(y|do(a), w_2)$  (38). Likewise,

---

<sup>10</sup>Called also adjustment formula or stratification.



applying Rule 2 to the second term in (37) leads to  $\mathbb{P}P(w_1|a, w_2)$ . Thus, the original problem reduces to identifying the term  $\mathbb{P}(y|do(a), w_2)$  in (38). Here, we cannot apply Rule 2 to exchange  $do(a)$  with  $a$  because  $G_{\underline{A}}$  (graph obtained by removing all emanating arrows from  $A$ ) contains a backdoor path from  $A$  to  $Y$ . Thus, to block that path, we need to condition and sum over all values of  $A$  as shown in Eq. (39) ( $\sum_{a'} \mathbb{P}(y|a', w_2, do(w_1))\mathbb{P}(a'|do(w_1), w_2)$ ). Now, applying Rule 2 to  $\mathbb{P}(y|a', w_2, do(w_1))$  and Rule 3 to  $\mathbb{P}(a'|do(w_1), w_2)$  and adding the other terms results in the final expression in (40). The problem of  $do$ -calculus is the difficulty of determining the correct order of application of the rules. Using the wrong order may hinder the identifiability or produce a very complex expression [63]. As an alternative to using the  $do$ -calculus, several contributions in the identifiability literature focused on defining graphical patterns and mapping them to simple and concise intervention-free expressions [49, 50, 51, 12].

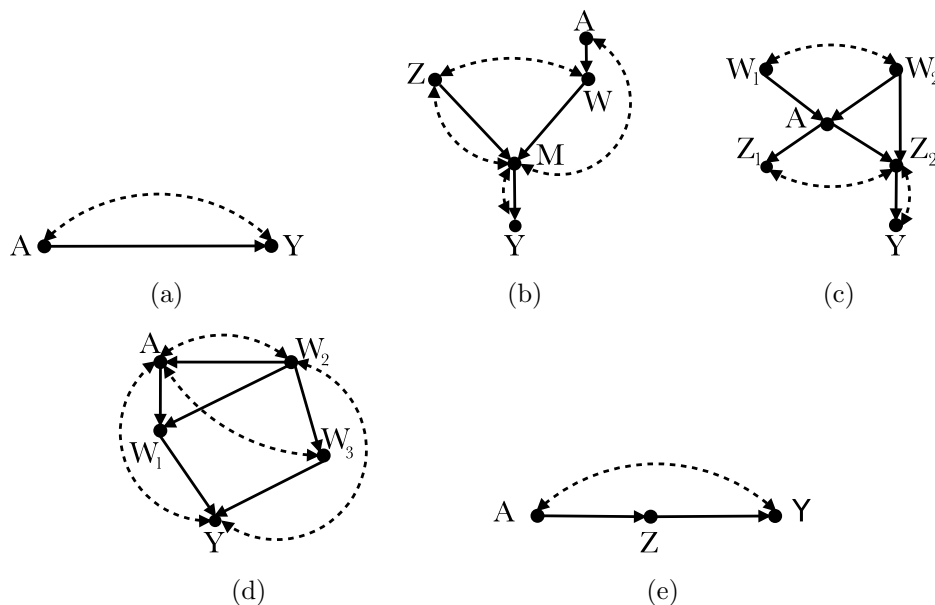


Figure 7: Figure 7(a) presents the “bow” graph, Figure 7(b) illustrates the structure of a  $c$ -tree, Figure 7(c) shows a semi-Markovian model where  $\mathbb{P}(y_a)$  is observable, Figure 7(d) presents a semi-Markovian model where  $\mathbb{P}(y_a)$  is identifiable and Figure 7(e) illustrates a simple example of the front-door criterion.

All graphical criteria can be generalized to the case where the sensitive attribute is not connected to any of its children through a confounding path. In such case,  $c$ -component factorization can be used. A  $c$ -component is a set of vertices in the graph such that every pair of vertices are connected by a confounding edge. The idea of  $c$ -

component factorization is to decompose the identification problem into smaller sub-problems, that is, a disjoint set of c-components in order to calculate  $\mathbb{P}(y_a)$ . For example, in Figure 7(c), there are three c-components:  $\{\{W1, W2\}, \{A\}, \{Z1, Z2, Y\}\}$ . Hence, as long as no confounding path connects  $A$  to any of its direct children,  $\mathbb{P}(y_a)$  is identifiable. C-component factorization is used in the ID algorithm [53] which is proven to be complete for causal effect identification.

In case there is an unobservable confounding between the sensitive attribute  $A$  and the outcome  $Y$ , all the above criteria will fail. However,  $\mathbb{P}(y_a)$  can still be identifiable using the front-door criterion. This criterion is satisfied in Figure 7(e) and consists in having a mediator variable  $Z$  such that:

- there are no backdoor paths from  $A$  to  $Z$
- all backdoor paths from  $Z$  to  $Y$  are blocked by  $A$ .

A backdoor path from  $A$  to  $Z$  is any path starting at  $A$  with a backward edge  $\leftarrow$  into  $A$  (e.g.,  $A \leftarrow \dots Z$ ). If such criterion is satisfied,  $\mathbb{P}(y_a)$  can be computed as follows:

$$\begin{aligned} \mathbb{P}(y_a) &= \sum_Z \mathbb{P}(y|do(z)) \mathbb{P}(z|do(a)) \\ &= \sum_Z \mathbb{P}(y|z, a) \mathbb{P}(a) \mathbb{P}(z|a) \end{aligned} \quad (41)$$

Shpitser and Pearl [12] proved that all the unidentifiable cases of the causal effect  $\mathbb{P}(y_a)$  boil down to a general graphical structure called the hedge criterion. Based on this criterion, they designed a complete identifiability algorithm called *ID*, which outputs the expression of  $\mathbb{P}(y_a)$  if it is identifiable or the reason for the unidentifiability otherwise.

The simplest graph in which the causal effect between  $A$  and  $Y$  is not identifiable is the “bow” graph (Figure 7(a)). This simple unidentifiability criterion can be generalized to a more complex graph called c-tree. A c-tree is a graph that is at the same time a tree<sup>11</sup> and a c-component. Figure 7(b) shows an example of a c-tree. If the causal graph is a c-tree rooted in the outcome variable  $Y$ ,  $\mathbb{P}(y_a)$  is unidentifiable [53].

### 6.1.2. Identifiability of counterfactuals

Most of the causal-based fairness notions in the disparate treatment framework (NDE (Eq. 14), path-specific effect (Eq. 16), counterfactual effects (Section 5.4),

---

<sup>11</sup>Notice that, in this paper, the direction of the arrows between nodes is reversed compared to the usual tree structure.

etc.) are defined in terms of counterfactual quantities. Hence, the applicability of those notions depends heavily on the identifiability of the counterfactuals composing them. In Markovian, as well as semi-Markovian models, if all parameters of the causal model are known (including  $\mathbb{P}(\mathbf{u})$ ), any counterfactual is identifiable and can be computed using the three steps abduction, action, and prediction (Theorem 7.1.7 in [7]).

Let  $P_* = \{P_{\mathbf{x}} | \mathbf{X} \subseteq \mathbf{V}, \mathbf{x} \text{ a value assignment of } \mathbf{X}\}$  be the set of all interventional distributions in a given causal model. While the identifiability of interventional probabilities  $\mathbb{P}(y_a)$  is characterized based on observational probabilities  $\mathbb{P}(\mathbf{v})$ , in this section, the identifiability of counterfactuals is characterized in terms of interventional probabilities  $P_*$ . Then, combining these results with the criteria of the previous section, a counterfactual can, in turn, be identified using observational probabilities  $\mathbb{P}(\mathbf{v})$ .

Given a causal graph  $G$  of a Markovian model and a counterfactual expression  $\gamma = v_x | e$  with  $e$  some arbitrary set of evidence, identifying and computing  $\mathbb{P}(\gamma)$  requires constructing a counterfactual graph which combines parallel worlds. Every world is represented by a model  $M_x$  corresponding to each subscript in the counterfactual expression. For example, given the causal graph in Figure 1 and the counterfactual expression  $y_{a_1} | a_0$ , the resulting counterfactual graph is shown in Figure 8(d). The counterfactual graph should be “reduced” by merging together vertices that share the same causal mechanism (**make-cg** algorithm in [53] automates this procedure). The resulting counterfactual graph can be considered as a typical causal graph for a larger causal model. Consequently, all the graphical criteria listed in Section 6.1.1 for the identifiability of causal effects apply on the counterfactual graph to identify counterfactual quantities, in particular, the c-component factorization of the counterfactual graph [54]. ID\* and IDC\* algorithms [53] automate the identifiability and computation of counterfactuals based on all the above criteria. Note that ID\* and IDC\* output expressions in terms of interventional probabilities  $P_*$ . Then, the ID algorithm is used to express those interventional probabilities in terms of observational probabilities.

The simplest unidentifiable counterfactual quantity is  $\mathbb{P}(y_{a'}, y'_a)$ , called the probability of necessity and sufficiency. Using the simple job hiring example, it reads as the probability of an individual being hired had he be female and being not hired had she be a male. is hiredThe corresponding counterfactual graph is the W-graph with the same structure as Figure 8(a). This simple criterion can be generalized to the zig-zag graph (Figure 8(b)) where the counterfactual  $\mathbb{P}(y_a, w_1, w_2, z')$  is not identifiable.

Pearl [7] proves two results about the identifiability of counterfactuals. First,

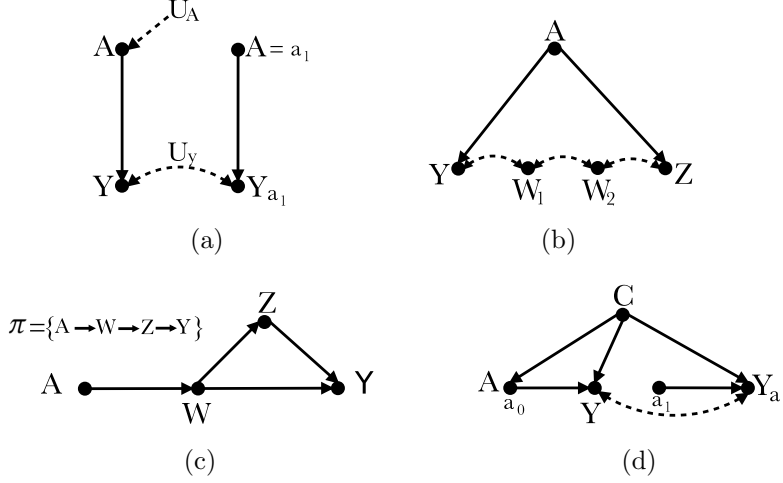


Figure 8: Causal graphs.

for linear causal models (i.e., the functions  $\mathbf{F}$  are linear), any counterfactual is experimentally (using  $P_*$ ) identifiable whenever the model parameters are identified. Second, in linear causal models, if some of the model parameters are unknown, any counterfactual of the form  $\mathbb{E}(Y_a|e)$  where  $e$  is some arbitrary set of evidence, is identifiable provided that  $\mathbb{E}(y_a)$  is identifiable. Finally, there is no single necessary and sufficient criterion for the identifiability of counterfactuals in semi-Markovian models [57].

**According to the toy example  $A=0$  corresponds to low level students.** To illustrate the computation of a counterfactual probability, consider the teacher firing example of Figure 1 and the counterfactual probability  $\mathbb{P}(y_{a_1}|a_0)$  which reads the probability of firing a teacher who is assigned a class with a high initial level of students ( $a_0$ ) had she been assigned a class with a low initial level of students ( $a_1$ ). Applying **make-cg** algorithm based on this counterfactual quantity produces the counterfactual graph in Figure 8(d), which combines two worlds: the actual world where the teacher has normally  $A = a_0$  and the counterfactual world where *the same* teacher is assigned  $A = a_1$ . Both variables  $C$  are reduced to a single variable, and  $Y$  and  $Y_{a_1}$  are connected by an unobservable confounder. The counterfactual graph is composed of three c-components  $\{C\}, \{A\}, \{Y, Y_{a_1}\}$ . Applying algorithm IDC\* [53] results in:

$$\mathbb{P}(y_{a_1}|a_0) = \frac{\sum_{y,c} Q(c) Q(a_0) Q(y, y_{a_1})}{\mathbb{P}(a_0)} \quad (42)$$

where  $Q(\mathbf{v}) = \mathbb{P}(\mathbf{v}|pa(\mathbf{V}))$  in the counterfactual graph. Hence,

$$\begin{aligned} \mathbb{P}(y_{a_1}|a_0) &= \frac{\sum_{y,c} \mathbb{P}(c) \mathbb{P}(a_0|c) \mathbb{P}(y, y_{a_1}|c)}{\mathbb{P}(a_0)} \\ &= \frac{\sum_c \mathbb{P}(c) \mathbb{P}(a_0|c) \mathbb{P}(y_{a_1}|c)}{\mathbb{P}(a_0)} \end{aligned} \quad (43)$$

$$\begin{aligned} &= \frac{\sum_c \mathbb{P}(c) \mathbb{P}(a_0|c) \mathbb{P}(y|a_1, c)}{\mathbb{P}(a_0)} \quad (44) \\ &= \frac{0.5 \times 0.8 \times 0.25 + 0.5 \times 0.2 \times 0.01}{0.5} \\ &= 0.202 \end{aligned}$$

$y$  in Eq. 43 is cancelled by summation while  $\mathbb{P}(y_{a_1}|c)$  in the same equation is transformed into  $\mathbb{P}(y|a_1, c)$  in Eq. 44 using Rule 2 of the *do*-calculus.

### 6.1.3. Identifiability of direct and indirect effects

In Markovian models, the average natural direct effect NDE and the average natural indirect effect NIE are always identifiable (from observational data) and can be computed as follows [39]:

$$\text{NDE}_{a_1, a_0}(Y) = \sum_{\mathbf{s}} \sum_{\mathbf{z}} \left( \mathbb{E}[Y|a_1, \mathbf{z}] - \mathbb{E}[Y|a_0, \mathbf{z}] \right) \mathbb{P}(\mathbf{z}|a_0, \mathbf{s}) \mathbb{P}(\mathbf{s}) \quad (45)$$

$$\text{NIE}_{a_1, a_0}(Y) = \sum_{\mathbf{s}} \sum_{\mathbf{z}} \mathbb{E}[Y|a_0, \mathbf{z}] \left( \mathbb{P}(\mathbf{z}|a_1, \mathbf{s}) - \mathbb{P}(\mathbf{z}|a_0, \mathbf{s}) \right) \mathbb{P}(\mathbf{s}) \quad (46)$$

where  $\mathbf{Z}$  is a set of mediator variables and  $\mathbf{S}$  is any set of variables satisfying the back-door criterion between the sensitive variable  $A$  and the mediator variables  $\mathbf{Z}$ , that is, (i) no variable in  $\mathbf{S}$  is a descendant of  $A$  and (ii)  $\mathbf{S}$  blocks all back-door paths between  $A$  and  $\mathbf{Z}$ . A simpler formulation can be used in case there is no confounding between  $A$  and  $\mathbf{Z}$ , where the need for  $\mathbf{S}$  is dropped altogether:

$$\text{NDE}_{a_1, a_0}(Y) = \sum_{\mathbf{z}} \left( \mathbb{E}[Y|a_1, \mathbf{z}] - \mathbb{E}[Y|a_0, \mathbf{z}] \right) \mathbb{P}(\mathbf{z}|a_0) \quad (47)$$

$$\text{NIE}_{a_1, a_0}(Y) = \sum_{\mathbf{z}} \mathbb{E}[Y|a_0, \mathbf{z}] \left( \mathbb{P}(\mathbf{z}|a_1) - \mathbb{P}(\mathbf{z}|a_0) \right) \quad (48)$$

In semi-Markovian models, NDE and NIE are not generally identifiable, even if we have the luxury to perform any experiment using *RCT*, because of the nested

counterfactuals  $\mathbb{P}(Y_{a_1}, \mathbf{Z}_{a_0})$  and  $\mathbb{P}(Y_{a_0}, \mathbf{Z}_{a_1})$  in Eq. 14 and Eq. 15, respectively. Nevertheless, these quantities are identifiable *from experimental data* provided that there is a set of variables  $\mathbf{W}$  which are parents of the outcome variable  $Y$  but non-descendants of  $A$  and  $Z$  such that  $Y_{a_0,z} \perp\!\!\!\perp Z_{a_0} | \mathbf{W}$  (reads:  $Y_{a_0,z}$  and  $Z_{a_0}$  are independent conditional on  $\mathbf{W}$ ). This condition can be easily checked from the causal graph as follows:  $\mathbf{W}$  d-separates  $Y$  and  $\mathbf{Z}$  in the graph formed by deleting all arrows emanating from  $A$  and  $Z$ , denoted simply as  $(Y \perp\!\!\!\perp \mathbf{Z} | \mathbf{W})_{G_{AZ}}$ .

If such graphical condition is satisfied, NDE and NIE can be computed from experimental quantities as follows:

$$\text{NDE}_{a_1, a_0}(Y) = \sum_{\mathbf{z}, \mathbf{w}} \left( \mathbb{E}[Y_{a_1, \mathbf{z}} | \mathbf{w}] - \mathbb{E}[Y_{a_0, \mathbf{z}} | \mathbf{w}] \right) \mathbb{P}(\mathbf{Z}_{a_0} = \mathbf{z} | \mathbf{w}) \mathbb{P}(\mathbf{w}) \quad (49)$$

$$\text{NIE}_{a_1, a_0}(Y) = \sum_{\mathbf{z}, \mathbf{w}} \mathbb{E}[Y_{a_0, \mathbf{z}} | \mathbf{w}] \left( \mathbb{P}(\mathbf{Z}_{a_1} = \mathbf{z} | \mathbf{w}) - \mathbb{P}(\mathbf{Z}_{a_0} = \mathbf{z} | \mathbf{w}) \right) \mathbb{P}(\mathbf{w}) \quad (50)$$

#### 6.1.4. Identifiability of path-specific effects

The identifiability of  $\text{PSE}_\pi(a_1, a_0)$  in Markovian models depends on whether  $\mathbb{P}(y | do(a_1 |_\pi, a_0 |_{\bar{\pi}}))$  is identifiable. Avin et al. [57] gave a single necessary and sufficient criterion for the identifiability of  $\mathbb{P}(y | do(a_1 |_\pi, a_0 |_{\bar{\pi}}))$  in Markovian models called recanting witness criterion. This criterion holds when there is a vertex  $W$  along the causal path  $\pi$  connected to  $Y$  through another causal path not in  $\pi$ . For instance, Figure 8(c) satisfies the recanting witness criterion when  $\pi = A \rightarrow W \rightarrow Z \rightarrow Y$  with  $W$  as witness. The corresponding graph structure is called “kite” graph. When this criterion is satisfied,  $\mathbb{P}(y | do(a_1 |_\pi, a_0 |_{\bar{\pi}}))$  is not identifiable, and consequently,  $\text{PSE}_\pi(a_1, a_0)$  is not identifiable. Shpitser [55] generalizes this criterion to semi-Markovian models known as recanting district criterion.

#### 6.2. Estimation based on full knowledge of the causal model parameters

The main reason behind the unidentifiability of causal quantities (causal effect, counterfactuals, etc.) is the presence of unobservable variables, namely, hidden latent variables. Some causal-based fairness notions, such as counterfactual fairness [8], can be assessed in the presence of such unobservable latent variables. However, The only requirement is knowledge of the causal model structure (skeleton). Based on the causal model, the latent/background variables are estimated using observable data. Then, the predictor is trained using both observable (non-descendants of the sensitive attributes) as well as the estimated latent variables. Such predictors tend to be fairer than typical predictors (trained using only observable variables) since they consider hidden bias captured by latent variables. Given the full causal model, counterfactual

fairness can be assessed by generating, for every observable data sample, a counterfactual data sample by simply changing the sensitive attribute value (e.g., turn male into female) and then using the three-step process (abduction, action, prediction) to compute the outcome. The predictor is considered fair if the predicted outcome distributions of both (protected and unprotected) groups are similar.

In this survey, and for clarity of presentation, the counterfactually fair learning and estimation approach described by Kusner et al.[8] is illustrated in Algorithm 1. The approach takes as input the labeled dataset  $\mathcal{D}$ . It assumes the structure of

---

**Algorithm 1:** Counterfactual learning and assessment

---

**input** : Labelled dataset  $\mathcal{D} \equiv \{(A^{(i)}, X^{(i)}, Y^{(i)})\} i = 1 \dots n$ ,  
Causal model  $\mathcal{M}$  with causal graph  $\mathcal{G}$

**output:** Predictor  $\hat{Y}$  w/ counterfactual fairness constraints,  
Estimation of counterfactual bias

- 1 Fit the causal model  $\mathcal{M}$  based on labelled dataset  $\mathcal{D}$ ;
  - 2 Estimate the posterior distribution of the latent variable(s)  $U$ :  $\mathbb{P}_{\mathcal{M}}(U|X, A)$ ;
  - 3 For every sample  $(x^{(i)}, a^{(i)})$  of  $\mathcal{D}$ , generate a counterfactual sample for every possible combination of sensitive attribute values  $a'^{(i)}$ ;
  - 4 For every counterfactual sample generated in Step 3, use the three inference steps: **Abduction**, **Action**, and **Prediction** to compute the values of the remaining variables  $x'^{(i)}$  as well as the label  $y'^{(i)}$  ;
  - 5 Train a predictor  $\hat{Y}$  using only variables non-descendants of  $A$  and the estimated latent variables  $U$ , that is,  $\hat{Y} \equiv h_{\theta}(U, X_A)$  where  $\theta$  represents the predictor parameters and  $X_A \subseteq X$  are non-descendants of  $A$ ;
  - 6 Use the trained predictor  $\hat{Y}$  to compute  $\hat{y}^{(i)}$  for every observed as well as generated counterfactual sample in the test set;
  - 7 Using predicted outcomes  $\hat{y}^{(i)}$  of observed and counterfactual samples, estimate the counterfactual bias;
- 

the causal model  $\mathcal{M}$  is available, including the relationships between all variables (observed and latent) along with the causal graph. For example, Kusner et al.[8] assumed a single latent variable  $U$  and a combination of Normal and Poisson distributions for their illustrated example. If the parameter values are unknown, they can be estimated using the observed data (Step 1). The approach has two objectives (1) learning a predictor  $\hat{Y}$  while taking into consideration counterfactual fairness constraints and (2) assessing whether the prediction is counterfactually fair. Given the causal model  $\mathcal{M}$ , Step 1 fits the model parameters to the observed data  $\mathcal{D}$ . This

allows to estimate the posterior distribution of the latent variable(s)  $U$  (Step 2). Step 3 is straightforward and consists in generating a counterfactual sample for every observed sample by assigning a different value to the sensitive attribute. For example, if the observed sample is  $(X = x_1, A = \textit{male})$ , the counterfactual sample would be  $(X = x_1, A = \textit{female})$ . Step 4 is the hard part of generating the counterfactual sample, which is, computing the “would-be” values of the remaining variables of these “made-up” samples. It tries to answer the core question of causal approaches: what would the output be, had the sensitive attribute value was different? This is achieved by the three steps process (Theorem 7.1.7 in [7]) to compute counterfactuals:

1. **Abduction:** compute the posterior distribution of  $U$  given the evidence, that is, the observed data  $(X = x^{(i)}, A = a^{(i)})$
2. **Action:** Substitute the observed sensitive attribute equation with the counterfactual value (e.g.,  $A = \textit{female}$ ) in the causal model  $\mathcal{M}$
3. **Prediction:** Compute the remaining variables values (including  $Y$ ) using the new equations.

Enforcing counterfactual fairness while training the predictor (Step 5) is achieved through the use of two principles. First, by using only variables non-descendants of the sensitive attribute  $A$  which is direct implication of counterfactual fairness definition (Lemma 1 in [8]). Second, unlike typical predictors, by considering latent variables  $U$  in the training. This allows to involve the background sources of the social bias encoded in the latent variables. The predictor is then used to generate predictions for observed and counterfactual samples in the test set<sup>12</sup>(Step 6). Finally, counterfactual bias can be estimated by comparing the output of observed samples and counterfactual samples (Step 7). In Kusner et al. [8] and in this survey, density distributions are used to estimate the counterfactual bias.

Algorithm 1 illustrates the steps for the (counterfactually) fair learning of a predictor and the estimation of counterfactual bias. However, typical real-world scenarios come with an already trained predictor and hence need only the second goal, that is, counterfactual fairness assessment of that predictor. Algorithm 2 is a modified version of Algorithm 1, which, given an already trained predictor, tries to tell how counterfactually fair the predictor is.

It is important to mention that assessing a given predictor’s counterfactual bias requires access to that predictor. This is needed to generate predictions for the counterfactual samples (Step 5).

---

<sup>12</sup>80% of the data is used for training while 20% of the data is used for testing.



---

**Algorithm 2:** Counterfactual fairness assessment

---

**input** : Labelled dataset  $\mathcal{D} \equiv \{(A^{(i)}, X^{(i)}, Y^{(i)})\} i = 1 \dots n$ ,

Predictor  $\hat{Y}$ ,

Causal model  $\mathcal{M}$  with causal graph  $\mathcal{G}$

**output:** Estimation of counterfactual bias in  $\hat{Y}$

- 1 Fit the causal model  $\mathcal{M}$  based on labelled dataset  $\mathcal{D}$ ;
  - 2 Estimate the posterior distribution of the latent variable(s)  $U$ :  $\mathbb{P}_{\mathcal{M}}(U|X, A)$ ;
  - 3 For every sample  $(x^{(i)}, a^{(i)})$  of  $\mathcal{D}$ , generate a counterfactual sample for every possible combination of sensitive attribute values  $a'^{(i)}$ ;
  - 4 For every counterfactual sample generated in Step 3, use the three inference steps: **Abduction**, **Action**, and **Prediction** to compute the values of the remaining variables  $x'^{(i)}$  as well as the label  $y'^{(i)}$  labels ;
  - 5 Use predictor  $\hat{Y}$  to compute  $\hat{y}'^{(i)}$  for every generated counterfactual sample in the test set;
  - 6 Using observed and counterfactual samples, estimate the counterfactual bias;
- 

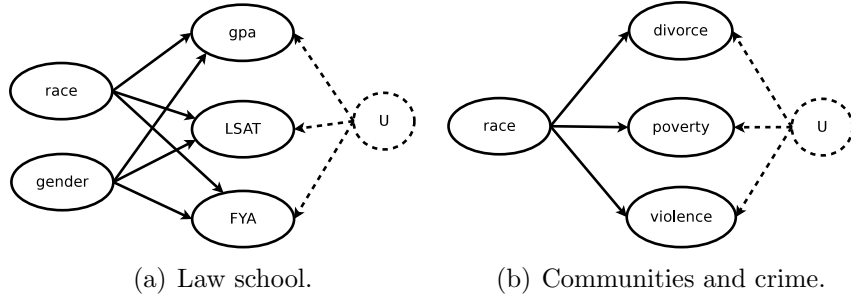


Figure 9: Causal graphs of the law school, and the communities and crime datasets.

Similarly to Kusner et al. [8], Stan programming language [64] is used for counterfactual learning and assessment. However, in addition to the *law school* [65] used in Kusner et al., a second dataset is used in the empirical analysis, namely, *communities and crime* [66]. The *law school* dataset [65] includes information on 21,790 law students, such as their LSAT, their GPA collected prior to law school, and their first-year average grade (FYA). Given this data, a school wishes to predict whether an applicant will have a high FYA. Race is considered a sensitive attribute. *Communities and crime* dataset [66] contains information relevant to per capita violent crime rates in different communities in the United States (e.g., percentage of people under the poverty level, median family income, percentage of the population who are

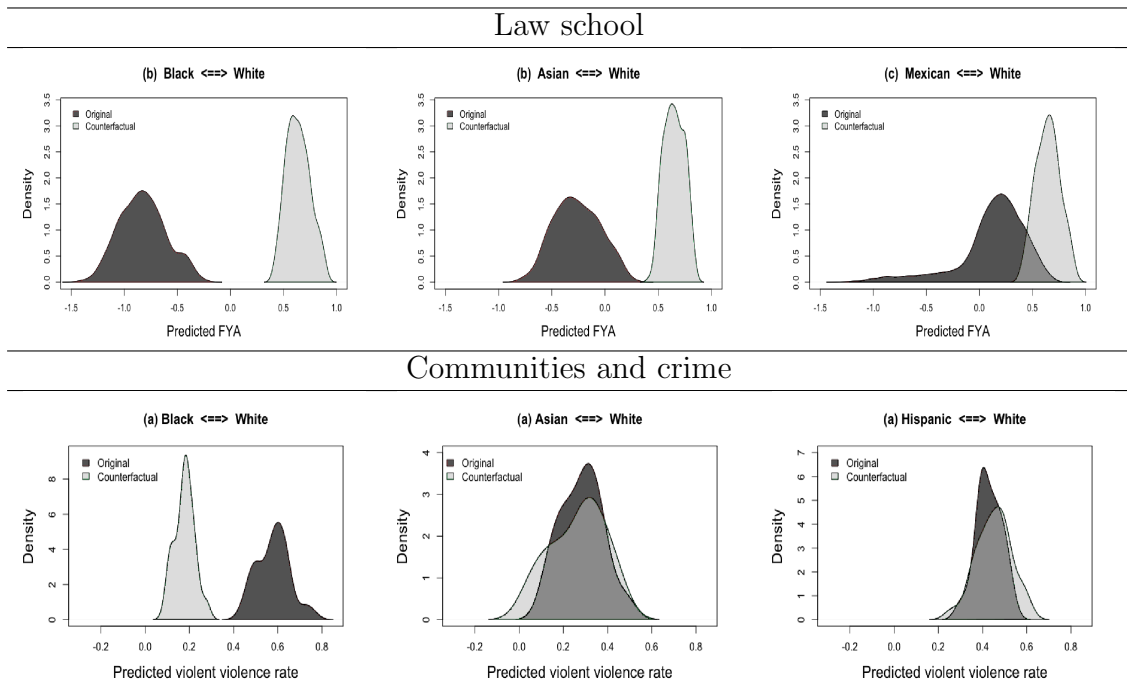


Table 6: Counterfactual learning and assessment of the law school, and communities and crime datasets.

divorced, etc.,). The goal is to predict the crime rate. Race is considered a sensitive attribute.

For both datasets, the experiment consists of training a baseline model using logistic regression<sup>13</sup>, then applying Algorithm 2 to assess counterfactual fairness to both the original and counterfactual sampled data and plot how the distribution of predicted FYA changes for that baseline model. Since sensitive attributes are not binary, the experiment is repeated for every counterfactual change (e.g., white vs. black, white vs. Asian, etc.). If both distributions are superposed, it means the prediction is counterfactually fair. Far-apart distributions indicate counterfactual bias.

Table 6 shows the density plots of  $\hat{Y}$  for both datasets. We assume that the true model of the world is given by the causal graphs presented in Figure 9. The causal graphs use a single latent/background variable  $U$ . For the *law school*,  $U$  can repre-

---

<sup>13</sup>The predictor uses all the variables in the dataset, including the sensitive variables, to make the predictions.

sent the knowledge of the student, while for the *communities and crime*, it refers to the criminality of the individual, which can capture other aspects of the individual that might have been used by the police. The red plots refer to the filtered original samples, while the green plots refer to the corresponding counterfactual samples. For example, in plot (a), the red plot represents the predicted first-year average ( $F\hat{Y}A$ ) for observed black students. In contrast, the green plot represents  $F\hat{Y}A$  for the corresponding counterfactual samples. Density plots for the law school corroborate the findings of Kusner et al. [8], that is, racial differences exhibit counterfactual unfairness in favor of whites. The unfairness is peaked when predicting the FYA for white versus black students. The race-based discrimination is present in the *communities and crime* result similarly: the counterfactual bias is peaked for white versus black communities. However, racial comparison can be considered counterfactually fair when comparing Hispanic/white and Asian/white communities.

### 6.3. Potential outcome estimation techniques

Causal inference in the potential outcome framework focuses on estimating the causal effect of a treatment variable  $A$  (e.g., the sensitive attribute) on an effect variable  $Y$  (e.g., the decision outcome). As mentioned in Section 4.2, three assumptions are typically made for causal effect estimation: SUTVA, ignorability, and positivity. In line with the potential outcome framework literature, this survey focuses on causal inference approaches that rely on the three assumptions [35, 13], namely, re-weighting [11], matching [34], and stratification [11].

#### 6.3.1. Re-weighting

One of the main challenges in causal inference is that the sensitive attribute is not randomly assigned in observational data. That is, the distribution in the observed dataset does not reflect the true distribution. Sample re-weighting methods try to overcome this discrepancy by assigning appropriate weights to sample units in the observational data. The aim is to generate a pseudo-population on which the distributions of the protected (e.g., female) and unprotected (e.g., male) groups are the same as in the original total population. This is achieved by defining a balancing score  $b(x)$  satisfying  $A \perp x \mid b(x)$ . The most common approach to balancing score is based on the propensity score [67], which is defined as the conditional probability of the sensitive attribute given background variables:

$$e(x) = \mathbb{P}(A = 1 \mid X = x) \tag{51}$$

Table 7: Estimation of ATE using inverse propensity weighting (IPW) on the job hiring example with propensity score  $e(c_i)$  and balancing score  $b(c_i)$ .

Female applicants (Treatment group)						Male applicants (Control Group)					
$i$	$A$	$C$	$Y$	$e(c_i)$	$b(c_i)$	$i$	$A$	$C$	$Y$	$e(c_i)$	$b(c_i)$
1:	1	0	1	2/3	3/2	13:	0	0	1	2/3	3
2:	1	0	1	2/3	3/2	14:	0	0	0	2/3	3
3:	1	0	0	2/3	3/2	15:	0	0	0	2/3	3
4:	1	0	0	2/3	3/2	16:	0	0	0	2/3	3
5:	1	0	0	2/3	3/2	17:	0	1	1	1/3	3/2
6:	1	0	0	2/3	3/2	18:	0	1	1	1/3	3/2
7:	1	0	0	2/3	3/2	19:	0	1	1	1/3	3/2
8:	1	0	0	2/3	3/2	20:	0	1	1	1/3	3/2
9:	1	1	1	1/3	3	21:	0	1	0	1/3	3/2
10:	1	1	1	1/3	3	22:	0	1	0	1/3	3/2
11:	1	1	1	1/3	3	23:	0	1	0	1/3	3/2
12:	1	1	0	1/3	3	24:	0	1	0	1/3	3/2

Propensity scores  $e(x)$  can be used to equate groups based on covariates  $X$ . In inverse propensity weighting (IPW), the balancing score  $b(x)$  for each sample is defined as:

$$b(x) = \frac{A}{e(x)} + \frac{1 - A}{1 - e(x)} \quad (52)$$

where  $A = 1$  corresponds to the protected group and  $A = 0$  corresponds to the unprotected group. The IPW estimator of ATE (Eq. 4) is defined as:

$$\hat{ATE}_{IPW} = \frac{1}{n} \sum_{i=1}^n \frac{A_i Y_i}{\hat{e}(x_i)} - \frac{1}{n} \sum_{i=1}^n \frac{(1 - A_i) Y_i}{1 - \hat{e}(x_i)} \quad (53)$$

Notice that the estimation of ATE is based only on the observable outcome (no counterfactual outcomes) and on the estimation of  $e(x_i)$ , that is,  $\hat{e}(x_i)$ .

Table 7 shows the values of propensity ( $e(c_i)$ ) as well as balance ( $b(c_i)$ ) scores for each unit  $i$  in the simple job hiring example. Using Eq. 53, the  $\hat{ATE}_{IPW}$  estimation of ATE is 0.25 indicating discrimination in favor of the female group.

When the propensity score is estimated, the normalized version of  $\hat{ATE}_{IPW}$  is preferred:

$$\hat{ATE}_{IPW}^{norm} = \left[ \frac{\sum_{i=1}^n \frac{A_i Y_i}{\hat{e}(x_i)}}{\sum_{i=1}^n \frac{A_i}{\hat{e}(x_i)}} - \frac{\sum_{i=1}^n \frac{(1 - A_i) Y_i}{1 - \hat{e}(x_i)}}{\sum_{i=1}^n \frac{(1 - A_i)}{1 - \hat{e}(x_i)}} \right] \quad (54)$$

$\widehat{\text{ATE}}_{IPW}^{norm}$  for the same example equals 0.125, which is a perfect estimation of ATE in this case as both values coincide ( $\text{ATE} = 0.125$ ).

The correctness of the IPW estimation relies heavily on the quality of the propensity score estimation ( $\hat{e}(X)$ ). A slight misspecification of propensity scores may lead to a significant discrepancy in the ATE estimation. In such cases, doubly robust (DR) estimation is recommended [68]. DR combines IPW estimation with outcome regression so that the estimation remains valid even if one of the approaches is incorrect (but not both). Another limitation of IPW can be observed if the propensity score  $e(X) = \mathbb{P}(A | X)$  for some value of  $X$  is small. In such a case, the estimation may suffer instability. To address this issue, trimming [69] is typically used. Trimming consists of removing the samples with a propensity score below a certain threshold.

### 6.3.2. Matching

Matching techniques [34, 70] focus on estimating the counterfactual outcome of units. The idea is to estimate the counterfactual outcomes  $Y_i^1|A = 0$  and  $Y_i^0|A = 1$  based on the matched neighbors of unit  $i$  in the opposite group. For example, given an observed female candidate  $f_k$ , estimating the counterfactual outcome (hiring decision) had she been a male is based on the units in the male group that are the most comparable to  $f_k$ . Hence, the first and main issue is to define a similarity metric between two given units (e.g.,  $x_i$  and  $x_j$ ). The most common approach is to rely on the propensity scores of units:

$$D(i, j) = |e(x_i) - e(x_j)| \tag{55}$$

and its logit version:

$$D(i, j) = |\text{logit}(e(x_i)) - \text{logit}(e(x_j))| \tag{56}$$

which is preferred as it has been proven to reduce the bias [71].

The second issue is the matching algorithm: how many neighbors to consider and how these neighbors are weighted to obtain the estimation. Matching algorithms include [34]:

- Exact matching: uses only identical matches. Typically, it is infeasible since many units will remain unmatched.
- Nearest neighbor matching (NNM): constructs the counterfactual using the closest neighbors according to a similarity metric. It can run with or without replacement, that is, by returning a matched unit to the pool or not. The no-replacement variant makes the estimation dependent on the units' matched order.

Table 8: Characteristics of the real-world datasets.

Dataset	Ref.	Sample size	Sensitive feature(s)	Covariate(s)	Mediator(s)	Outcome
Communities and crimes	[66]	1994	race	age unemp. rate	poverty rate divorce rate	crime rate
Compas	[72]	5915	race gender	age	priors	recidivism
German credit	[19]	1000	gender marital status	age	emp. length	default
Berkeley	[73]	4526	gender	department	department	admission

- Caliper matching: a version of NNM that restricts matches to those within a caliper, and in case there are no matches within the caliper, some units may not receive matches because no neighbors fall within their caliper. A hybrid variant consists of using a caliper, and in case of no neighbors, select a matching neighbor outside the caliper.
- Radius matching: the same as caliper, but matching is done with replacement.
- Kernel matching: an extension to all the above algorithms where to match a unit  $i$ , all units in the opposite group are used. Each unit is weighted according to the distance to the unit  $i$ .

### 6.3.3. Stratification

Stratification [11] uses the same principle underlying the identifiability approach (Section 6.1), that is, adjusting on confounders. The aim is to split the entire observed data into consistent groups such that the units in the same group can be considered as sampled from data under RCT. The two ingredients of stratification are the splitting of groups and then the combination of the created groups. The stratification estimator of ATE can be defined generically as:

$$\hat{ATE}^{strat} = \sum_{k=1}^K m(k) [\bar{Y}_1(k) - \bar{Y}_0(k)] \quad (57)$$

where  $K$  is the number of stratification groups,  $m(k)$  is the portion of units in group  $k$  to the total number of units  $N$ , (Eq. 11) for groups  $A = 1$  and  $A = 0$ , respectively.  $\hat{ATE}^{strat}$  expression has the same structure as the back-door formula (Eq. 35).

If all variables needed for the stratification are observed and the available data is infinitely large,  $ATE^{strat}$  can lead to a consistent and unbiased estimator of ATE. However, in typical datasets, stratification may result in strata with few or no units.

Consequently, some CATE estimates cannot be calculated using the available data. Propensity score can be used to address this data sparseness problem. The main idea is the following: “strata with identical propensity scores can be combined into more coarse strata” [34]. In other words, propensity score can be considered a single stratifying variable that usually results in larger strata. The same idea is used in the SCM framework to address the sparseness of data when computing identifiable expressions.

Other estimation methods in the potential outcome framework include tree-based methods [74], representation learning methods [75], and meta-learning methods [76].

#### 6.4. Estimating causal effects on benchmark datasets

We conduct experiments on four real-world datasets which are commonly used in discrimination-discovery literature, namely: *communities and crime* [66], *Compas* [72], *German credit* [19] and *Berkely admission* [73].

*Communities and crime* dataset [66] contains information relevant to per capita violent crime rates in 1994 different communities in the United States, and the goal is to predict that crime rate. Race is kept as a sensitive attribute. However, in these experiments, we transformed the label into a binary feature by thresholding<sup>14</sup> where 1 corresponds to a high violent crime rate and 0 corresponds to a low violent crime rate. *Compas* dataset [72] includes information of 5915 individuals from Broward County, Florida, initially compiled by ProPublica [72] and the goal is to predict the two-year violent recidivism. That is, whether a convicted individual would commit a violent crime in the following two years (1) or not (0). We consider race as a sensitive feature. *German credit* dataset comprises data of 1000 individuals applying for loans. The goal is to predict whether an individual will default on the loan (1) or not (0). The gender and the personal status of an individual are considered sensitive features. Finally, *Berkeley admission* [73] dataset consists of 4526 applicants to UC-Berkeley’s graduate programs in Fall 1973. The goal is to predict whether a student is admitted (1) or not (0). Gender is treated as a sensitive feature. *Berkeley admission* is commonly used to illustrate the Simpson’s Paradox [3].

Table 8 summarizes the main characteristics of the four benchmark datasets. In the potential outcome framework, estimation techniques rely on selecting a set of covariates and mediators. Column four lists the covariate(s), while column five lists the mediator variables for each dataset. A covariate can be a confounder or any other variable that might distort the causal effect estimation. A mediator is a covariate that

---

<sup>14</sup>The median value of the violent crime rate in the dataset is used as a threshold.

is simultaneously influenced by the sensitive feature and influences the outcome<sup>15</sup>. In addition to total variation TV (statistical parity) serving as a baseline measure and also to illustrate the need for causality, we use estimation techniques for seven causal-based fairness notions defined in the potential outcome framework, namely, *ATE\_IPW*, *ATE\_match*, *ATE\_DR*, *ATE\_strat*, *TE\_imp*, *DE\_imp*, and *IE\_imp*. The four first notions are estimations of ATE using four different approaches. *ATE\_IPW* uses inverse propensity weighting as explained in Section 6.3.1. The variables used to estimate the propensity score are the covariates listed in column 4 of Table 8. *ATE\_match* is the matching estimation of ATE (Section 6.3.2). Recall that this approach matches each unit (individual) with one or more “similar” units in the other group. Propensity score matching is used. *ATE\_DR* is a double robust estimation of ATE, which requires at least propensity score estimation or outcome regression to be correct, as explained in Section 6.3.1. *ATE\_strat* is a stratification-based estimation of ATE (Section 6.3.3). For all these ATE estimations, we used the *causallib* package implementations [77]. The three remaining measures, namely, *TE\_imp*, *DE\_imp*, and *IE\_imp*, are imputation-based estimations of total, direct, and indirect effects. The main idea behind imputation approach is to provide  $K + 1$  outcome models that characterize  $\mathbb{E}[Y|\mathbf{X}, A]$ ,  $\mathbb{E}[Y|\mathbf{X}, A, M_1]$ ,  $\dots$ ,  $\mathbb{E}[Y|\mathbf{X}, A, M_k]$ , where  $\mathbf{X}$  is the set of the covariates (column 4 in Table 8) and  $M_1, \dots, M_k$  are the mediators (column 5 in Table 8). The imputation approach involves modeling the conditional means of the outcome variable  $Y$ , given the sensitive feature  $A$ , pretreatment confounders  $X$ , and varying sets of mediators  $M$ . The *paths* package implementation [78] is used to compute these estimations.

Figure 10 shows the causal effects (of the sensitive attribute  $A$  on the outcome  $Y$  values according to the seven estimators on the four benchmark datasets. The values are between  $-1$  (full discrimination against one group) and  $1$  (full discrimination against the second group). A value of  $0$  indicates fairness of the outcome with respect to both groups.

For the three first datasets (*Communities and crimes*, *Compas*, and *German credit*), TV overestimates the discrimination as it results in higher values compared to all causal-based notions (except *ATE\_match* in *German credit* where the results are reverted. In other words, the discrimination is now against male applicants:  $-0.167$ ). For *Berekley*, TV concludes significant discrimination against women ( $-0.14$ ) while causal-based notions show either the absence of discrimination ( $0$  for *TE\_match* and *DE\_imp*) or discrimination in favor of women which confirms the Simpson’s

---

<sup>15</sup>Graphically, a mediator is a variable in the causal path between the sensitive feature and the outcome, i.e., an explaining or redlining (proxy) variable.



paradox. For the three last estimations ( $TE_{imp}$ ,  $DE_{imp}$ , and  $IE_{imp}$ ), notice that TE is approximately the total of DE and IE which is expected as the total effect is composed of direct and indirect effects.

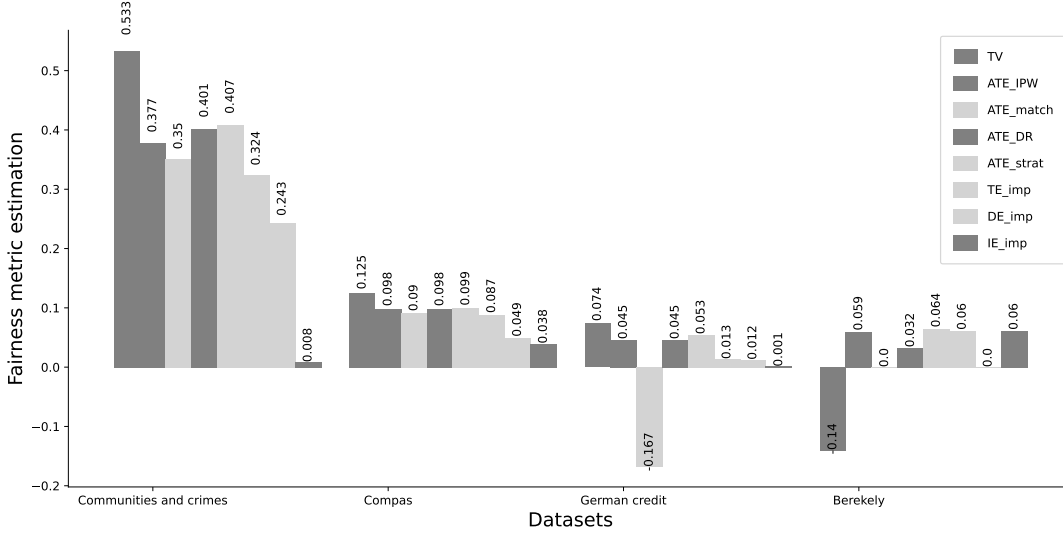


Figure 10: Estimation of causal effects on real-world datasets.

As mentioned earlier, estimation techniques aim to balance the control and the treatment groups in order to remove (or mitigate) potential selection bias. For the particular case of re-weighting based estimation of ATE ( $ATE_{IPW}$ ), one way to assess whether the balancing worked is to use the absolute standard mean difference (**abs-smd**) [79] between the treatment and the control groups for each considered covariate (Column 4 in Table 8). That is, measuring the difference in means between groups divided by the (pooled) standard deviation. Figure 11 shows the **abs-smd** for each covariate in the four datasets prior (unweighted) and posterior (weighted) to applying IPW re-weighting. One can observe the same pattern in all benchmark datasets. The unweighted **abs-smd** across control and treatment groups is at least one order of magnitude higher than the weighted **abs-smd**. For instance, the unweighted **abs-smd** across control and treatment groups reached 0.68 standard deviations for the *Unemployment rate* covariate in the *communities and crimes* dataset. After IPW re-weighting, the **abs-smd** is significantly reduced (0.078).

## 7. Suitability and applicability

Section 5 lists 19 causal-based fairness notions. Given a real-world scenario, selecting which fairness notion to use is a challenging and error-prone task, as using

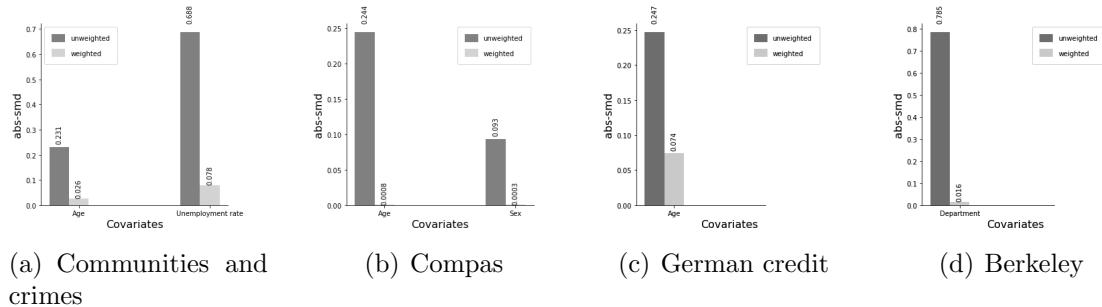


Figure 11: Absolute standard mean difference of covariate values of different groups (e.g., male vs female) prior and posterior to IPW re-weighting.

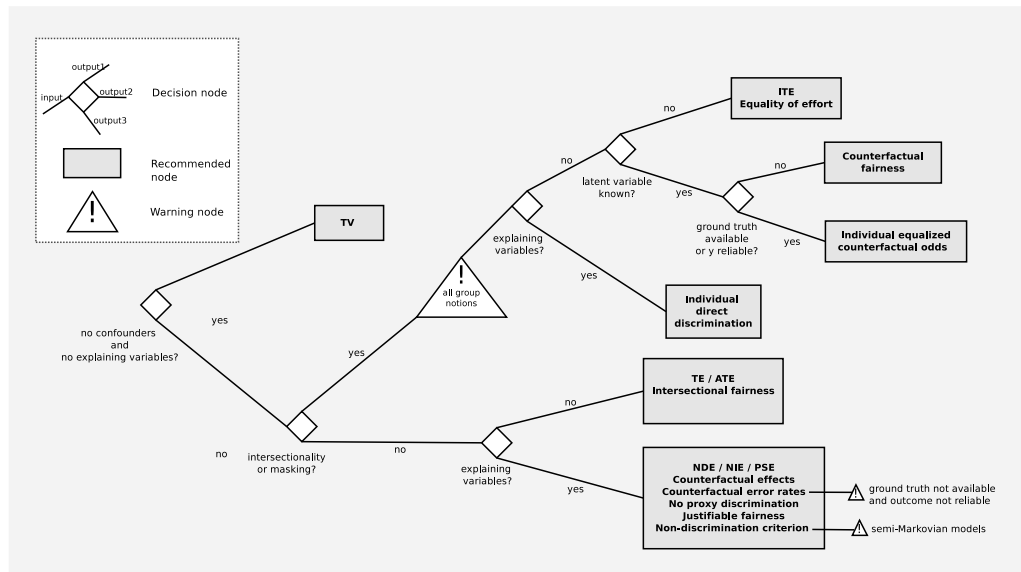


Figure 12: Guideline for causal-based fairness notions selection

the wrong fairness notion may indicate unfairness in an otherwise fair scenario or the opposite (failing to detect unfairness in an unfair scenario). On the other hand, according to Pearl’s SCM framework, computing causal quantities (interventions and counterfactuals) depends on their identifiability. Hence, even if a fairness notion is appropriate in some setups, it might not be applicable because of identifiability issues. The two following subsections address the suitability and the applicability of causal-based fairness notions.

### 7.1. Suitability

In this section, we try to systemize the selection process by considering the subtleties of each causal-based fairness notion and defining 6 criteria that correspond to characteristics of the real-world scenario at hand. We check whether each criterion holds in the scenario at hand or not. Then, these answers will be used to recommend the most suitable causal-based fairness notion. The criteria are listed and briefly described as follows.

- **Presence of confounding:** A variable that is a common cause of two or more other variables.
- **Presence of explaining variable:** A variable that is correlated with the sensitive attribute such that any discrimination that is explained using that variable is considered legitimate and acceptable.
- **Likelihood of intersectionality:** A specific type of bias due to the combination of sensitive attributes. An individual might not be discriminated against based on race or based on gender only, but she might be discriminated against because of a combination of both.
- **Likelihood of masking:** A form of intentional discrimination that allows decision-makers with prejudicial views to discriminate against individuals or groups while masking their intentions.
- **Latent variables are known:** Latent (background) variables are not observable. However, they are identified in some scenarios, and their relationships with observable variables are known.
- **Ground truth or reliable outcome:** the label in the training data can or cannot be reliable. In several scenarios, the outcome is inferred by humans (job hiring, college admission, etc.) and hence can encode bias. The most reliable outcome is when the ground truth is available<sup>16</sup>.

---

<sup>16</sup>An example of a scenario where ground truth is available is when predicting whether an individual has a disease. For example, the ground truth value is observed by submitting the individual to a blood test (Assuming the blood test is flawless). An example of a scenario where ground truth is unavailable is predicting whether a job applicant is hired. The outcome in the training data is inferred by a human decision-maker, which is often a subjective decision, no matter how hard she tries to be objective.

The diagram in Figure 12 can be used as a guideline to select an appropriate causal-based fairness notion given a real-world scenario.

Confounding variables result in backdoor paths between the sensitive attribute ( $A$ ) and the outcome ( $Y$ ). For example, path  $A \leftarrow C \rightarrow Y$  in Figure 6 is a backdoor path. Backdoor paths are not causal paths, but they contribute to the association between the  $A$  and  $Y$ . Therefore, they are the reason why it is said that “correlation is different than causation”. In the absence of confounding, the total causal effect (TE and ATE) coincides with the difference in conditional probabilities  $TV = \mathbb{P}(y|a_1) - \mathbb{P}(y|a_0)$ , which corresponds to statistical parity. On the other hand, if there are no explaining variables in the model representation of the world, both direct and indirect causal paths are discriminatory<sup>17</sup>. Consequently, assessing unfairness/bias due to the sensitive attribute does not require considering separately the different causal paths (direct, indirect, and path-specific). In such cases (absence of confounding and explaining variables), causal inference is not needed to assess fairness appropriately.

Any unintentional type of bias can also be “orchestrated” intentionally by decision-makers with prejudicial views. To appropriately assess the bias in the presence of such masking attempts, it is recommended to avoid group-based notions as they can more easily be gamed by prejudicial decision-makers. Intersectionality is similar to masking as both lead to discrimination, which is difficult to detect at the group level and hence requires more fine-grained measures. Therefore, individual causal-based fairness notions are recommended in the presence of one of those criteria. For individual notions, in the presence of explaining variables, it is recommended to use individual direct discrimination (Section 5.6) as it is the only individual notion listed in Section 5 that distinguishes direct from indirect discrimination. Counterfactual fairness (Section 5.3) and individual equalized counterfactual odds (Section 5.10) are recommended to be used in case the latent variables are known. If the ground truth is not available or the outcome  $Y$  is not reliable, individual equalized counterfactual odds is not recommended.

For the group causal-based fairness notions, if there are no explaining variables, there is no need to consider the different causal paths and hence TE, ATE, or interventional fairness (Section 5.9) can be safely used. In the presence of explaining variables, the remaining causal-based fairness notions are appropriate to use with two exceptions. First, the non-discrimination criterion is misleading if the causal model is semi-Markovian because the variables  $A$  can remain dependent even after conditioning on all observable variables because of the hidden confounders. Second,

---

<sup>17</sup>Indirect causal paths all go through proxy variables.

as counterfactual error rates (Section 5.5) are expressed in terms of the true outcome  $Y$ , they are not recommended in case the ground truth is not available and the true outcome is not reliable.

Finally, note that ETT, ATT, and ATC are not generally used in fairness scenarios because, typically, the bias can be observed in both directions: when considering a disadvantaged group/individual as advantaged or the opposite. ETT is relevant when studying the effect of a treatment medicine on patients. For example, if a patient agrees to take the medicine and it turns out to be painful, she may be wondering about the chances of recovery if she did not take the treatment or if she took it with a lower dose. The opposite direction (the effect of treating an individual in the control group) is irrelevant in this case.

In his book, *The Book of Why* [14], Pearl describes a causation ladder with three rungs: statistical observations (seeing), intervention (doing), and counterfactual (imagining). In this section, all causal-based fairness notions defined in Pearl’s SCM framework (all notions in Section 5 except ATE, ATT, ATC, ITE, and equality of effort) are placed in the causation ladder, which will help us address the problem of their applicability in real-scenarios. The causation ladder is structured in such a way that a quantity at a certain rung can be identified in terms of quantities at the rung just below it. Consequently, the higher the rung, the more challenging the problem of identifiability is; hence, the less applicable a fairness notion is defined at that rung.

The diagram in Figure 13 shows the causation ladder and indicates at which rung every causal-based fairness notion stands.  $TV$  which is the only non-causal fairness notion covered in this paper is at rung 1. It is always applicable, provided that a set of observations (dataset) is available. No unresolved and non-discrimination criteria are placed midway between rungs 1 and 2 as they are applicable provided that the causal graph is available along the dataset. The non-discrimination criterion, however, requires the Markov property to be applicable because causal dependence through unobservable paths cannot be blocked. It also has an exponential complexity since it considers all combinations of values of the parent variables of the outcome  $Y$ . A relaxation is described by the authors [45], but the notion remains computationally intractable.

## 7.2. Applicability

Fairness notions at rung 2 (TE, No-proxy discrimination, interventional and justifiable fairness, and individual direct discrimination) are applicable in any scenario where either experiments (RCT) are possible or hypothetical interventions are identifiable. As mentioned in Section 6.1.1, in Markovian models any intervention prob-

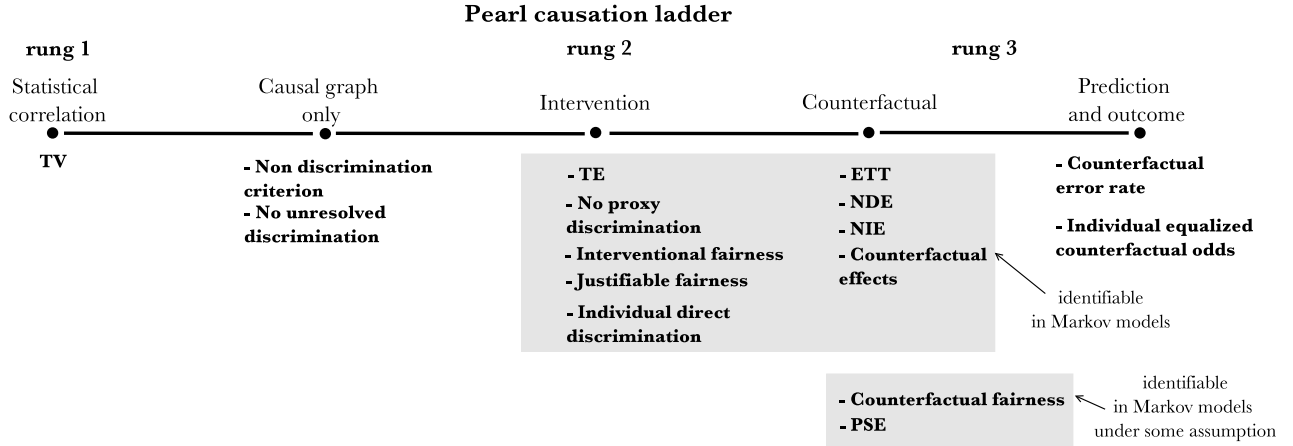


Figure 13: Classification of causal-based fairness notions according to Pearl causation ladder [14]

ability is identifiable from observational data. Hence, these fairness notions are always applicable to Markovian models. In semi-Markovian models, the applicability of these rung 2 notions depends on the identifiability of the intervention terms used in their respective definitions. For instance, the term in question for individual direct discrimination is  $CE(q_k, q'_k)$  in Eq. 27.

The bulk of causal-based fairness notions are defined in terms of counterfactual quantities and hence are placed in rung 3 of the causation ladder. In Figure 13, the counterfactual notions are ranked from top to bottom according to their degree of applicability. For instance, counterfactual effects are placed on top of counterfactual fairness to indicate that the former is applicable in more scenarios than the latter. In Markovian models, the top 4 notions (ETT, NDE, NIE, and counterfactual effects) are always identifiable and hence applicable. That is, specific formulas are already available to compute each counterfactual term used in their definitions.

In Markovian models, the identifiability of counterfactual fairness depends on the identifiability of the term  $\mathbb{P}(y_{a_1} | \mathbf{X} = \mathbf{x}, A = a_0)$  which is only identifiable if  $\mathbf{X}$  does not contain any variable which is at the same time descendant of  $A$  and ancestor of  $Y$ , that is,  $\mathbf{X} \cap \mathbf{B} = \emptyset$  where  $\mathbf{B} = An(Y) \cap De(A)$ [56]. *PSE* is applicable provided that the model is Markovian and the recanting witness criterion is not satisfied. In semi-Markovian models, unless all model parameters are known (including  $P(\mathbf{u})$ )<sup>18</sup>, the identifiability of rung 3 fairness notions depends on the criteria discussed in

<sup>18</sup>In that case, it is possible to use the three steps abduction, action, and prediction [7].

Section 6.1.2, which rarely hold in practice.

Finally, counterfactual error rate and individual equalized counterfactual odds are special cases of rung 3 fairness notions as they are the only notions that condition on the true outcome  $Y$  to assess the fairness of the prediction  $\hat{Y}$  (Eq. 23,24, 25, and 33). Such conditioning has an important implication on identifiability since  $Y$  is a collider, and conditioning on a collider creates a dependence between the previous variables [7]. This leads to unobservable confounding between the causes of  $Y$ . Hence, even if the causal model is Markovian, applying both notions turns it into a semi-Markovian model. Zhang and Bareinboim [42] define an identifiability criterion for counterfactual error rate in Markovian models called the explanation criterion.

## 8. Discussion

As a summary, this paper describes the situations where the causality approach is relevant for evaluating fairness and how it can be achieved using causal-based fairness notions and approximation techniques. However, the applicability of the causal framework is limited because of its reliance on prior knowledge and often untestable assumptions. Indeed, many causal requirements can be achieved by designing appropriate experiments (RCT [80]). However, in fairness scenarios, it is often not a plausible option. Consequently, discrimination is usually evaluated from observational datasets. In the following, we list and discuss further requirements for applying causal inference that are most relevant for fairness applications.

First, the goal of causal approaches to fairness is to measure the impact of a sensitive attribute (e.g. gender) on the outcome (e.g. job hiring). Most definitions of causal effect are based on a notion of intervention or manipulation of a cause variable (exposure) [81, 82]. This makes it hard to justify causal claims related to “immutable” quantities, such as sensitive attributes, for example, race or gender. One alternative is shifting attention from an actual manipulation to changes in perception [83]. For example, instead of changing the gender of candidates to estimate the effect on a hiring decision, one could manipulate the perception of gender by the employer. For example, this could be done by submitting two analogous resumes but varying the name or title of the candidate. This methodology has been adopted in social experiments on the impact of race or gender of an applicant on hiring decisions [84].

Second, the SUTVA [85] (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 [86]. 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 [87]. In fairness evaluation, treatment is replaced by the sensitive attribute, which is often a social construct such as race or gender. 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 in [88].

Third, one of the important impediments to the use of causal-based approaches is the unavailability of the true causal graph which indicates the causal relations between variables [27, 13, 28]. Causal graphs can be set manually by experts in the field, but very often generated using experiments (called also interventions). The process of identifying the causal graph is called causal discovery or structure learning. The gold standard of causal discovery is to perform randomized controlled trials (RCT) [80]. However, RCT is generally not feasible because of practical, ethical, and scalability reasons. For instance, assessing the causal relation of the gender  $A$  on the hiring decision  $Y$  requires changing the gender of a randomly selected job candidate. Even if such intervention/experiment was possible, there are practical concerns about the number of experiments to be performed. As an alternative to RCT, causal discovery is typically carried out using statistical tests on observable data. A large number of causal discovery algorithms exist in the literature. The most commonly used algorithms are PC [89] and its FCI extension [89] which belong to the constraint-based category, GES [90] which belongs to the score-based category, and directLiNGAM [91] which belong to the category of procedures that exploit semi-parametric assumptions. In the context of causal discovery for fairness, Binkyte et al. [92] show how different causal discovery approaches may result in different causal models and, most importantly, how even slight differences between causal models can have significant impact on fairness/discrimination conclusions.



## 9. Conclusion

Notions of fairness that are inconsistent with the causal relationships in the data can lead to misleading conclusions about bias and discrimination of the outcomes. In particular, using causal reasoning to tackle the problem of fairness in machine learning has at least three advantages. First, it appropriately measure discrimination in presence of statistical anomalies (e.g., Simpson’s paradox). Second, it provides natural interpretation of causal relationships between variables in support of discrimination claims. This is particularly important in the disparate treatment legal framework. Third, it makes it possible to break down the dependence between the sensitive attribute and the outcome into different paths (direct, indirect, etc.) which allows to assess fairness more accurately in presence of acceptable and unacceptable discrimination.

Most of the causal-based notions of fairness examined in this paper rely on the availability of the causal graph. The issue of generating causal graphs consistent with the observed data is a known problem in the causal inference literature. Studying it for the specific context of machine learning fairness is a relevant direction for future work.

## 10. Acknowledgements

This work was supported by the European Research Council (ERC) project HYPATIA under the European Union’s Horizon 2020 research and innovation programme. Grant agreement n. 835294.

## References

- [1] R. B. Darlington, Another look at “cultural fairness”, *Journal of educational measurement* 8 (2) (1971) 71–82.
- [2] M. Hardt, E. Price, N. Srebro, Equality of opportunity in supervised learning, in: *Advances in neural information processing systems*, Spain, 2016, pp. 3315–3323.
- [3] E. H. Simpson, The interpretation of interaction in contingency tables, *Journal of the Royal Statistical Society: Series B (Methodological)* 13 (2) (1951) 238–241.
- [4] P. J. Bickel, E. A. Hammel, J. W. O’Connell, Sex bias in graduate admissions: Data from berkeley, *Science* 187 (4175) (1975) 398–404.

- [5] J. R. Loftus, C. Russell, M. J. Kusner, R. Silva, Causal reasoning for algorithmic fairness, arXiv preprint arXiv:1805.05859 (2018).
- [6] S. Barocas, A. D. Selbst, Big data’s disparate impact, *Calif. L. Rev.* 104 (2016) 671.
- [7] J. Pearl, *Causality*, Cambridge university press, 2009.
- [8] M. J. Kusner, J. Loftus, C. Russell, R. Silva, Counterfactual fairness, in: *Advances in neural information processing systems, USA*, 2017, pp. 4066–4076.
- [9] J. Zhang, E. Bareinboim, Fairness in decision-making—the causal explanation formula, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018.
- [10] B. Salimi, L. Rodriguez, B. Howe, D. Suci, Interventional fairness: Causal database repair for algorithmic fairness, in: *Proceedings of the 2019 International Conference on Management of Data*, 2019, pp. 793–810.
- [11] G. W. Imbens, D. B. Rubin, *Causal inference in statistics, social, and biomedical sciences*, Cambridge University Press, 2015.
- [12] I. Shpitser, J. Pearl, Identification of conditional interventional distributions, in: *Proceedings of the Twenty-Second Conference on Uncertainty in Artificial Intelligence*, 2006, pp. 437–444.
- [13] R. Guo, L. Cheng, J. Li, P. R. Hahn, H. Liu, A survey of learning causality with data: Problems and methods, *ACM Computing Surveys (CSUR)* 53 (4) (2020) 1–37.
- [14] J. Pearl, D. Mackenzie, *The book of why: the new science of cause and effect*, Basic Books, 2018.
- [15] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, R. Zemel, Fairness through awareness, in: *Proceedings of the 3rd innovations in theoretical computer science conference*, ACM, USA, 2012, pp. 214–226.
- [16] A. Chouldechova, Fair prediction with disparate impact: A study of bias in recidivism prediction instruments, *Big data* 5 (2) (2017) 153–163.
- [17] I. Zliobaite, A survey on measuring indirect discrimination in machine learning, arXiv preprint arXiv:1511.00148 (2015).

- [18] R. Berk, H. Heidari, S. Jabbari, M. Kearns, A. Roth, Fairness in criminal justice risk assessments: The state of the art, *Sociological Methods & Research* (2018) 0049124118782533.
- [19] S. Verma, J. Rubin, Fairness definitions explained, in: 2018 IEEE/ACM International Workshop on Software Fairness (FairWare), IEEE, 2018, pp. 1–7.
- [20] A. Asuncion, D. Newman, UCI machine learning repository (2007).
- [21] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, A. Galstyan, A survey on bias and fairness in machine learning, *ACM Computing Surveys (CSUR)* 54 (6) (2021) 1–35.
- [22] S. Mitchell, E. Potash, S. Barocas, A. D’Amour, K. Lum, Algorithmic fairness: Choices, assumptions, and definitions, *Annual Review of Statistics and its Application* 8 (2021) 141–163.
- [23] K. Makhoul, S. Zhioua, C. Palamidessi, Machine learning fairness notions: Bridging the gap with real-world applications, *Information Processing & Management* 58 (5) (2021) 102642.
- [24] N. Kilbertus, M. R. Carulla, G. Parascandolo, M. Hardt, D. Janzing, B. Schölkopf, Avoiding discrimination through causal reasoning, in: *Advances in Neural Information Processing Systems*, 2017, pp. 656–666.
- [25] R. Binkytė, S. Zhioua, On the need and applicability of causality for fair machine learning, arXiv preprint arXiv:2207.04053 (2022).
- [26] W. Huan, Y. Wu, L. Zhang, X. Wu, Fairness through equality of effort, in: *Companion Proceedings of the Web Conference 2020*, ACM, USA, 2020, pp. 743–751.
- [27] P. Spirtes, K. Zhang, Causal discovery and inference: concepts and recent methodological advances, in: *Applied Informatics*, Vol. 3, SpringerOpen, 2016, pp. 1–28.
- [28] D. Malinsky, D. Danks, Causal discovery algorithms: A practical guide, *Philosophy Compass* 13 (1) (2018) e12470.
- [29] R. Nabi, I. Shpitser, Fair inference on outcomes, in: *Proceedings of the... AAAI Conference on Artificial Intelligence*. AAAI Conference on Artificial Intelligence, Vol. 2018, NIH Public Access, 2018, p. 1931.

- [30] M. Rhee, Impact: The dcps evaluation and feedback system for school-based personnel (2019).
- [31] K. Quick, The unfair effects of impact on teachers with the toughest jobs, The Century Foundation (2015).
- [32] C. O’Neill, Weapons of math destruction, How Big Data Increases Inequality and Threatens Democracy (2016).
- [33] J. Pearl, Causal analysis in theory and practice, <http://causality.cs.ucla.edu/blog/index.php/2012/12/03/judea-pearl-on-potential-outcomes> (2012).
- [34] S. L. Morgan, C. Winship, Counterfactuals and causal inference, Cambridge University Press, 2015.
- [35] L. Yao, Z. Chu, S. Li, Y. Li, J. Gao, A. Zhang, A survey on causal inference, arXiv preprint arXiv:2002.02770 (2020).
- [36] C. Glymour, K. Zhang, P. Spirtes, Review of causal discovery methods based on graphical models, *Frontiers in genetics* 10 (2019) 524.
- [37] A. Khademi, S. Lee, D. Foley, V. Honavar, Fairness in algorithmic decision making: An excursion through the lens of causality, in: *The World Wide Web Conference*, ACM, USA, 2019, pp. 2907–2914.
- [38] A. Castelnovo, R. Crupi, G. Greco, D. Regoli, I. G. Penco, A. C. Cosentini, A clarification of the nuances in the fairness metrics landscape, *Scientific Reports* 12 (1) (2022) 4209.
- [39] J. Pearl, Direct and indirect effects, in: *Proceedings of the Seventeenth conference on Uncertainty in Artificial Intelligence*, 2001, pp. 411–420.
- [40] S. Chiappa, Path-specific counterfactual fairness, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33, PKP Publishing Services Network, 2019, pp. 7801–7808.
- [41] Y. Wu, L. Zhang, X. Wu, H. Tong, Pc-fairness: A unified framework for measuring causality-based fairness, in: *Advances in Neural Information Processing Systems*, 2019, pp. 3404–3414.

- [42] J. Zhang, E. Bareinboim, Equality of opportunity in classification: A causal approach, in: *Advances in Neural Information Processing Systems*, 2018, pp. 3671–3681.
- [43] L. Zhang, Y. Wu, X. Wu, Situation testing-based discrimination discovery: A causal inference approach., in: *IJCAI*, Vol. 16, 2016, pp. 2718–2724.
- [44] M. Bendick, Situation testing for employment discrimination in the united states of america, *Horizons stratégiques* (3) (2007) 17–39.
- [45] L. Zhang, Y. Wu, X. Wu, Achieving non-discrimination in data release, in: *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2017, pp. 1335–1344.
- [46] A. Romei, S. Ruggieri, A multidisciplinary survey on discrimination analysis, *The Knowledge Engineering Review* 29 (5) (2014) 582–638.
- [47] S. R. Pfohl, T. Duan, D. Y. Ding, N. H. Shah, Counterfactual reasoning for fair clinical risk prediction, in: *Machine Learning for Healthcare Conference*, 2019, pp. 325–358.
- [48] D. Galles, J. Pearl, Testing identifiability of causal effects, in: *Proceedings of the Eleventh conference on Uncertainty in Artificial Intelligence*, ACM, 1995, pp. 185–195.
- [49] J. Tian, J. Pearl, A general identification condition for causal effects, in: *Aaai/iaai*, 2002, pp. 567–573.
- [50] J. Tian, J. Pearl, On the identification of causal effects, 2015.  
URL <https://api.semanticscholar.org/CorpusID:16596796>
- [51] J. Tian, Identifying linear causal effects, in: *AAAI*, 2004, pp. 104–111.
- [52] Y. Huang, M. Valtorta, Identifiability in causal bayesian networks: A sound and complete algorithm, in: *Proceedings of the national conference on Artificial Intelligence*, Vol. 21, Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, AAAI Press, London, 2006, p. 1149.
- [53] I. Shpitser, J. Pearl, Complete identification methods for the causal hierarchy, *Journal of Machine Learning Research* 9 (Sep) (2008) 1941–1979.
- [54] I. Shpitser, J. Pearl, What counterfactuals can be tested, in: *23rd Conference on Uncertainty in Artificial Intelligence, UAI 2007*, 2007, pp. 352–359.

- [55] I. Shpitser, Counterfactual graphical models for longitudinal mediation analysis with unobserved confounding, *Cognitive science* 37 (6) (2013) 1011–1035.
- [56] Y. Wu, L. Zhang, X. Wu, Counterfactual fairness: Unidentification, bound and algorithm., in: *IJCAI*, 2019, pp. 1438–1444.
- [57] C. Avin, I. Shpitser, J. Pearl, Identifiability of path-specific effects, in: *Proceedings of the 19th international joint conference on Artificial Intelligence*, 2005, pp. 357–363.
- [58] L. Zhang, Y. Wu, X. Wu, A causal framework for discovering and removing direct and indirect discrimination, in: *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, 2017, pp. 3929–3935.
- [59] L. Zhang, X. Wu, Anti-discrimination learning: a causal modeling-based framework, *International Journal of Data Science and Analytics* 4 (1) (2017) 1–16.
- [60] D. Malinsky, I. Shpitser, T. Richardson, A potential outcomes calculus for identifying conditional path-specific effects, *Proceedings of machine learning research* 89 (2019) 3080.
- [61] R. A. Fisher, *Statistical methods for research workers*, in: *Breakthroughs in statistics*, Springer, 1992, pp. 66–70.
- [62] S. Tikka, J. Karvanen, Enhancing identification of causal effects by pruning, *The Journal of Machine Learning Research* 18 (1) (2017) 7072–7094.
- [63] S. Tikka, J. Karvanen, Simplifying probabilistic expressions in causal inference, *The Journal of Machine Learning Research* 18 (1) (2017) 1203–1232.
- [64] S. D. Team, et al., Rstan: the r interface to stan, *R package version 2* (1) (2016) 522.
- [65] L. F. Wightman, Lsac national longitudinal bar passage study. lsac research report series. (1998).
- [66] M. Redmond, A. Baveja, A data-driven software tool for enabling cooperative information sharing among police departments, *European Journal of Operational Research* 141 (3) (2002) 660–678.
- [67] P. R. Rosenbaum, D. B. Rubin, The central role of the propensity score in observational studies for causal effects, *Biometrika* 70 (1) (1983) 41–55.

- [68] M. J. Funk, D. Westreich, C. Wiesen, T. Stürmer, M. A. Brookhart, M. Davidian, Doubly robust estimation of causal effects, *American journal of epidemiology* 173 (7) (2011) 761–767.
- [69] B. K. Lee, J. Lessler, E. A. Stuart, Weight trimming and propensity score weighting, *PloS one* 6 (3) (2011) e18174.
- [70] J. M. Alvarez, S. Ruggieri, Counterfactual situation testing: Uncovering discrimination under fairness given the difference, in: *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, 2023, pp. 1–11.
- [71] E. A. Stuart, Matching methods for causal inference: A review and a look forward, *Statistical science: a review journal of the Institute of Mathematical Statistics* 25 (1) (2010) 1.
- [72] J. Angwin, J. Larson, S. Mattu, L. Kirchner, Machine bias. *propublica*, See <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing> (2016).
- [73] D. Freedman, R. Pisani, R. Purves, *Statistics*, 3rd edn., pp. a-107 (1998).
- [74] S. Athey, G. Imbens, Recursive partitioning for heterogeneous causal effects, *Proceedings of the National Academy of Sciences* 113 (27) (2016) 7353–7360.
- [75] S. Ben-David, J. Blitzer, K. Crammer, F. Pereira, et al., Analysis of representations for domain adaptation, *Advances in neural information processing systems* 19 (2007) 137.
- [76] S. R. Künzel, J. S. Sekhon, P. J. Bickel, B. Yu, Metalearners for estimating heterogeneous treatment effects using machine learning, *Proceedings of the national academy of sciences* 116 (10) (2019) 4156–4165.
- [77] Y. Shimoni, E. Karavani, S. Ravid, P. Bak, T. H. Ng, S. H. Alford, D. Meade, Y. Goldschmidt, An evaluation toolkit to guide model selection and cohort definition in causal inference, *arXiv preprint arXiv:1906.00442* (2019).
- [78] X. Zhou, T. Yamamoto, Tracing causal paths from experimental and observational data (2020).
- [79] C. Andrade, Mean difference, standardized mean difference (smd), and their use in meta-analysis: As simple as it gets, *The Journal of clinical psychiatry* 81 (5) (2020) 0–0.

- [80] R. A. Fisher, Design of experiments, *British Medical Journal* 1 (3923) (1936) 554.
- [81] T. J. VanderWeele, M. 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* (2012) 101–113.
- [82] P. W. Holland, Statistics and causal inference, *Journal of the American statistical Association* 81 (396) (1986) 945–960.
- [83] A. Rahmattalabi, A. Xiang, Promises and challenges of causality for ethical machine learning, *arXiv preprint arXiv:2201.10683* (2022).
- [84] D. Neumark, Experimental research on labor market discrimination, *Journal of Economic Literature* 56 (3) (2018) 799–866.
- [85] D. B. Rubin, Statistics and causal inference: Comment: Which ifs have causal answers, *Journal of the American Statistical Association* 81 (396) (1986) 961–962.
- [86] S. L. Morgan, C. Winship, *Counterfactuals and causal inference*, Cambridge University Press, 2015.
- [87] L. Keele, The statistics of causal inference: A view from political methodology, *Political Analysis* 23 (3) (2015) 313–335.
- [88] L. Laffers, G. Mellace, Identification of the average treatment effect when sutva is violated, *Discussion Papers on Business and Economics*, University of Southern Denmark 3 (2020).
- [89] P. Spirtes, C. Meek, T. Richardson, An algorithm for causal inference in the presence of latent variables and selection bias, *Computation, causation, and discovery* 21 (1999) 211–252.
- [90] A. Hauser, P. Bühlmann, Characterization and greedy learning of interventional Markov equivalence classes of directed acyclic graphs, *Journal of Machine Learning Research* 13 (2012) 2409–2464.  
URL <https://jmlr.org/papers/v13/hauser12a.html>



- [91] S. Shimizu, T. Inazumi, Y. Sogawa, A. Hyvärinen, Y. Kawahara, T. Washio, P. O. Hoyer, K. Bollen, Directlingam: A direct method for learning a linear non-gaussian structural equation model, *The Journal of Machine Learning Research* 12 (2011) 1225–1248.
- [92] R. Binkytė-Sadauskienė, K. Makhlouf, C. Pinzón, S. Zhioua, C. Palamidessi, Causal discovery for fairness, *arXiv preprint arXiv:2206.06685* (2022).