The Zoo of Fairness Metrics in Machine Learning

Alessandro Castelnovo

Intesa Sanpaolo

Riccardo Crupi

Intesa Sanpaolo

Greta Greco

Intesa Sanpaolo

Daniele Regoli (✉ daniele.regoli@sns.it)

Intesa Sanpaolo

Ilaria Giuseppina Penco

Intesa Sanpaolo

Andrea Claudio Cosentini

Intesa Sanpaolo

Research Article

Keywords: zoo, fairness, Machine Learning, Artificial Intelligence

Posted Date: December 29th, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-1162350/v1>

License: This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

The zoo of fairness metrics in Machine Learning

Alessandro Castelnovo^{1,+}, Riccardo Crupi^{1,+}, Greta Greco^{1,+}, Daniele Regoli^{1,*,+}, Ilaria Giuseppina Penco¹, and Andrea Claudio Cosentini¹

¹Data Science and Artificial Intelligence, Intesa Sanpaolo S.p.A., Torino, Italy

*daniele.regoli@intesasanpaolo.com

+these authors contributed equally to this work

ABSTRACT

In recent years, the problem of addressing fairness in Machine Learning (ML) and automatic decision making has attracted a lot of attention in the scientific communities dealing with Artificial Intelligence. A plethora of different definitions of fairness in ML have been proposed, that consider different notions of what is a “fair decision” in situations impacting individuals in the population. The precise differences, implications and “orthogonality” between these notions have not yet been fully analyzed in the literature. In this work, we try to make some order out of this zoo of definitions.

Introduction

The issue of discrimination bias in Artificial Intelligence (AI), in particular in Machine Learning (ML), has gained a lot of momentum in the scientific community in the last few years¹⁻⁴. The establishment of dedicated conferences such as the ACM Fairness, Accountability and Transparency (FAccT) conference (Formerly ACM FAT*, grown out of the FAT/ML workshops held from 2014 to 2018), has fostered the emergence of a new community of researchers working on topics that are sometimes roughly collected under the label of Trustworthy AI^{5,6} (also Responsible AI or Ethical AI) — topics such as fairness and bias in ML, and explainability of “black-box” ML models, that focus on possible risks and drawbacks of the use of ML at scale in automated decision making.

Contextually, many research institutions, private companies and public sector organizations have started issuing and taking positions on ethical principles and guidelines for the responsible use of AI systems⁷. Indeed, besides the benefits, many ethical concerns have emerged in recent years with the widespread use of such systems: bias amplification, data privacy, lack of transparency, human oversight, accountability, etc.^{5,8-13}. The European Commission has recently published a proposal for what is going to be the first attempt ever to insert AI systems and their use in a coherent legal framework¹⁴: the proposal explicitly refers to the risk of bias discrimination of AI systems. Critics to the “non-responsible” employment of AI systems and to the possible harms for people and society have also reached a broader audience thanks to popular books and documentaries¹⁵⁻¹⁸, raising awareness on these topics.

It is now widely recognized that data-driven decision making is not *per se* safe from producing *unfair* or *biased* decisions, either for amplifications of biases already present in the data they learn from, or for algorithmic inaccuracies^{4,16,19-21}. The scientific literature addressing the problem of fairness in ML has focused in the recent years on two main aspects:

- i) how to measure and assess fairness (or, equivalently, bias), and
- ii) how to mitigate bias in models when necessary.

Regarding the first aspect, in the last few years an incredible number of definitions have been proposed, formalizing different perspectives from which to assess and monitor fairness in decision making processes. A popular tutorial presented at the Conference on Fairness, Accountability, and Transparency in 2018 was titled “21 fairness definitions and their politics”²². The number has grown since then. The proliferation of fairness definitions is not *per se* an issue: it reflects the evidence that fairness is a multi-faceted concept, and concentrates on itself different meanings and nuances, in turn depending in complex ways on the specific situation considered. As it is often the case with moral and ethical issues, conflicts are present between different and typically equally reasonable positions²³⁻²⁵.

Researches, in proposing definitions, have focused on different intuitive notions of “unfair decisions”, often considered as the ones impacting people in different ways on the basis of some personal characteristics, such as gender, ethnicity, age, sexual or political or religious orientations, considered to be *protected*, or *sensitive*. Relationships and interdependence of these sensitive variables with other features useful for making decisions lie at the very heart of these ambiguities and nuances, and entangle in complex ways the ultimate aim of any model, i.e. making efficient decisions, and the desired goal of not allowing unfair discrimination to impact people’s lives.

Fairness notions proposed in the literature are usually classified in broad areas, such as: definitions based on parity of statistical metrics across groups identified by different values in protected attributes (e.g. male and female individuals, or people in different age groups); definitions focusing on preventing different treatment for individuals considered similar with respect to a specific task; definitions advocating the necessity of finding and employing causality among variables in order to really disentangle unfair impacts on decisions.

These three broad classes can be further seen as the result of two important distinctions:

- (a) observational vs. causality-based criteria;
- (b) group (or statistical) vs. individual (or similarity-based) criteria.

Distinction (a) discriminates criteria that are purely based on observational distribution of the data from criteria that try to first unveil causal relationships among the variables at play in a specific situation (mainly through a mixture of domain knowledge and opportune inference techniques) and then assess fairness. Distinction (b) discriminates criteria that focus on equality of treatment among groups of people from criteria requiring equality of treatment among couples of similar individuals.

Regarding aspect ii), i.e. that of removing bias when present, several approaches have been introduced in order to provide bias mitigation in data-driven decisions. Roughly speaking, they fall into three families, depending on the specific point along the ML pipeline in which they operate: pre-processing methods, that try to remove bias directly from data; in-processing techniques, imposing fairness either as a constraint or as an additional loss term during training optimization; post-processing methods, working directly on the outcomes of the model.

The present work focuses on the first of these aspects, and in particular deals with the aforementioned fact that the number of different metrics for fairness introduced in the literature has boomed. The researcher or practitioner first approaching this facet of ML may easily feel confused and somehow lost in this zoo of definitions. These multiple definitions capture different aspects of the concept of fairness but, at the best of our knowledge, there is still no clear understanding of the overall landscape where these metrics live. This work aims to take a step in the direction of analyzing the relationship among the metrics, and trying to put order in the fairness landscape. Table 1 provide a rough schematic list of fairness metrics discussed throughout the paper. As general references for the definition of fairness in ML we refer to the book by Barocas, Hardt and Narayanan¹ and to more compact surveys^{2-4,26-29}.

In this work, we shall mainly deal with the most common problem in supervised ML, the binary classification with a single sensitive attribute, and we shall only make a brief reference to the additional subtleties arising when multiple sensitive attributes are present. Despite this huge simplifications, the landscape of fairness definitions is nevertheless extremely rich and complex.

In the literature, some attempts to survey existing fairness definitions in ML are present^{1,3,4,26}, but, to the best of our knowledge, none has focused solely on metrics trying to disentangle and analyze in depth the interdependence and relationships among them. The contribution of this work lies not so much in the exhaustiveness of the taxonomy as in the attempt to clarify the nuances in the fairness metrics landscape with respect to the twofold dimensions listed above, namely group vs. individual notions and observational vs. causality-based notions.

The rest of this paper is organized as follows: section 1 proposes a concise account of the most important sources of bias in ML models; section 2 introduces individual fairness; section 3 is devoted to the most populated set of definitions, namely the group (or statistical) metrics; causality-based criteria are discussed in section 4.

1 The problem of bias in data-driven decisions

Most sources of bias in data-driven decision making lie in the data itself and in the way in which they are collected. It is out of the scope of this manuscript to list and discuss all the possible shades of biases that can hide in the data^{4,31-34}, we want nevertheless to say something about the main cases that can be encountered in many real-life scenarios.

In our view, the most important qualitative distinction lies between *statistical* or *representation biases*^{3,4,19}, i.e. — roughly speaking — when the model learns from biased sampling from the population, and *historical* and *societal biases*⁴, i.e. when the model learns from data impacted by decisions made on the basis of (cultural and social) prejudices.

1.1 Statistical bias

Generally speaking, *statistical bias* occurs whenever the data used for model training are not representative of the true population. This can be due to a form of *selection bias*, i.e. when the individuals appearing in the data come from a non-random selection of the full population. This happens, for example, in the context of credit lending, where the information of the repayment is known only for people that were granted the loan.

Another way in which statistical bias can enter the data is via *systematic measurement errors*. This happens when the record of past errors and performance is systematically distorted, especially in the case of different amount of distortion for

Table 1. Fairness metrics. Qualitative schema of the most important fairness metrics discussed throughout the paper.

	notion	use of Y	condition
group fairness	Demographic Parity	-	equal acceptance rate across groups
	Conditional Demographic Parity	-*	equal acceptance rate across groups in any strata
	error parity	Equal Accuracy	✓ equal accuracy across groups
		Equality of Odds	✓ equal FPR and FNR across groups
		Predictive Parity	✓ equal precision across groups
individual fairness	FTU/Blindness	-	no explicit use of sensitive attributes
	Fairness Through Awareness	-*	similar people are given similar decisions
causality-based fairness	Counterfactual Fairness	-	an individual would have been given the same decision if she had had different values in sensitive attributes
	path-specific Counterfactual Fairness	-	same as above, but keeping fixed some specific attributes

* there are exceptions to these cases where Y is actually employed, e.g. CDP conditioning on Y becomes Equality of Odds, and there are notions of individual fairness that use a similarity metric defined on the target space³⁰.

different groups of people. Similarly, it may happen that data are systematically missing or poorly recorded for entire strata of the population.

One aspect of paramount importance, but often omitted, is the following: in the context of fairness, *even in samples perfectly representative of the population* it may be that some protected group happens to be a minority group. All in all, minorities exist. In this case, especially when the minority group is highly unbalanced, a ML model can learn less accurately or may learn to discard errors on that group simply because it is less important in terms of crude counting. Again, notice that this is not a form of statistical bias *tout court*, since in this case the “under-representativeness” of the protected group is a true feature of the population. Nevertheless, this is something very common, and one of the most important source of unfair discrimination against protected groups that happen to be minorities³⁴.

This last issue may be reasonably considered as a source of biased decisions due to poor modeling rather than to a problem in data. The boundary is subtle, and in any case this is a matter of nomenclature: the important thing is to be aware of the risk.

1.2 Historical or societal bias

Even when the data are free from statistical bias, i.e. they truly represent the population, take into account minorities and there is no systematic error in recording, still it may be that bias exists simply because data reflect biased decisions.

In most cases, this is due to a form of *labelling bias*, i.e. a systematic favour/disfavour towards groups of people at the moment of creating the target variable from which the model is going to learn. If the recorded outcomes are somehow due to human decisions (e.g. a model for granting loans may be trained on loan officers’ past decisions), then we cannot in general trust their objectiveness and “fairness”.

Other forms of historical bias may be even more radical: gender bias has a rather long history, and is embedded in all sorts of characteristics and features in such a way that it is difficult or even impossible to evaluate its impact and disentangle its dependence on other variables. Think for example of income or profession disparities, just to name a few out of many. Thus, this is a situation in which long-lasting biases cause systematic differences in features pertaining different groups of people. Again, this is not a form of un-representativeness of the sample, it is a bias present in the full population.

Finally, bias can lurk in the data in a variety of additional ways^{1,4,19,34}, e.g. a poor selection of features may result in a loss of important information in disproportionate ways across groups¹⁹ — among many others.

2 Individual fairness

2.1 Toy examples and notation

To be as clear as possible, in what follows we shall make several references to 2 toy examples: credit lending decisions, and job recruiting. We call A the categorical random variable representing the protected attribute, that for reference we shall take to be gender, we label with X all the other (non-sensitive) random variables that the decision maker is going to use to provide its yes/no decisions $\hat{Y} = f(X, A) \in \{0, 1\}$; while we label $Y \in \{0, 1\}$ the ground truth target variable that needs to be estimated/predicted – typically by minimizing some loss function $\mathcal{L}(Y, \hat{Y})$. $\tilde{X} = (\tilde{X}_1, \dots, \tilde{X}_d) = (X, A)$ collectively represent all the d features involved in the problem. We denote with lowercase letters the specific realizations of random variables, e.g. $\{(x_1, y_1), \dots, (x_n, y_n)\}$ represent a dataset of n independent realizations of (X, Y) . We employ calligraphic symbols to refer to domain spaces, namely $\mathcal{X}$ denotes the space where features X live. (Technically, since we employ uppercase letters to denote random variables, it would be more proper to say that $\mathcal{X}$ is the image space of $X : \Omega \rightarrow \mathcal{X}$, where Ω is the event space, and that $x \in \mathcal{X}$. We shall nevertheless use this slight abuse of notation for sake of simplicity).

2.2 Similarity-based criteria

Individual fairness is embodied in the following principle: *similar individuals should be given similar decisions*. This principle deals with the comparison of single individuals rather than focusing on groups of people sharing some characteristics. On the other hand, group fairness starts from the idea that there are *groups of people* potentially suffering biases and unfair decisions, and thus tries to reach equality of treatment for groups instead of individuals.

The first attempt to deal with a form of individual fairness is by Dwork, Hardt, Pitassi, Reingold and Zemel³⁵, who introduce the concept as a Lipschitz condition of the map f from the feature space to the model space:

$$\text{dist}_Y(\hat{y}_i, \hat{y}_j) < L \times \text{dist}_{\tilde{X}}(\tilde{x}_i, \tilde{x}_j), \quad (1)$$

where dist_Y and $\text{dist}_{\tilde{X}}$ denote a suitable distance in target space and feature space, respectively, and L is a constant. Loosely speaking, *a small distance in feature space (i.e. similar individuals) must correspond to a small distance in decision space (i.e. similar outcomes)*. They also propose an approach to enforce such kind of fairness by introducing a constrained optimization at training time, once a distance on the feature space is given.

The concept of individual fairness is straightforward, and it certainly resonates with our intuitive notion of “equality”. Formula (1) provides also an easy way to assess whether the decision model $\hat{Y} = f(\tilde{X})$ satisfies such concept. However, the major drawback of this type of definitions lies in the subtle concept of “similar individuals”. Indeed, defining a suitable distance metric $\text{dist}_{\tilde{X}}$ on feature space embodying the concept of similarity on “ethical” grounds alone is almost as tricky as defining fairness in the first place.

Indeed, while for target space the natural choice $\text{dist}_Y(\hat{Y}_i, \hat{Y}_j) = |\hat{Y}_i - \hat{Y}_j|$ will do, for feature space the natural choice of the euclidean distance on the space $\mathcal{A} \times \mathcal{X}$ implies that any smooth function f shall satisfy condition (1), but this is not a very interesting notion of fairness. The point is that one should come up with a distance that captures the essential features for determining that target, and that do not mix with sensitive attributes. But, again, this is not very different from defining what is fairness in a specific situation. Take, e.g., the job recruiting framework: what identifies the couples of individuals that should be considered similar and thus given the same chance of being recruited? Maybe the ones with same level of skills and experience irrespective of anything else?

One possibility is to define similar individuals as couples belonging to different groups with respect to sensitive features (e.g. male and female) but with the same values for all the other features. With this choice, what we are requiring is that its outcome should be unchanged if we take an observation and we only change its protected attribute A . This concept is usually referred to as *Fairness Through Unawareness (FTU)* or *blindness*²⁶, and it is expressed as the *requirement of not explicitly employing protected attributes in making decisions*.

Notice that this idea is very likely one of the first that one may think of when asked for, e.g., a decision making process that does not discriminate against gender: not to explicitly use gender to make decisions, i.e. $\hat{Y} = f(X)$. Indeed, this concept is also referred to as *disparate treatment*, i.e. there is disparate treatment whenever two individuals sharing the same values of non-sensitive features but differing on the sensitive ones are treated differently^{19,36}.

Unfortunately, despite its compelling simplicity, fairness through unawareness comes not without flaws. Firstly, if it is true that reaching FTU is straightforward, it is not that easy to assess. The problem of *bias assessment* in a model consists in measuring whether the model decisions are biased given a set of realizations of $(A, X, \hat{Y})$, namely $\{(a_1, x_1, \hat{y}_1), \dots, (a_n, x_n, \hat{y}_n)\}$ – the corresponding values of Y are needed as well for some criteria. In this setting, it is tricky to measure whether FTU holds, the main reason being that it is more a request on how the model works rather than a request on properties of the output

decisions. One possible candidate metric is the following³⁷:

$$\text{consistency} = 1 - \frac{1}{n} \left(\sum_{i=1}^n \left| \hat{y}_i - \frac{1}{k} \sum_{x_j \in kNN(x_i)} \hat{y}_j \right| \right). \quad (2)$$

Namely, for each observation $(x_i, \hat{y}_i)$ it measures how much the decision $\hat{y}_i$ is close to the decisions given to its k nearest neighbors $kNN(x_i)$ in the $\mathcal{X}$ space (notice that also in the computation of kNN one has to choose a distance function on $\mathcal{X}$). It may happen that the k neighbors of, say, a male individual are all males: in this case consistency (2) would in fact be equal to 1, but this does not prevent the model from explicitly using A in making decisions. Another possibility would be to compute³⁰

$$\frac{1}{n_1 n_2} \sum_{\substack{a_i=1, \\ a_j=0}} e^{-\text{dist}(x_i, x_j)} |\hat{y}_i - \hat{y}_j|, \quad (3)$$

which measures the difference in decisions among men and women weighted by their similarity in feature space: the higher its value the higher the difference in treatment for couples of similar males and females (the term $e^{-\text{dist}(x_i, x_j)}$ can be substituted with any measure of similarity of the points x_i and x_j). On the other hand, it is much easier to assess FTU if we also have access to the model: we could create a synthetic dataset by flipping A (this holds if A is binary, but it is easy to come up with generalizations to the multiclass case): $\{(a'_1, x_1), \dots, (a'_n, x_n)\}$, feeding it to the model to get the corresponding outcomes $\hat{y}'_1, \dots, \hat{y}'_n$, and then compute the average $\frac{1}{n} \sum_{i=1}^n |\hat{y}_i - \hat{y}'_i|$.

However, the main drawback of FTU is the following: it does not take into account the possible interdependence between A and X . Other features may contain information on the sensitive attribute, thus explicitly removing the sensitive attribute is not sufficient to remove its information from the dataset. Namely, it may be that in the actual dataset there is a very low chance that a male and a female have similar values in all the (other) features, since gender is correlated with some of them. One may or may not decide that (some of) these correlations are legitimate (e.g. gender may be correlated with income, but, depending on the problem, one may decide that the use of income, even if correlated with gender, is not a source of unfair discrimination; in section 4 we will dwell more on this issue, namely that there may be some information correlated to the sensitive attribute but still considered “fair”). This is one of the reasons why the definition of a issue-specific distance is crucial for a more refined notion of individual fairness.

One straightforward way to deal with correlations is to develop a model that is blind not only with respect to the sensitive attribute, but also to all the other variables with sufficiently high correlation with it. This is a method known as *suppression*³⁸. Apart from the obvious issue in defining a good threshold for correlation above which a predictor should be removed, this approach has the main drawback in the potentially huge loss of legitimate information that may reside in features correlated with the sensitive attribute.

One line of research has tried to solve this issue by “cleaning” the training dataset in order to remove both the sensitive attribute *and* the information coming from the sensitive attribute “lurking” inside other features, while *keeping the features*. This can be done, e.g., by using residuals of regressions of the (correlated) features on the sensitive attribute^{39,40}, i.e. projecting the feature space in a subspace orthogonal to A . However, it is tricky to account for dependencies on feature interactions. Another way of keeping information while removing “unfairness” due to the sensitive attribute is to learn a *fair representation* of the training dataset^{37,41–43}, i.e. to learn a new set of variables Z , such that they are able to reconstruct X with as few errors as possible, while at the same time being as independent as possible of A (in the sense that it is hard to reconstruct A from Z). The idea is then to use any ML model on this “cleaned” representation of the original dataset obtaining fair-by-design outcomes. Notice, however, that in most situations, when trying to remove from the data the dependence of A and X , individual fairness will be spoiled. We shall discuss this in section 3.6.

Counterfactual frameworks^{44,45}, that will be more thoroughly discussed in section 4, provide a clear way in which this similarity should be thought of: a male individual is similar to himself in the counterfactual world where he is a woman. Notice that this is crucially different from fairness through unawareness approach: a male individual transported in the counterfactual world where he is a woman will have differences in other features as well, and such differences are precisely due to the causal structure among the variables (very roughly speaking, this is the “causal way” to account for correlations). As an example, in the now popular Judea Pearl’s assessment of the 1973 UC Berkeley admission case^{46–48} where there was different admission rate between men and women, being a woman “causes” the choice of higher demanding departments, thus impacting the admission rate (in causal theory jargon, department choice is said to be a *mediator* from gender to admission)¹. In this case, a male transported in the counterfactual world where he is a woman would have himself chosen higher demanding departments, thus this feature would change as well. In this case, a simple gender-blind model would not guarantee individual fairness. Indeed attribute flipping in general does not produce valid counterfactuals⁴⁹. More details on this will be given in section 4.

Turning to the general similarity-based definition (1), some work has been done to address the problem of finding a suitable similarity metrics in feature space. E.g.³⁵ and^{50,51} introduce the possibility to learn a issue-specific distance from data and from the contributions of domain experts. Indeed, the simple idea of using standard similarities, e.g. related to the euclidean distance on feature space, does not take into account the trivial fact that some feature are more important than other in determining the relationship of an individual to specific target. Namely, for two applicants for a loan, the difference in income is much more important than the difference in, say, age, or even profession. Thus, judging what does it mean to be similar *with respect to a specific task* is not that simple, and is in some sense connected also to the ground truth target variable Y .

Indeed, some have proposed³⁰ a notion of individual fairness as a penalty function – to be added to the risk minimization fitting – which shifts the concept of similarity from the feature space to the target space: namely two individuals are deemed similar if they have a similar value of the target variable. Thus, this notion of individual fairness relies directly on the target attribute to define a task-specific distance. Of course, this definition is prone to biases possibly present in the target.

Broadly speaking, the Lipschitz notion (1) is sometimes referred to as Fairness Through Awareness (FTA), as opposed to Fairness Through Unawareness: in fact, even if they share the same principle of treating equally similar individuals, FTA is generally meant to use similarity metrics that are problem and target specific, i.e. that derive from an “awareness” of the possible impact, while FTU is a simple recipe that does not depend on the actual scenario.

3 Group fairness

Group fairness criteria are typically expressed as conditional independence statements among the relevant variables in the problem. There are three main broad notions of observational group fairness, *independence*, *separation*, *sufficiency*¹. Independence is strictly linked to what is known as *Demographic Parity* or *Statistical Parity*, separation is related to *Equality of Odds* and its relaxed versions, while sufficiency is connected to the concept of *calibration* and *Predictive Parity*.

There is a crucial aspect that discriminates independence criteria from the others: *independence criteria rely only on the distribution of features and decisions, namely on $(A, X, \hat{Y})$, while separation and sufficiency criteria deal with error rate parities, thus making use of the target variable Y as well*. This is an important thing to have in mind when trying to find your way in the zoo of fairness criteria in a specific case study: when using separation-based or sufficiency-based criteria one must be careful to check whether the target variable Y is itself prone to sources of bias and unfairness.

In the example of credit lending, if Y represents the decision of a loan officer, than separation-based and sufficiency-based criteria must be used with particular care, since Y can be itself biased against some groups of people. Actually, even if Y stands for the repayment (or lack of it) of the loan, a form of *selection bias* is very likely at work: we have that information only on applicants that received the loan in the first place, and these are almost surely not representative of the whole population of applicants.

Sometimes the term *disparate mistreatment* is used to refer to all group fairness criteria relying on disparity of errors, thus, in general, all group metrics dealing with comparisons among decisions $\hat{Y}$ and true outcomes Y ³⁶.

Moreover, independence is said to be a *non-conservative* measure of fairness⁵² — meaning that it forces to change the *status quo* — since it is in general not satisfied by the perfect classifier $\hat{Y} = Y$ (apart from the trivial case when $Y \perp\!\!\!\perp A$). Error rate parities, on the other hand, are *conservative*, since they trivially hold for the perfect predictor.

3.1 Independence

The criterion of independence¹ states that the *decisions should be independent of any sensitive attribute*:

$$\hat{Y} \perp\!\!\!\perp A. \quad (4)$$

This can be also expressed as follows — in the binary classification assumption:

$$P(\hat{Y} = 1 \mid A = a) = P(\hat{Y} = 1 \mid A = b), \quad \forall a, b \in \mathcal{A}, \quad (5)$$

i.e. the ratio of loans granted to men should be equal to the ratio of loans granted to women.

The ratio of favorite outcomes is sometimes known as positive prediction ratio (ppr), thus independence is equivalent to requiring the same positive prediction ratio across groups identified by the sensitive features. This form of independence is usually known as *Demographic Parity (DP)*, *statistical parity*, or sometimes as *group fairness*. DP is a very common concept in the literature on fairness in ML^{1,53}, and it was actually employed even before the now common nomenclature was introduced⁵⁴. Figure 1 shows a very simple visualization of a model reaching demographic parity among men and women.

If some group has a significantly lower positive prediction ratio with respect to others, we say that demographic parity is not satisfied by the model decisions $\hat{Y}$. In order to have a single number summarizing the amount of disparity, it is common to use either maximum possible difference or the minimum possible ratio of positive prediction ratios: a difference close to 0 or

Figure 1. example of demographic parity in gender in credit lending toy model.

a ratio close to 1 indicate a decision system *fair* with respect to A in the sense of DP. Typically, some tolerance is considered by employing a threshold below (or above) which the decisions are still considered acceptable.

In the literature, it is common to cite the “4/5 rule” or “80% rule”: the selection rate of any group should be not less than 4/5 than the one of the group with the highest selection rate, i.e DP ratio greater than 80%. This is a reference to the guidelines by the US Equal Employment Opportunity Commission (EEOC)^{55,56}, and it is often cited as one of the few examples of a legal framework relying on a specific definition of fairness, in particular on the notion of Demographic Parity. However, it is usually overlooked that the guidelines explicitly state that this should be considered only as a rule of thumb, with many possible exceptions, and in particular they prescribe that business or job-related necessities can justify a lower DP ratio, depending on the circumstances. This is actually a step away from pure DP towards more individual notions, such those of Conditional Demographic Parity and error parities, that shall be discussed in sections 3.1.2, 3.2, 3.3.

3.1.1 Subtleties of Demographic Parity

The meaning of DP is intuitive only to a superficial analysis. For example, one may at first think that removing the sensitive attribute from the decision making process is enough to guarantee independence and thus demographic parity. In general, this is not the case. Take the credit lending example and assume that, for whatever reason, women tend to actually pay back their loans with higher probability with respect to men. If this is the case, it is reasonable to assume that a rating variable that we call R will be higher for women than for men, on average. In this scenario, the sensitive attribute gender (A) is correlated with rating. If the model uses only rating (but not gender) to compute its decision, there will be a higher rate of loans granted to women than to men, resulting in a demographic *disparity* among these groups. In this case, if you want to reach DP, the model needs to *favor men over women*, granting loans to men with a lower rating threshold with respect to women.

Thus, because of interdependence of X and A , not only it is not enough to remove the sensitive attribute from the decision making process, but if you want to have DP you need, in general, to *treat different groups in different ways*, precisely in order to compensate for the (unwanted) effect of this dependence. This is somehow the opposite of an intuitive notion of fairness! Figure 2 displays this example: to reach equal ppr, one must use a different rating threshold for each group.

This, in turn, reveals another subtlety: even if your dataset and your setting does not apparently contain any sensitive features, discrimination could lurk in via correlations to sensitive features that you are not even collecting.

We want to stress here what we think is a crucial aspect: another possibility to reach DP would be to use *neither gender nor rating in making decisions*, i.e. trying to remove all gender information from the dataset. Notice that this could be problematic for the accuracy of your decisions, since it’s plausible that, by removing all variables correlated to A , information useful to estimate the target is lost as well. This is called *suppression*, and was discussed already in section 2, together with the concept of *fair representation*, by which one tries to remove all sensitive information from the dataset while keeping as much useful information as possible.

Another possible flaw is the following: if it is true that women repay their loans with higher probability, is it really fair to have demographic parity between men and women? Why the bank should agree to grant loans with the same rate to groups

Figure 2. example of a subtlety of demographic parity: in order to reach demographic parity between men and women and still using rating as fundamental feature, one must use a different threshold between the groups, thus manifestly treating differently men and women.

that actually pay back their loans with different probabilities? Should we stick to actual repayment rates or we should ask *why* these probabilities are different and if this is possibly due to gender discrimination in the first place?

We try to summarize, in a non exhaustive list, a set of scenarios in which it might be reasonable to take into account demographic parity, among other metrics:

- when you want to actively *enforce* some form of equality between groups, irrespective of all other information. Indeed there are some characteristics that are widely recognized to be independent *in principle* of sensitive attributes, e.g. intelligence and talent, and there may be the need to enforce an independence in problems where the decisions is fundamentally linked to those characteristics; more in general, there may be reasons to consider unfair any relation among A and Y , even if the data (objective data as well) is telling differently;
- (intertwined with the previous) when you deem that in your specific problem there are hidden historical biases that impact in a complex way the entire dataset;
- when you cannot trust the objectivity of the target variable Y , then demographic parity still makes sense, while many other metrics don't (e.g. separation).

Justifying the use of DP as a metric with the idea of enforcing an equality that should hold in principle has been discussed in the literature^{52,57,58}. We would like to spend here a word of care in this respect, namely on the fact that favouring a group does indeed guarantee that the equality is reached in that specific use case, but it may trigger mechanisms that actually *amplify* the bias that we intended to break. In the example of job recruiting, imposing the decision maker to increase the acceptance rate on a group of people with (relatively) low skill levels, will increase the likelihood that applicants hired from this group will perform poorly relative to the average, effectively perpetuating the bias in the way in which this group is perceived by others. In general, assessing the future impact of imposing a fairness condition, i.e. considering fairness as a dynamical process rather than a one-shot condition, is a subtle issue by itself^{59–61}.

3.1.2 Conditional demographic parity

Another version of the independence criteria is that of *Conditional Demographic Parity (CDP)*⁶². In the example given above, we may think that a fairer thing to do, with respect to full independence, is to require independence of the decision on gender only for men and women with the same level of rating. In other words, if a man and a woman have both a certain level of rating, we want them to have the same chance of getting the loan. This goes somewhat in the direction of being an individual form of fairness requirement, since parity is assessed into smaller groups with respect to the entire sample.

Formally, this results in requiring $\hat{Y}$ independent of A given R ,

$$\hat{Y} \perp\!\!\!\perp A \mid R, \quad (6)$$

or, in other terms:

$$P(\hat{Y} = 1 \mid A = a, R = r) = P(\hat{Y} = 1 \mid A = b, R = r), \quad \forall a, b \in \mathcal{A}, \forall r. \quad (7)$$

This seems a very reasonable requirement in many real-life scenarios. For example, if you think of a recruitment setting where you don't want to bias women against men, but still you want to recruit the most skilled candidates, you may require your decision to be independent of gender but conditional on a score based on the curriculum and past work experiences: among people with an “equivalent” set of skills, you want to recruit men and women with the same rate. In other words, the only disparities that you are willing to accept between male and female candidates are those justified by curriculum and experience.

Unfortunately, also in this case one must be really careful, because the variable that you are conditioning on *might itself be a source of unfair discrimination*. For example, it might well be that rating is higher for women not because they actually pay back their debts more likely than men, but because the rating system is biased against men. And this may be due to self-fulfilling prophecy: if men have lower rating they may receive loans with higher interest rates, and thus have a higher probability of not paying them back, in a self-reinforcing loop.

Moreover, it may be not so straightforward to select the variables to condition on: why condition on rating and not, e.g., on level of income, or profession, etc...?

Finally, in line with what outlined above for demographic parity, one may argue that women's curricula and work experiences tend to be on average different than those of male candidates for historical reasons and for a long-lasting (and die-hard) man-centered society. This may suggest that plain demographic parity could be more appropriately enforced in this case to reach a “true” equality.

Incidentally, notice that pushing to the extreme the notion of conditional demographic parity, i.e. conditioning on *all the (non-sensitive) variables*, one has

$$\hat{Y} \perp\!\!\!\perp A \mid X, \quad (8)$$

i.e.

$$P(\hat{Y} = 1 \mid A = a, X = x) = P(\hat{Y} = 1 \mid A = b, X = x), \quad \forall a, b \in \mathcal{A}, \forall x \in \mathcal{X}; \quad (9)$$

meaning that a male and a female *with the same value for all the other features* must be given the same outcome. This criterion is reachable by a gender-blind model, thus is strictly connected to the notion of FTU. As mentioned above – and in more details in section 3.6 – conditioning is equivalent to restrict the groups of people among which we require parity, thus is a way to go in the direction of obtaining an individual criterion.

3.2 Separation

Independence and conditional independence do not make use of the true target Y . What if, instead of conditioning over rating R we condition on the target Y ? This is equivalent to requiring the independence of the decision $\hat{Y}$ and gender A separately for individuals that actually repay their debt and for individuals that don't. Namely, among people that repay their debt (or don't), we want to have the same rate of loan granting for men and women.

This concept has been called separation¹:

$$\hat{Y} \perp\!\!\!\perp A \mid Y. \quad (10)$$

In other terms

$$P(\hat{Y} = 1 \mid A = a, Y = y) = P(\hat{Y} = 1 \mid A = b, Y = y), \quad \forall a, b \in \mathcal{A}, y \in \{0, 1\}. \quad (11)$$

Equivalently, *disparities in groups with different values of A (male and female) should be completely justified by the value of Y (repayment or not)*.

As in the conditional independence case, this seems a very reasonable fairness requirement, *provided that you can completely trust the target variable*. Namely, one should be extremely careful to check whether the target Y is not itself a source of bias.

For example, if Y instead of reflecting true repayment was the outcome of loan officers' decision on whether to grant the loans, it could incorporate bias, thus we it would be risky to assess fairness with direct comparisons with Y . Moreover, as

Figure 3. Example of Equality of Odds between men and women: false negative and false positive rates must be equal across groups.

said above, even in the objective case where Y is the actual repayment, a form of selection bias would likely distort the rate of repayment.

We can express separation in terms of what are known in statistics as *type I* and *type II* errors. Indeed, it is easy to see that the two conditions in equation (11) (one for $y = 1$ and one for $y = 0$) are equivalent to requiring that the model has *the same false positive rate and false negative rate across groups identified via A* . False positives and false negatives are precisely type I and type II errors, respectively. Namely, individuals that are granted loans but are not able to repay, and individuals that are able to repay but are not granted loans. This is known as *Equality of Odds*⁶³, and is thus the requirement of having the same type I and type II error rates across relevant groups, as displayed in Figure 3.

There are two relaxed version of this criterion:

- *Predictive Equality*: equality of false positive rate across groups,

$$P(\hat{Y} = 1 \mid A = a, Y = 0) = P(\hat{Y} = 1 \mid A = b, Y = 0), \quad \forall a, b \in A,$$

- *Equality of Opportunity*: equality of false negative rate across groups,

$$P(\hat{Y} = 0 \mid A = a, Y = 1) = P(\hat{Y} = 0 \mid A = b, Y = 1), \quad a, b \in \mathcal{A}.$$

While demographic parity, and independence in general, focuses on equality in terms of acceptance rate (loan granting rate), Equality of Odds, and separation in general, focuses on equality in terms of error rate: the model is fair if it is as efficient in one group as it is in the other.

The difference between Predictive Equality and Equality of Opportunity is the perspective from which equality is required: Predictive Equality takes the perspective of people that won't repay the loan, while Equality of Opportunity takes the one of people that will repay. Depending on the problem at hand, one may consider either of these two perspectives as more important. For example, Predictive Equality may be considered when we want to minimize the risk of innocent people from being erroneously arrested: in this case it may be reasonable to focus on the parity of among innocents. Both in the credit lending and job listing examples, on the other hand, it may be reasonable to focus on Equality of Opportunity, i.e. on the parity among people that are indeed deserving.

Here follows a non exhaustive set of situations in which separation criteria may be suitable:

- when your target variable Y is an objective ground truth;
- when you are willing to make discrimination as long as they are justified by actual trustable data;

- when you do not want to actively enforce an “ideal” form of equality, and you want to be as equal as possible given the data.

We can summarize by saying that separation, being a concept of parity given the ground truth outcome, is a notion that takes the point of view of people that are subject to the model decisions, rather than that of the decision maker. In the next subsection, instead, we shall take into account the other side of the coin, i.e. parity given the model decision.

3.3 Sufficiency

Sufficiency¹ takes the perspective of people that are given the same model decision, and requires parity among them irrespective of sensitive features.

While separation deals with error rates in terms of fraction of errors over the ground truth, e.g. the number of individuals whose loan request is denied among those who would have repaid, sufficiency takes into account the number of individuals who won't repay among those who are given the loan.

Mathematically speaking, this is the same distinction you have between recall (or true positive rate) and precision, i.e. $P(\hat{Y} = 1 | Y = 1)$ and $P(Y = 1 | \hat{Y} = 1)$, respectively.

A fairness criterion that focuses on this type of error rate is called *Predictive Parity*⁵³, also referred to as *outcome test*^{3,26}:

$$P(Y = 1 | A = a, \hat{Y} = 1) = P(Y = 1 | A = b, \hat{Y} = 1), \quad \forall a, b \in \mathcal{A}, \quad (12)$$

i.e. the model should have the same precision across sensitive groups. If we require condition (12) to hold for the case $Y = 0$ as well, then we get the following conditional independence statement:

$$Y \perp\!\!\!\perp A | \hat{Y},$$

which is referred to as *sufficiency*¹.

Predictive Parity, and its more general form of sufficiency, focuses on error parity among people who are given the same decision. In this respect, Predictive Parity takes the perspective of the decision maker, since they group people with respect to the decisions rather than the true outcomes. Taking the credit lending example, the decision maker is indeed more in control of sufficiency rather than separation, since parity given decision is something directly accessible, while parity given truth is known only in retrospect. Moreover, as we have discussed above, the group of people who are given the loan ($\hat{Y} = 1$) is less prone to selection bias than the group of people who repay the loan ($Y = 1$): indeed we can only have the information of repayment for the $\hat{Y} = 1$ group, but we know nothing about all the others ($\hat{Y} = 0$).

As you may notice, going along a similar reasoning, one can define other group metrics, such as Equality of Accuracy across groups: $P(\hat{Y} = Y | A = a) = P(\hat{Y} = Y | A = b)$, for all $a, b \in \mathcal{A}$, i.e. focusing on unconditional errors, and others²⁶.

3.4 Group fairness on scores

In most cases, even in classification setting, the actual output of a model is not a binary value, but rather a *score* $S \in \mathbb{R}$, estimating the probability, for each observation, to have the target equal to the favorable outcome (usually labelled 1). Then, the final decision is made by the following t -threshold rule:

$$\hat{Y} = \begin{cases} 1 & S \geq t, \\ 0 & S < t. \end{cases} \quad (13)$$

Most of the things we have outlined for group fairness metrics regarding $\hat{Y}$ can be formulated for the joint distribution (A, Y, S) as well:

$$\begin{aligned} \text{independence} & \quad S \perp\!\!\!\perp A, \\ \text{separation} & \quad S \perp\!\!\!\perp A | Y, \\ \text{sufficiency} & \quad Y \perp\!\!\!\perp A | S. \end{aligned} \quad (14)$$

These formulations provide stronger constraints on the model with respect to the analogous with $\hat{Y}$, e.g. the condition $S \perp\!\!\!\perp A$ is sometimes referred to as *strong demographic parity*^{2,64}.

For instance, all three criteria in their form (5), (11), (12), i.e. with constraints on the joint distribution $(A, Y, \hat{Y})$ can effectively be satisfied by defining group dependent thresholds t on black-box model outcomes S (a technique that goes under the name of *postprocessing*¹), while this is not the case for (14).

Instead of requiring conditions on the full distribution of S as in (14), something analogous to the “binary versions” of group fairness criteria have been defined simply by requiring parity of the *average score*⁶⁵. *Balance of the Negative Class*, defined as

$$E(S \mid Y = 0, A = a) = E(S \mid Y = 0, A = b), \quad \forall a, b \in \mathcal{A}, \quad (15)$$

corresponds to Predictive Equality, while *Balance of the Positive Class* (same as (15) with $Y = 1$) corresponds to Equality of Opportunity. Notice that these two last definitions fall into the one given in subsection 3.2 when $S = \hat{Y}$. Requiring both balances is of course equivalent to requiring EO.

AUC parity, namely the equality of the area under the ROC for different groups identified by A , can be seen as the analogous of the equality of accuracy.

Finally, notice that the score formulation of sufficiency is connected to the concept of *calibration*. Calibration holds when

$$P(Y = 1 \mid S = s) = s, \quad (16)$$

i.e. if the model assigns a score s to 100 people then, on average, $100 \times s$ of them will actually be positive. Writing down the condition for sufficiency in score:

$$P(Y = 1 \mid S = s, A = a) = P(Y = 1 \mid S = s, A = b), \quad \forall a, b \in \mathcal{A}, \forall s, \quad (17)$$

it is clearly related to what can be called *Calibration within Groups*⁶⁵:

$$P(Y = 1 \mid A = a, S = s) = s \quad (18)$$

– actually a consequence of it. This condition is in general not so hard to achieve, and it is often satisfied “for free” by most models¹.

3.5 Incompatibility statements

It is interesting to analyse the relationships among different criteria of fairness. In subsection 3.6 we shall discuss in detail the connections between individual and group notions, while here we focus on differences among various group criteria. We have already seen that each of them highlights one specific aspect of an overall idea of fairness, and we may wonder what happens if we require to satisfy multiple of them at once. The short answer is that it is not possible except in trivial or degenerate scenarios, as stated by the following propositions drawn from the literature^{1,52,53,65,66}:

1. if Y is binary, $Y \not\perp S$ and $Y \not\perp A$, then **separation and independence** are incompatible.

In other words, to achieve both separation and independence, the only possibility is that either the model is completely useless ($Y \perp S$), or the target is independent of the sensitive attribute ($Y \perp A$), which implies an equal base rate for different sensitive groups. Namely, if there is an imbalance in groups identified by A , then you cannot have both EO and DP holding.

2. Analogously: if $Y \not\perp A$, then **sufficiency and independence** cannot hold simultaneously.

Thus, if there is an imbalance in base rates for groups identified by A , then you cannot impose both sufficiency and independence.

3. Finally, if $Y \not\perp A$ and the distribution (A, S, Y) is strictly positive, then **separation and sufficiency** are incompatible.

Meaning that separation and sufficiency can both hold either when there is no imbalance in sensitive groups (i.e. the target is independent of sensitive attributes), or when the joint probability (A, S, Y) is degenerate, i.e. — for binary targets — when there are some values of A and S for which only $Y = 1$ (or $Y = 0$) holds, in other terms when the score exactly resolves the uncertainty in the target (as an example, the perfect classifier $S = Y$ always trivially satisfies both sufficiency and separation).

Notice that proposition 3 reduces to the following, more intuitive statement if the classifier is also binary (e.g. when $S = \hat{Y}$): if $Y \not\perp A$, Y and $\hat{Y}$ are binary, and there is at least one false positive prediction, then separation and sufficiency are incompatible. Moreover, it has been shown⁶⁵ that *Balance of Positive Class*, *Balance of Negative Class* and *Calibration within Groups* can hold together only if either there is no imbalance in groups identified by A or if each individual is given a perfect prediction (i.e. $S \in \{0, 1\}$ everywhere).

When dealing with incompatibility of fairness metrics, the literature often focuses on the 2016 COMPAS — a recidivism prediction instrument developed by Northpoint Inc. — recidivism case²⁰, now become a case study in the fairness literature. Indeed, the debate on this case is a perfect example to highlight the fact that there are *different and non-compatible notions of fairness*, and that this may have concrete consequences on people. While we refer to the literature for a thorough discussion of the COMPAS case^{53,67}, we here just point out that in the debate there were two parties, one stating that the model predicting recidivism was *fair* since it satisfied Predictive Parity by ethnicity, while the other claiming it was *unfair* since it had different false positive and false negative rates for black and white individuals. Chouldechova showed that⁵³, if $Y \not\perp A$, i.e. if the true recidivism rate is different for black and white people, then Predictive Parity and Equality of Odds cannot both hold, thus implying that a reflection on which of the two (in general of the many) notions is more important to be pursued in that specific case must be carefully considered.

Summarizing, apart from trivial or peculiar scenarios, the three families of group criteria above presented are not mutually compatible.

3.6 Group vs individual

The most common issue with group fairness definitions is the following: since group fairness requires to satisfy conditions only on average among groups, it leaves room to bias discrimination *inside the groups*. As we argued in the example of section 3 referring to figure 2, one way of reaching DP is to use a different rating threshold for men and women: this means that there will be a certain range of ratings for which men will receive the loan, and women won't. More formally, conditionally on rating there is no independence among A and $\hat{Y}$. In general, to reach group fairness one may “fine tune” the interdependence of A and X to reach parity on average, but effectively producing differences in subgroups of A .

Notice that this is precisely what individual fairness is about. In the example above, a men and a woman that have the same rating may be treated differently, thus violating the individual notion of fairness.

As already discussed in section 3, this is only one possibility: one may as well reach DP by using neither gender nor rating, and grant loan on the basis of other information, provided it is independent of A . In this case, there will be no group discrimination *and* there won't be any subgroup discrimination as well.

However, we can say that, if we want to reach DP by using as much information of Y as possible contained in (X, A) , i.e. minimizing the risk $E\mathcal{L}(f(X, A), Y)$, then it is unavoidable to have some form of disparate treatment among people in different groups with respect to A whenever $X \not\perp A$. *Fair affirmative action*³⁵ is the name given to the process of requiring DP while trying to keep as low as possible the amount of disparate treatment between people having similar X .

To clarify the general picture, we can put the different notions of (observational) fairness in a plane with two *qualitative* dimensions (see figure 4): 1. to what extent a model is fair at the individual level, 2. how much information of A is retained in making decisions. The first dimension represents to what extent two individuals with similar overall features X are given similar decisions: the maximum value is reached by models blind on A (FTU). These are the models that are also using all information in X , irrespective of the interdependence with A , thus FTU-compliant models will use all information contained in $\tilde{X}$ apart from the information that is contained in A only. The minimum value in this dimension is reached by models that satisfy DP. Models using suppression methods, being blind to both A and other features with high correlation with A , are individually fair in the sense of preventing disparate treatment. In so doing, they can exploit more or less information of A with respect to general DP-compliant models depending on how many correlated variables are discarded. However, the price to pay for discarding variables is in terms of errors in approximating Y , which is not highlighted in this plot. Notice that, of course, full suppression – i.e. removing all variables dependent on A – trivially satisfies the condition $\hat{Y} \perp A$, i.e. it is DP-compliant as well. In figure 4, we label with DP a general model that tries to maximize performance while satisfying a DP constraint, without any further consideration.

Models satisfying CDP are somewhat in-between, of course depending on the specific variables considered for conditioning. They guarantee less disparate treatment than unconditional DP, and they use more information of A by controlling for other variables possibly dependent on A .

Notice that approaches such as fair representation (see section 2), where you try to remove all information of A from X to get new variables Z which are as close as possible to being independent of A , produce decision systems $\hat{Y} = f(Z)$ that are not, in general, individually fair. This is due to the simple fact that, precisely to remove the interdependence of A and X while keeping as much information of X as possible, two individuals with same X and different A will be mapped in two distinct points on $\mathcal{Z}$, thus having, in general, different outcomes. Referring again to the credit lending example, suppose that we have $R = g(A) + U$, with g a complicated function encoding the interdependence of rating and gender, and U some other factor independent of A representing other information in R “orthogonal” to A . In this setting, the variable Z that we are looking for is precisely U . Notice that U is indeed independent of A , thus any decision system $\hat{Y} = f(U)$ satisfies DP, but given two individuals with $R = r$ nothing prevents them from having different values of U . In other terms, you *need* to have some amount of disparate treatment to guarantee DP *and* employ as much information as possible to estimate Y .

Figure 4. Landscape of observational fairness criteria with respect to the group-vs-individual dimension and the amount of information of A used (via X).

Figure 5. Landscape of observational fairness criteria with respect to the model performance dimension and the amount of information of A used (via X).

Figure 5 shows a *qualitative* representation of observational metrics with respect to the amount of information of A (through X) that is used by the model, and the predictive performance⁶⁸. Notice that DP can be reached in many ways: e.g. a constant score model, namely a model accepting with the same chance all the individuals irrespective of any feature, is DP-compliant (incidentally, it is also individual), a model in which all the variables dependent on A have been removed (a full suppression), or a model where DP is reached while trying to maximize performance (e.g. through fair representations). All these ways differ, in general, in terms of the overall performance of the DP-compliant model.

FTU-compliant models, on the other hand, by employing all information in X will be, in general, more efficient in terms of model performance.

Incidentally, notice that this discussion is to be taken at a qualitative level, one can come up with scenarios in which, e.g., models satisfying DP have higher performances than models FTU-compliant (think, e.g., of a situation in which $Y \perp\!\!\!\perp A$ and $X \not\perp\!\!\!\perp A$).

3.7 Multiple sensitive features

Even if a detailed discussion of the problem of multiple sensitive features is out of the scope of this manuscript, we shall nevertheless give a brief overview.

Generally speaking, all the definitions and results we gave in previous sections are subject to the fact that the sensitive feature A is represented by a single categorical variable. If, for a given problem, we identify more than one characteristic that we need to take into account as sensitive or protected – say $(A_1, \dots, A_I)$ – we can easily assess fairness on each of them separately. Unfortunately this approach — sometimes called *independent group fairness*⁶⁹ — is in general not enough: even if fairness is achieved (in whatever sense) separately on each sensitive variable A_i , it may happen that some subgroups given by the intersection of two or more A_i ’s undergo unfair discrimination with respect to the general population. This is sometimes

referred to as *intersectional bias*⁷⁰, or, more specifically, *fairness gerrymandering*⁷¹.

To prevent bias from occurring in *all the possible subgroups identified by all A_i 's* one can simply identify a new feature $A = (A_1, \dots, A_I)$, whose values are the collection of values on all the sensitive attributes, and require fairness constraints on A — *intersectional group fairness*⁶⁹.

This “trick” indeed solves the problem of intersectional bias, at least theoretically. Still, issues remain at a computational and practical level, whose two main reasons are:

- the exponential increase of the number of subgroups when adding sensitive features,
- the fact that, with finite samples, many of the subgroups will be empty or with very few observations.

These two aspects imply that assessing (group) fairness with respect to multiple sensitive attributes may be unfeasible in most practical cases. Since the presence of many sensitive features is more of a norm than it is an exception, this actually represents a huge problem, that the literature on fairness in ML has barely begun to address^{69,71–73}.

4 Causality-based criteria

Another important distinction in fairness criteria is the one between *observational* and *causality-based* criteria. As we have seen, observational criteria rely only on observed realizations of the distribution of data and predictions. In fact, they focus on enforcing equal metrics (acceptance rate, error rate, etc. . .) for different groups of people. In this respect, they don't make further assumptions on the mechanism generating the data and suggest to assess fairness through statistical computation on observed data.

Causality-based criteria, on the other hand, try to employ domain and expert knowledge in order to come up with a casual structure of the problem, through which it becomes possible to answer questions like “what would have been the decision if that individual had a different gender?”. While counterfactual questions like this seem in general closer to what one may intuitively think of as “fairness assessment”, the observational framework is on the one hand easier to assess and constrain on, and on the other more robust, since counterfactual criteria are subject to strong assumptions about the data and the underlying mechanism generating them^{74,75}.

As we argued above (section 2), answering to counterfactual questions is *very different* from taking the feature vector of, e.g., a male individual and just flip the gender label and see the consequences in the outcome. The difference lies precisely in the causal chain of “events” that this flip would trigger. If there are some features related, e.g., to the length of the hair, or the height, then it is pretty obvious that the flip of gender should come together with a change in these two variables as well, with a certain probability. And this may be the case for other — less obvious but more relevant — variables. This also suggests why counterfactual statements involve *causality relationships* among the variables. In general, to answer counterfactual questions, one needs to know the causal links underlying the problem.

However, as major drawback, once given the casual structure there are many counterfactual models compatible with that structure (actually infinite), and the choice of one of them is in general not falsifiable with observations^{75,76}, since, by definitions, they deal with claims that have not observable consequences.

Indeed, causality-based criteria can be formulated at least at two different levels, with increasing strength of assumptions: at the level of *interventions* and at the level of *counterfactuals*.

4.1 Causal models

Before introducing fairness definitions based on the underlying causal structure of a problem, we briefly introduce the necessary theoretical framework of *causal graphs* and *Structural Causal Models* (SCMs) or Structural Equation Models^{1,44,47,48,76}.

We model the underlying causal relationships among features by means of a Directed Acyclic Graph (DAG) $G = (V, E)$, with V set of vertices (or nodes) and E set of directed edges (or links) — see figure 6. Nodes of the graph G represent the variables $\tilde{X}$ used as predictors in the model. Moreover, we denote with U exogenous or *latent* variables, representing factors not accounted for by the features $\tilde{X}$. In causal graph theory, edges in the graph represent not only conditional dependence relations, but are interpreted as the causal impact that the source variable has on the target variable. To work at the level of *interventions* this is all that is needed. Modeling causal knowledge is complex and challenging since it requires an actual understanding of the relations, beyond statistical evidence. Different causal discovery algorithms have been proposed to identify causal relationships from observational data through automatic methods⁷⁷.

To work at the level of *counterfactuals* one needs much more than the sole graph structure: a Structural Causal Model is a triplet $(\tilde{X}, U, F)$ where F are a set of assignments of the form

$$\tilde{X}_i = f_i(\text{pa}(\tilde{X}_i), U_i), \quad i = 1, \dots, d; \quad (19)$$

Figure 6. Example of causal graph with 3 endogenous variables (left). Intervention on A is expressed via a different graph (right) where all incoming edges in A are removed and the variable A is set to the value a .

where the functions f_i represent the precise way in which the parents of each node variable $\tilde{X}_i$ ($\text{pa}(\tilde{X}_i)$) — namely the variables having a direct causal impact over X_i in the assumed (or inferred) graph G — together with the latent variable U_i , influence the value of X_i . These relations are called Structural Equations (SE) and, besides describing which variables causally impacts which (that is already encoded in the graph G), they also determine *how* these relations work. One typical simplifying assumption on SCM is to work with Additive Noise Models⁷⁶:

$$\tilde{X}_i = g_i(\text{pa}(\tilde{X}_i)) + U_i, \quad i = 1, \dots, d.$$

In general, assumptions like this are needed in order to be able to perform computations using SCM starting from observational data. A detailed presentation of the assumptions underlying the use of graphical models and SCMs (e.g. modularity, markovianity, causal sufficiency, etc. . .) is beyond the scope of this survey^{47,74,76,78}.

As mentioned above, one of the limits of counterfactual models is that, since they make “predictions” regarding something that has not happened, the choice of a specific counterfactual model is in general not falsifiable with observable data^{75,76}.

4.2 Fairness metrics in causality-based settings

Fairness at the level of interventions can be formally expressed as follows⁷⁹:

$$P(\hat{Y} = 1 \mid \text{do}(A = a), X = x) = P(\hat{Y} = 1 \mid \text{do}(A = b), X = x), \quad \forall a, b \in \mathcal{A}, x \in \mathcal{X}. \quad (20)$$

In words, if you take a random individual, force it to be, e.g., female ($\text{do}(A = a)$) and she happens to have $X = x$, you want to give him the same chance of acceptance as for a random individual forced to be male ($\text{do}(A = b)$) that also happens to have $X = x$. We refer to it as *Intervention Fairness*. Incidentally, notice that it may be unfeasible to compute quantities in (20) on the basis of observational data and the graph, without assuming an SCM (or without physically performing the experiment); indeed it may be that the events whose probability needs to be computed in order to get the post-intervention distribution have no observations in the dataset at hand^{47,80}.

The same requirement can be set at the counterfactual level, and is known as *Counterfactual Fairness (CFF)*⁴⁴:

$$P(\hat{Y}_{A \leftarrow a} = 1 \mid A = a, X = x) = P(\hat{Y}_{A \leftarrow b} = 1 \mid A = a, X = x), \quad \forall a, b \in \mathcal{A}, x \in \mathcal{X}, \quad (21)$$

which, in words, reads: if you take a random individual with $A = a$ and $X = x$ and the same individual if she had $A = b$, you want to give them the same chance of being accepted.

The difference between the two levels is subtle but important: roughly speaking, when talking about interventions one is considering the average value over exogenous factors U that, after the intervention, are compatible with the conditioning, while counterfactuals consider only the values of U that are compatible with the factual observation (namely, the distribution $P(U \mid A = a, X = x)$) to begin with, and then perform the intervention and consider the consequences. In other words, counterfactuals consider only events that take into account actual observed value of A (and X as well). As a further clarification⁴⁴, suppose the following structural equation: $X = A + U$. Then equation (20) compares two individuals with $U = x - a$ and $U' = x - b$, i.e. two different individuals that, with the interventions $\text{do}(A = a)$ and $\text{do}(A = b)$ both happen to have the observed value $X = x$. Equation (21) instead take the individual with $U = x - a$ and then perform the interventions $\text{do}(A = a)$ (which has no consequences) and $\text{do}(A = b)$ which implies on $X = b + (x - a) \neq x$.

Equations (20), (21) make use of Pearl’s *do*-calculus^{47,48,80}, where $\text{do}(A = a)$ — the *intervention* — consists in modifying the causal structure of the problem by exogenously setting $A = a$, thus removing any causal paths impacting on A — the theoretical equivalent of randomized experiments (see figure 6 (right)). Effectively, post-intervention variables have a different distribution $P_{\text{do}(A=a)}(X, A, Y)$, and samples (or observations) from this distribution can be drawn only if the intervention

$do(A = a)$ can be physically carried out (e.g. in randomized experiments). In many real life situations this is not feasible, and *do*-calculus provide a mathematical machinery that allows, in some specific cases, to compute quantities related to the (non observable) post-intervention distribution $P_{do(A=a)}(X, A, Y)$ employing only quantities related to the (observable) distribution $P(X, A, Y)$, given the graph — a concept called *identifiability*^{47,48,76,78,80}.

The notation $\hat{Y}_{A \leftarrow b}$, on the other hand, stands for the three steps of counterfactual calculus, *abduction*, *action*, *prediction*^{47,48}:

- *abduction* is where you account for observed values and compute the distribution of $U \mid \{A = a, X = x\}$;
- *action* corresponds to implementing the intervention $do(A = b)$;
- *prediction* consists in using the new causal structure and the exogenous conditional distribution $P(U \mid A = a, X = x)$ to compute the posterior of $\hat{Y}$.

As in the observational setting, causality-based criteria have a group and an individual version: equations (20), (21) correspond to the individual form, but nothing prevents from taking the version unconditional on X , i.e. holding on average on all the individuals, or even to condition on only some subset of X , as in CDP.

4.3 Group vs. individual fairness in causality-based setting

We call *Expectation Intervention Fairness* the condition:

$$P(\hat{Y} = 1 \mid do(A = a)) = P(\hat{Y} = 1 \mid do(A = b)) \quad \forall a, b \in \mathcal{A}, \quad (22)$$

i.e. requiring that, on average, an individual taken at random from the whole population and forced to be woman ($do(A = a)$) should have the same chance of being accepted as a random individual forced to be man. Analogously, we can define *Expectation Counterfactual Fairness (ECFF)* as:

$$P(\hat{Y}_{A \leftarrow a} = 1 \mid A = a) = P(\hat{Y}_{A \leftarrow b} = 1 \mid A = a) \quad \forall a, b \in \mathcal{A}, \quad (23)$$

which states that, on average, the acceptance rate for a random woman ($A = a$) should be the same as the one given to a random woman forced to be a man. Similar definitions can be given conditioning on partial information R .

Summarizing, we can visualize CFF as the following process: the conditioning on some values $\{A = a, R = r\}$ represents the group of people that we take into account (a single individual for $R = X$), and any consideration we do on them must hold on average over that group; given that group, we force a flipping of the sensitive attribute from $A = a$ to $A = b$ (and here we are going in a new, non-observable, distribution), and this will trigger a cascade of causal consequences on the other features X (namely, on the descendants of A in the causal graph). Then, we compare the model outcomes averaged on the observed group and on the counterfactual group.

Intervention Fairness is slightly different: you take again all the individuals compatible with the conditioning – which this time is $R = r$ without fixing the sensitive A – then force them first to have $A = a$, and then to have $A = b$ (with all the causal consequences of these interventions), and require them to have, on average, the same acceptance rate.

This may resonate with the notion of FTU: when flipping A for an observation you want the outcome of the model not to change. Indeed *Fairness Through Unawareness is a causality-based notion of fairness*, at least in its formulation of not explicitly using A . The very concept of flipping A is nothing but an intervention. Notice, incidentally, that the flipping of A *without any impact on other variables* corresponds to assuming a causal graph where A has no descendant, or, better, corresponds to assuming that all the changes caused by the flipping of A on other variables are legitimate, i.e. considered fair. Moreover, the fact that FTU is difficult to be measured without having access to the model, is due to its nature of being a non-observational notion, but it requires fictitious data to be assessed – a dataset with A flipped, i.e. *a dataset not sampled from the real distribution of $(A, X, \hat{Y})$* . This is similar to the way in which CFF can be assessed: compare the predictions over a dataset with the predictions on the same dataset but with A flipped *and* with all the changes caused by this flipping (knowledge of the underlying SCM is thus needed). This, in turn, reveals the subtle difference among CDP in the form $\hat{Y} \perp\!\!\!\perp A \mid X$ and FTU: CDP is an observational notion, i.e. can be measured, in principle, by having access to (realizations from) the distribution of $(\hat{Y}, A, X)$, while FTU does not. We say “in principle” since it’s very likely that in real-world datasets there will be very few observations corresponding to $X = x$ (typically one), thus resulting in a poor estimation of the distribution of $\hat{Y} \mid X$.

Causality-based notions are richer than the observational ones, and permit a further possibility: to select which causal paths from A to $\hat{Y}$ are considered legitimate and which, instead, we want to forbid. In the job listing example, we may consider that the type of degree of the applicants is fundamental for the job position, and thus crucial to making decisions. But it could well be that the type of degree is correlated, and even “caused” by the gender of the applicants: women and men may have different attitudes towards choosing the preferred degree. In this situation, if we require CFF, we would compare a man with

Figure 7. Landscape of causality-based fairness criteria.

some degree to his “parallel self” in the counterfactual world where he is a woman. In that world, however, it’s very likely that her degree is going to be different. Thus, requiring CFF means to somehow prevent the decision maker to employ the degree type for the assessment. Taking into account situations like this one is not very difficult. Supplementary Material of⁴⁴ and⁴⁵ introduce the following definition, known as *path-specific Counterfactual Fairness (pCFF)*:

$$P(\hat{Y}_{A \leftarrow a, X_F \leftarrow x_1} = 1 \mid A = a, X_F = x_1, X_F^c = x_2) = P(\hat{Y}_{A \leftarrow b, X_F \leftarrow x_1} = 1 \mid A = a, X_F = x_1, X_F^c = x_2), \quad \forall a, b \in \mathcal{A}, \forall x_1, x_2, \quad (24)$$

where X_F correspond to the variables descendants of A that we consider fair mediators to make decisions (e.g. the degree type), and X_F^c is its complement with respect to X , namely $X = (X_F, X_F^c)$. In words, take a man ($A = a$) with a specific degree x_1 and other features x_2 , and force him to be a woman, letting all the causal consequences of this flipping to happen *with the exception of the degree, kept fixed at x_1* , then compare their outcomes.

Incidentally, notice that if we allow X_F to contain all the descendants of A , then we end up with a notion that we may call *direct Counterfactual Fairness (dCFF)*, namely the only path that we are concerned of is the direct causal path from gender A to the decision $\hat{Y}$. In this case we don’t let any causal consequences of the gender flipping to happen when assessing fairness. This is, again, strictly connected to FTU.

Figure 7 is a schematic and qualitative visualization of many of the points discussed above: the dimension of group vs. individual is “controlled” by how much information we condition on, while the dimension of the causal impact of A on the decision is controlled by the fraction of paths causally connecting A to $\hat{Y}$ that are considered fair. Figure 7 is meant to be the causal analogue of figure 4. Of course there are many possible cases that, for simplicity, we have omitted from figure 7, namely all the Intervention notions, and the path-specific notions valid on average in broader groups. For instance, one could think of a measure taking into account only the A flip, without causal consequences, and valid on average over X as well (i.e. a group notion), namely $P(\hat{Y}_{A \leftarrow a, X \leftarrow X} = 1 \mid A = a) = P(\hat{Y}_{A \leftarrow b, X \leftarrow X} = 1 \mid A = a), \quad \forall a, b \in \mathcal{A}$.

If we focus on CFF, it can be reached by training a model on the space (Σ, U_A, Y) where Σ are the *non-descendants* of A , i.e. variables not caused by A , *directly or indirectly*; U_A denotes the information inside descendants of A that is not attributable to A . A decision system $\hat{Y} = f(\Sigma, U)$ is counterfactually fair by design (21). This closely reflects what we discussed in section 2 about fair representations: train a model on a new space “cleaned” from all A information. Indeed, in that case we searched for statistical independence of $\hat{Y}$ on A while here we look for “causal independence” of $\hat{Y}$ on A . In this specific sense, counterfactual fairness can be reached via a form of *preprocessing* of the dataset, and is very similar, in spirit, to the concept of fair representation.

We may then wonder why the fair representation approach, reaching DP, is a notion of group rather than individual fairness, while CFF is considered an individual notion: the point is simply that *we employed two different concepts of individual fairness in the observational and in the counterfactual setting*. Namely, in the observational setting we expressed disparate treatment as the event in which two individuals with same X but different A are given different outcomes. In general, fair representation is not guaranteed to prevent such scenario. In the counterfactual context, instead, we consider disparate treatment when an individual and her “counterfactual self” with A flipped are given different outcomes. CFF is designed precisely to prevent this. Notice that this concept, translated in the observational setting, is analogous to requiring a similar treatment for two

individuals similar in $\mathcal{Z}$, not in $\mathcal{X}$. Indeed, if you take an individual with some $\{A = a, X = x\}$ and flip her gender *while taking into account the interdependence of X and A* , she will be transformed in the same point Z .

This is an example of the fact that the concept of individual fairness is strongly dependent on the concept of similarity that one decides to consider.

Summarizing, CFF and fair representation are very similar in the way in which they deal with disparities, namely they both try to remove all (causal) information of A from the feature space, and then use the “cleaned” space to make decisions. In this respect, CFF can be seen as both a “causal analogue” of a DP-compliant methods, in the sense that all the way in which A may impact the decision are forbidden, and an individually fair notion in the sense that it imposes condition on individual basis.

5 Conclusions

The notion of fairness in decision making has many nuances, that have been reflected in the high number of proposed mathematical and statistical definitions. Notice, moreover, that this aspect is not limited to ML or artificial intelligence: the problem of how to define and assess bias discrimination in decision making processes is present largely independently of “who” is making the decision. The growing attention to this issue in the domain of automated data-driven decisions can be imputed to the fact that these processes can potentially amplify biases *at scale*, and could possibly do that without *human oversight*. It is well known that AI systems usually compute their outcomes leveraging very complicated relationships among input features, making not only the users but often also the developer blind with respect to the reasons underlying those outcomes. This issue, that has triggered in the recent years the flourishing stream of research of Explainable AI^{81,82}, is intertwined to that of bias detection and fairness in ML, since the problem of discrimination becomes even subtler when carried out by a “black-box” algorithm whose rationale is obscure or ambiguous.

Even if a lot of work has been done in this respect, still there is confusion on the interplay among different fairness notions and metrics to assess them. In this paper, we tried to highlight some important aspects about the relationships between fairness metrics, in particular with respect to the clash “individual vs. group” and “observational vs. causality-based” and we have argued that we are facing rather a landscape of intertwined possibilities than clear-cut distinctions.

Some have expressed critics and doubts about the possibility of capturing the complexity of notions such as equity and fairness with quantitative methods⁸³. Even if these doubts are reasonable, we believe that quantitative research in the domain of fairness notions can provide an important contribution to the more general issue of bias discrimination, at least in terms of understanding and comprehension of the problems and possibly to highlight shortcomings and subtleties – let alone to foster a boost in the attention of the scientific community on it. Thus, we strongly welcome more research both on the quantitative aspects, dealing with assessment metrics and algorithmic bias mitigation, and on the societal and legal aspects, and hopefully in the fruitful interplay between them.

Additional information

The views and opinions expressed are those of the authors and do not necessarily reflect the views of Intesa Sanpaolo, its affiliates or its employees.

References

1. Barocas, S., Hardt, M. & Narayanan, A. *Fairness and Machine Learning* (fairmlbook.org, 2019). <http://www.fairmlbook.org>.
2. Oneto, L. & Chiappa, S. Fairness in machine learning. In *Recent Trends in Learning From Data*, 155–196 (Springer, 2020).
3. Mitchell, S., Potash, E., Barocas, S., D’Amour, A. & Lum, K. Algorithmic fairness: Choices, assumptions, and definitions. *Annu. Rev. Stat. Its Appl.* **8**, 141–163 (2021).
4. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K. & Galstyan, A. A survey on bias and fairness in machine learning. *ACM Comput. Surv. (CSUR)* **54**, 1–35 (2021).
5. High-level expert group on artificial intelligence (HLEG, AI). Ethics guidelines for trustworthy AI. Tech. Rep., European Commission (2019).
6. Thiebes, S., Lins, S. & Sunyaev, A. Trustworthy artificial intelligence. *Electron. Mark.* **31**, 447–464 (2021).
7. Jobin, A., Ienca, M. & Vayena, E. The global landscape of AI ethics guidelines. *Nat. Mach. Intell.* **1**, 389–399 (2019).
8. Royal Society (Great Britain). *Machine learning: The power and promise of computers that learn by example* (2017).

9. Fjeld, J., Achten, N., Hilligoss, H., Nagy, A. & Srikumar, M. Principled artificial intelligence: Mapping consensus in ethical and rights-based approaches to principles for AI. *Berkman Klein Cent. Res. Publ.* (2020).
10. Campolo, A., Sanfilippo, M. R., Whittaker, M. & Crawford, K. AI Now 2017 report (2017).
11. Floridi, L. *et al.* AI4Peoplean ethical framework for a good AI society: opportunities, risks, principles, and recommendations. *Minds Mach.* **28**, 689–707 (2018).
12. Floridi, L. Establishing the rules for building trustworthy ai. *Nat. Mach. Intell.* **1**, 261–262 (2019).
13. Floridi, L. Translating principles into practices of digital ethics: Five risks of being unethical. In *Ethics, Governance, and Policies in Artificial Intelligence*, 81–90 (Springer, 2021).
14. The European Commission. Proposal for a Regulation of the European Parliament and of the Council laying down harmonised rules on Artificial Intelligence (Artificial Intelligence Act) and amending certain Union legislative acts (2021). <https://digital-strategy.ec.europa.eu/en/library/proposal-regulation-laying-down-harmonised-rules-artificial-intelligence>.
15. Pasquale, F. *The black box society* (Harvard University Press, 2015).
16. O’neil, C. *Weapons of math destruction: How big data increases inequality and threatens democracy* (Crown, 2016).
17. Eubanks, V. *Automating inequality: How high-tech tools profile, police, and punish the poor* (St. Martin’s Press, 2018).
18. Kantayya, S. Coded bias. Sundance Film Festival (2020).
19. Barocas, S. & Selbst, A. D. Big data’s disparate impact. *Calif. L. Rev.* **104**, 671 (2016).
20. Angwin, J., Larson, J., Mattu, S. & Kirchner, L. Machine bias: Theres software used across the country to predict future criminals, and its biased against blacks. *ProPublica* (2016).
21. Mulligan, D. K., Kroll, J. A., Kohli, N. & Wong, R. Y. This thing called fairness: disciplinary confusion realizing a value in technology. *Proc. ACM on Human-Computer Interact.* **3**, 1–36 (2019).
22. Narayanan, A. Translation tutorial: 21 fairness definitions and their politics. In *Proc. Conf. Fairness Accountability Transp., New York, USA*, vol. 2, 6–2 (2018).
23. Harty, J. F. Reasoning with moral conflicts. *Nous* **37**, 557–605 (2003).
24. Brink, D. O. Moral conflict and its structure. *The Philos. Rev.* **103**, 215–247 (1994).
25. Thomson, J. J. Killing, letting die, and the trolley problem. *The Monist* **59**, 204–217 (1976).
26. Verma, S. & Rubin, J. Fairness definitions explained. In *2018 IEEE/ACM International Workshop on Software Fairness (FairWare)*, 1–7 (IEEE, 2018).
27. Chouldechova, A. & Roth, A. The frontiers of fairness in machine learning. *arXiv preprint arXiv:1810.08810* (2018).
28. Chouldechova, A. & Roth, A. A snapshot of the frontiers of fairness in machine learning. *Commun. ACM* **63**, 82–89 (2020).
29. Makhoulf, K., Zhioua, S. & Palamidessi, C. On the applicability of machine learning fairness notions. *ACM SIGKDD Explor. Newsl.* **23**, 14–23 (2021).
30. Berk, R. *et al.* A convex framework for fair regression. *arXiv preprint arXiv:1706.02409* (2017).
31. Friedman, B. & Nissenbaum, H. Bias in computer systems. *ACM Transactions on Inf. Syst. (TOIS)* **14**, 330–347 (1996).
32. Crawford, K. The hidden biases in big data. *Harv. business review* **1** (2013).
33. Baeza-Yates, R. Bias on the web. *Commun. ACM* **61**, 54–61 (2018).
34. Hardt, M. How big data is unfair: Understanding unintended sources of unfairness in data driven decision making (Medium) (2014).
35. Dwork, C., Hardt, M., Pitassi, T., Reingold, O. & Zemel, R. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, 214–226 (2012).
36. Zafar, M. B., Valera, I., Gomez Rodriguez, M. & Gummadi, K. P. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In *Proceedings of the 26th international conference on world wide web*, 1171–1180 (2017).
37. Zemel, R., Wu, Y., Swersky, K., Pitassi, T. & Dwork, C. Learning fair representations. In *International Conference on Machine Learning*, 325–333 (2013).

38. Kamiran, F. & Calders, T. Data preprocessing techniques for classification without discrimination. *Knowl. Inf. Syst.* **33**, 1–33 (2012).
39. Berk, R. The role of race in forecasts of violent crime. *Race social problems* **1**, 231 (2009).
40. Johndrow, J. E., Lum, K. *et al.* An algorithm for removing sensitive information: application to race-independent recidivism prediction. *The Annals Appl. Stat.* **13**, 189–220 (2019).
41. Louizos, C., Swersky, K., Li, Y., Welling, M. & Zemel, R. The variational fair autoencoder. *arXiv preprint arXiv:1511.00830* (2015).
42. McNamara, D., Ong, C. S. & Williamson, R. C. Provably fair representations. *arXiv preprint arXiv:1710.04394* (2017).
43. Calmon, F., Wei, D., Vinzamuri, B., Ramamurthy, K. N. & Varshney, K. R. Optimized pre-processing for discrimination prevention. In *Advances in Neural Information Processing Systems*, 3992–4001 (2017).
44. Kusner, M. J., Loftus, J., Russell, C. & Silva, R. Counterfactual fairness. In *Advances in neural information processing systems*, 4066–4076 (2017).
45. Chiappa, S. Path-specific counterfactual fairness. In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 7801–7808 (2019).
46. Bickel, P. J., Hammel, E. A. & O’Connell, J. W. Sex bias in graduate admissions: Data from berkeley. *Science* **187**, 398–404 (1975).
47. Pearl, J., Glymour, M. & Jewell, N. P. *Causal inference in statistics: A primer* (John Wiley & Sons, 2016).
48. Pearl, J. & Mackenzie, D. *The book of why: the new science of cause and effect* (Basic Books, 2018).
49. Black, E., Yeom, S. & Fredrikson, M. Fliptest: fairness testing via optimal transport. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 111–121 (2020).
50. Jung, C. *et al.* Eliciting and enforcing subjective individual fairness. *arXiv preprint arXiv:1905.10660* (2019).
51. Ilvento, C. Metric learning for individual fairness. *arXiv preprint arXiv:1906.00250* (2019).
52. Rüz, T. Group fairness: Independence revisited. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 129–137 (2021).
53. Chouldechova, A. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big data* **5**, 153–163 (2017).
54. Kamiran, F. & Calders, T. Classifying without discriminating. In *2009 2nd International Conference on Computer, Control and Communication*, 1–6 (IEEE, 2009).
55. The U.S. Equal Employment Opportunity Commission (EEOC). Uniform Guidelines on employee selection procedures. Tech. Rep., EEOC (1979).
56. Feldman, M., Friedler, S. A., Moeller, J., Scheidegger, C. & Venkatasubramanian, S. Certifying and removing disparate impact. In *proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*, 259–268 (2015).
57. Friedler, S. A., Scheidegger, C. & Venkatasubramanian, S. On the (im) possibility of fairness. *arXiv preprint arXiv:1609.07236* (2016).
58. Hertweck, C., Heitz, C. & Loi, M. On the moral justification of statistical parity. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 747–757 (2021).
59. Jabbari, S., Joseph, M., Kearns, M., Morgenstern, J. & Roth, A. Fairness in reinforcement learning. In *International conference on machine learning*, 1617–1626 (PMLR, 2017).
60. Liu, L. T., Dean, S., Rolf, E., Simchowitz, M. & Hardt, M. Delayed impact of fair machine learning. In *International Conference on Machine Learning*, 3150–3158 (PMLR, 2018).
61. Hu, L. & Chen, Y. A short-term intervention for long-term fairness in the labor market. In *Proceedings of the 2018 World Wide Web Conference*, 1389–1398 (2018).
62. Kamiran, F., Žliobaitė, I. & Calders, T. Quantifying explainable discrimination and removing illegal discrimination in automated decision making. *Knowl. information systems* **35**, 613–644 (2013).
63. Hardt, M., Price, E. & Srebro, N. Equality of opportunity in supervised learning. In *Advances in neural information processing systems*, 3315–3323 (2016).

64. Jiang, R., Pacchiano, A., Stepleton, T., Jiang, H. & Chiappa, S. Wasserstein fair classification. In *Uncertainty in Artificial Intelligence*, 862–872 (PMLR, 2020).
65. Kleinberg, J., Mullainathan, S. & Raghavan, M. Inherent trade-offs in the fair determination of risk scores. *arXiv preprint arXiv:1609.05807* (2016).
66. Simoiu, C., Corbett-Davies, S. & Goel, S. The problem of infra-marginality in outcome tests for discrimination. *The Annals Appl. Stat.* **11**, 1193–1216 (2017).
67. Washington, A. L. How to argue with an algorithm: Lessons from the compas-propublica debate. *Colo. Tech. LJ* **17**, 131 (2018).
68. Berk, R., Heidari, H., Jabbari, S., Kearns, M. & Roth, A. Fairness in criminal justice risk assessments: The state of the art. *Sociol. Methods & Res.* **50**, 3–44 (2021).
69. Yang, F., Cisse, M. & Koyejo, O. O. Fairness with overlapping groups; a probabilistic perspective. *Adv. Neural Inf. Process. Syst.* **33** (2020).
70. Crenshaw, K. W. Mapping the margins. *The public nature private violence* 93–118 (1994).
71. Kearns, M., Neel, S., Roth, A. & Wu, Z. S. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In *International Conference on Machine Learning*, 2564–2572 (PMLR, 2018).
72. Kearns, M., Neel, S., Roth, A. & Wu, Z. S. An empirical study of rich subgroup fairness for machine learning. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 100–109 (2019).
73. Buolamwini, J. & Gebru, T. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, 77–91 (PMLR, 2018).
74. Dawid, A. P. Beware of the dag! In *Causality: objectives and assessment*, 59–86 (PMLR, 2010).
75. Dawid, A. P. Causal inference without counterfactuals. *J. Am. statistical Assoc.* **95**, 407–424 (2000).
76. Peters, J., Mooij, J. M., Janzing, D. & Schölkopf, B. Causal discovery with continuous additive noise models. *J. Mach. Learn. Res.* **15** (2014).
77. Glymour, C., Zhang, K. & Spirtes, P. Review of causal discovery methods based on graphical models. *Front. genetics* **10**, 524 (2019).
78. Guo, R., Cheng, L., Li, J., Hahn, P. R. & Liu, H. A survey of learning causality with data: Problems and methods. *ACM Comput. Surv. (CSUR)* **53**, 1–37 (2020).
79. Kilbertus, N. *et al.* Avoiding discrimination through causal reasoning. In *Advances in Neural Information Processing Systems*, 656–666 (2017).
80. Pearl, J. *Causality* (Cambridge university press, 2009).
81. Guidotti, R. *et al.* A survey of methods for explaining black box models. *ACM computing surveys (CSUR)* **51**, 1–42 (2018).
82. Burkart, N. & Huber, M. F. A survey on the explainability of supervised machine learning. *J. Artif. Intell. Res.* **70**, 245–317 (2021).
83. Green, B. & Hu, L. The myth in the methodology: Towards a recontextualization of fairness in machine learning. In *Proceedings of the machine learning: the debates workshop* (2018).