

Investigating Bias with a Synthetic Data Generator: Empirical Evidence and Philosophical Interpretation

Alessandro Castelnovo^{1,*}, Riccardo Crupi¹, Nicole Inverardi¹, Daniele Regoli¹ and Andrea Cosentini¹

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

Abstract

Machine learning applications are becoming increasingly pervasive in our society. Since these decision-making systems rely on data-driven learning, risk is that they will systematically spread the bias embedded in data. In this paper, we propose to analyze biases by introducing a framework for generating synthetic data with specific types of bias and their combinations. We delve into the nature of these biases discussing their relationship to moral and justice frameworks. Finally, we exploit our proposed synthetic data generator to perform experiments on different scenarios, with various bias combinations. We thus analyze the impact of biases on performance and fairness metrics both in non-mitigated and mitigated machine learning models.

Keywords

Synthetic Data, Bias, Fairness, Worldview, Machine Learning

1. Introduction

As society grows more digital, a greater amount of data becomes accessible for decision-making. In this context, machine learning techniques are increasingly being adopted by businesses, governments, and organizations in many important domains that affect people's lives every day. However, algorithms, like humans, are susceptible to biases that might lead to unfair outcomes [1].

Bias is not a recent problem; rather, it is ingrained in human society and, as a result, it is reflected in data [2]. The risk is that the adoption of machine learning algorithms could amplify or introduce biases that will reoccur in society in a perpetual cycle [3].

Several projects and initiatives have been launched in recent years aimed at bias mitigation and the development of fairness-aware machine learning models. Following to [2], we divide these works in three main categories:


- *Understanding bias.* Approaches that help to understand how bias is generated in society and manifest in data. This category contains studies of differences among biases, as well as their definition and formalization.
- *Accounting for bias.* Approaches discussing how to manage bias depending on the context, the regulation, the vision and the strategy on fairness [4, 5]. As discussed in [6], different

1st Workshop on Bias, Ethical AI, Explainability and the role of Logic and Logic Programming, BEWARE-22, co-located with AIXIA 2022, November 28 - December 2, 2022, University of Udine, Udine, Italy

* Corresponding author.



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

definitions of fairness and their implementations correspond to different axiomatic beliefs about the world (or *worldviews*), in general mutually incompatible.

- *Mitigating bias*. Technical approaches aimed to the development of machine learning models devoted to bias reduction with performance optimization. Depending on the stage of the machine learning pipeline where the bias is mitigated, these methods are typically divided into: *pre-processing* [7, 8], *in-processing* [9] and *post-processing* [10, 11].

One common approach to investigate the nature of bias is to conduct experiments on *ad hoc* scenarios through the generation of synthetic data [12, 13]. The benefits of this strategy include the possibility of examining circumstances not available with real-world data but that may occur, and—even when real-world data is available—to precisely control and understand the data generation mechanism. Moreover, it is indisputable that making data, and related challenges, accessible to the research community for analysis could be of help for the development of sound policy decisions and benefit society [13].

1.1. Contribution

With this work, we aim to contribute to the literature on fairness in machine learning in each of the three areas discussed above through the use of synthetic data.

In particular, we contribute to *understanding bias* by introducing a model framework for generating synthetic data with specific types of bias. Our formalisation of these various types of bias is based on the theoretical classifications present in the relevant literature, such as the surveys on bias in machine learning by Mehrabi et al. and Ntoutsis et al.

Against the background of the stream of literature about the relation between moral worldviews and biases, and in particular following [6, 14, 15], we analyze the worldviews related with each bias that our framework is able to generate, thus providing some insights in the discussion on *accounting for bias*.

Finally, about *mitigating bias*, we leverage our framework to generate twenty-five different scenarios characterized by the presence of various bias combinations. In each setting, we investigate the behavior and effects of traditional machine learning mitigation strategies [16].

An open source implementation of the proposed model framework is available at github.com/rcrupiISP/ISParity.

2. Related Works

2.1. Synthetic data

Synthetic data generation is a relevant practice for both businesses and the scientific community. As a result, the literature has given it a lot of attention.

Main directions behind the generation of synthetic data are: the *emulation* of certain key information in real dataset while preserving privacy [13, 17] and; the *generation* of different testing scenarios for evaluating phenomenon not covered by the available data [12]. Assefa et al. [17] presented basic use cases with specific examples in the financial domain like: internal data use restrictions, data sharing, tackling class imbalance, lack of historical data and training

advanced Machine Learning models. Moreover, the authors defined *Privacy preserving*, *Human readability* and *Compactness* as desirable properties for synthetic representation.

It is important to remark that synthetic data generation is not a possible technique for data anonymisation [18], but an alternative data sanitisation method to data masking for preserving privacy in published data [19]. In fact, synthetic data are typically randomly generated with constraints to protect sensitive private information and retain valid inferences with the attributes in the original data [13]. The provided synthetic data are generally classified into: *fully synthetic data*, *partially synthetic data* and *hybrid synthetic data* (see [20] for further details). We refer to [19] for a detailed overview of the techniques for generating synthetic data sets.

As introduced in the previous section, we are interested in producing fully synthetic data to replicate some common biases that could affect the data, and investigate fairness-related issues that arise from the development of machine learning models on them. In this regard, there are numerous works in the literature that generate synthetic datasets to simulate desired scenarios and, from these, testing discrimination-aware methods [21, 22, 23, 24]. For examples, Reddy et al. evaluate different fairness methods trained with deep neural networks on synthetic dataset. In these data were present different imbalanced and correlated data configuration to verify the limits of the current models and better understand in which setups they are subject to failure.

We contribute in this field of literature by introducing a modeling framework to generate synthetic data presenting specific forms of bias. Our formalisation of these different kinds of bias builds upon theoretical classifications present in the relevant literature, such as the works of Mehrabi et al. and Ntoutsi et al.. We leverage our proposed method to actually generate several datasets, each characterized by a specific combination of biases, and perform experiments on them to examine the effects of such biases on state-of-the-art mitigation approaches.

2.2. Bias and moral framework in decision-making

There is little consensus in the literature regarding bias classification and taxonomy. Moreover, the very notion of bias depends on profound ethical and philosophical considerations, which is likely one of the very causes for the lack of consensus. Different understandings of bias and fairness depend on the assumption of a belief system beforehand. Friedler et al. and Hertweck et al. talk about *worldviews*. In particular, [6] outlines two extreme cases, referred to as *What You See Is What You Get* (WYSIWYG) and *We are All Equal* (WAE).

Starting from the definition of three different metric spaces, these two perspectives differ because of the way they consider the relations in between. The first space is named *Construct Space* (CS) and represents all the unobservable realized characteristics of an individual, such as intelligence, skills, determination or commitment. The second space is the *Observable Space* (OS) and contains all the measurable properties that aim to quantify the unobservable features, think e.g. of IQ or aptitude tests. The last space is the *Decision Space* (DS), representing the set of choices made by the algorithm on the basis of the measurements available in OS. Note that shades of ambiguity are already detectable at this level, because the mappings between spaces are susceptible of distortions. Moreover, CS is by definition unobservable, thus we can only make *assumptions* on it.

According to WYSIWYG, CS and OS are essentially equal and any distortion between the two

is altogether irrelevant for the fairness of the decision resulting in DS. Contrarily, WAE doesn't make assumptions about the similarity of OS and CS, and moreover assumes that we are all equal in CS, i.e. that any difference between CS and OS is due to a biased observation process that results in an unfair mapping between CS and OS. With the distance between worldviews in mind, the notion of fairness inspired by [25] affirming that individuals that are close in CS shall also be close in DS (commonly known as individual fairness) appears diversified and differently achievable. If WYSIWYG is chosen, non-discrimination is guaranteed as soon as the mapping between OS and DS is fair, since $CS \approx OS$. On the other hand, according to WAE the mapping between CS and OS is distorted by some bias whenever a difference among individuals emerges (this difference is named Measurement Bias in [14]); therefore, to obtain a fair mapping between CS and DS those biases should be mitigated properly.

Building on [6], Hertweck et al. describe a more realistic and nuanced scenario by introducing the notion of *Potential Space* (PS): individuals belonging to different groups may indeed have different realized talents (i.e. they actually differ in CS), and these may be accurately measured by resumes (i.e. $CS \approx OS$), but, if we assume that these groups have the same *potential* talents (i.e. they are equal in PS), then the realized difference must be due to some form of unfair treatment of one group, that is referred to as *life bias*. Hertweck et al. call this view *We Are All Equal in Potential Space* (WAEPS). Actually, as argued in [14], we can effectively think of the WAE assumption as a family of assumptions, depending on the point time in which the equality is assumed to hold: the more we go back in time in assuming equality between individuals, the more the consideration of life circumstances becomes strong, and thus the less discrimination between individuals is considered legitimate.

These extreme worldviews amount on the one hand to accept the situation as it is observed (WYSIWYG), on the other to infer some form of unfairness whenever there is some observed disparity (WAE). To avoid such extreme scenarios, philosophical theories around *Equality of Opportunity* (EO) offer some suggestions and interpretive tools for approaching biases in different situations [26, 27, 28]. In this sense, western political philosophy and algorithmic fairness literature encounter themselves in the formulation of fairness around the concept of equal opportunities for all members of society. Heidari et al. list three different EO conceptions, going from more permissive to more stringent: *Libertarian EO*, by which individuals are held accountable for any characterizing feature, sensitive one included; *Formal EO*, by which individuals are not held accountable for differences in sensitive features only; *Substantive EO*, by which there is a set of individual characteristics which are due to *circumstances* and others that are a consequence of individual *effort*, and people should be held accountable only on the basis of the latter. The choice of which characteristics fall in the level of circumstances and which can be considered as individual effort is far from obvious. Depending on the EO framework that one is willing to embrace, observed disparities may be seen as "just" or "unjust" forms of bias.

In the following section, we shall describe the most common biases, explaining how they relate to these fundamental concepts.

3. Fundamental types of bias

Considering that the assumptions about the worldview and the EO framework affect the conception as well as the assessment of biases, in what follows we focus on what we consider the fundamental building blocks of most types of biases, namely: *Historical bias*, *Measurement bias*, *Representation bias*, *Omitted variable bias*.

Historical bias—sometimes referred to as *social bias*, *life bias*, or *structural bias* [3, 2, 14]—occurs whenever a variable of the dataset relevant to some specific goal or task is dependent on some sensitive characteristic of individuals, *but in principle it should not*. An example of this bias is the different average income among men and women, which is due to long-lasting social pressures in a man-centered society, and does not reflect intrinsic differences among sexes. Following [3], we can talk of a form of bias going *from users to data*: this type of bias affects directly the actual phenomenon generating the data. A similar situation may arise when the dependence of sensitive individual characteristics is present with respect to the variable that we are trying to estimate or predict. For instance, there are cases in which the target of model estimation is itself prone to some form of bias, e.g. because it is the outcome of some human decision. Think e.g. to try to build a data-driven decision process to decide whether to grant or not a loan on the basis of past loan officers' decisions, and not of actual repayments. Note that the actual presence of historical bias is conditioned by the previous assumption of the WAE worldview. Indeed, arguing that in principle there should be no dependence on some sensitive features makes sense only if a moral belief of substantial equity is required to begin with. Otherwise, according to WYSIWYG, CS is fairly reported in OS and therefore structural differences between individuals are legitimate sources of inequality. Moreover, accepting the Libertarian EO or the Substantive EO frameworks would involve the legitimate use of some sensitive features, respectively because they are a property of the self and because we should be aware of their influence on the values of non-sensitive features. Ultimately, the presence of historical bias depends on the assumption of WAE at the initial time of life. As argued in [14], interpreting bias as historical means conceiving equality at the level of PS, that describes the innate/native potential of each individual.

Measurement bias occurs when a proxy of some variable relevant to a specific goal or target is employed, and that proxy happens to be dependent on some sensitive characteristics. For instance, one may argue that IQ is not a “fair” approximation of actual “intelligence”, and it might systematically favour/disfavour specific groups of individuals. Statistically speaking, this type of bias is not very different from historical bias—since it results in employing a variable correlated with sensitive attributes—but the underlying mechanism is nevertheless different, and in this case the bias needs not to be present in the phenomenon itself, but rather it may be a consequence of the choice of data to be employed. In other words, this is an example of bias *from data to algorithm* in the taxonomy of [3], i.e. a bias due to data availability, choice and collection. Incidentally, notice that this form of bias—using a biased proxy of a relevant variable—might as well happen with the target variable. In this situation, it is the quantity that we need to estimate/predict that is somehow “flawed”. The fact that measurement bias depends also on a choice component, which is to say the choice of the dataset, can extend its occurrence in the WYSIWYG worldview as well. Indeed, the choice of “what to measure” determines and modifies what is made observable. The eventuality of awareness of biased measurements would

probably require mitigation also in WYSIWYG. Alternatively, according to WAE, measurement bias may lie in the mapping between construct and observed space, which is to say between a “real” ability of an individual and an observable quantity that tries to measure it.

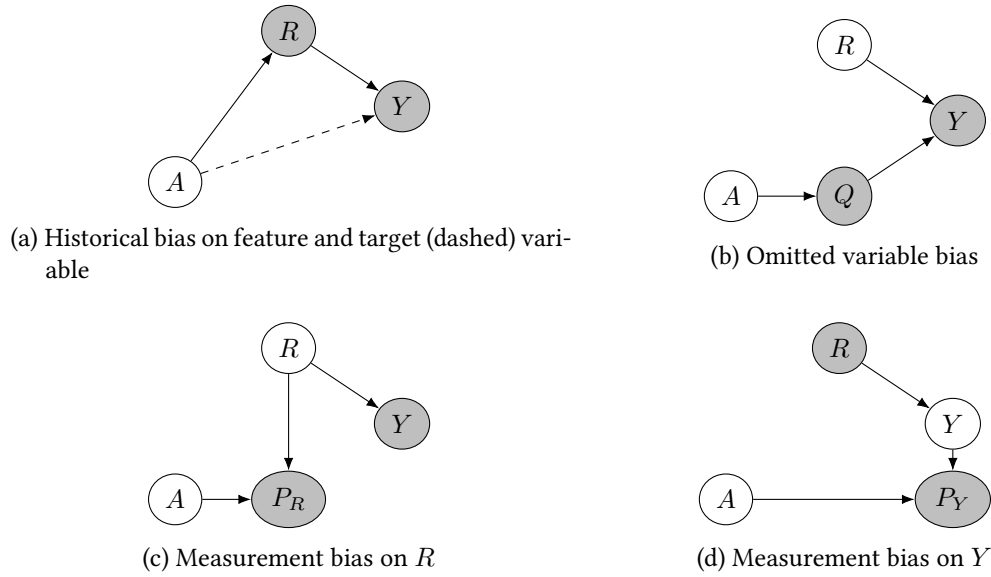


Figure 1: Atomic biases. Grey-filled circles represent variables employed by the model \hat{f} .

Representation bias occurs when, for some reasons, data are not representative of the world population. One subgroup of individuals, e.g. identified by some sensitive characteristic such as ethnicity, age, etc. may be heavily underrepresented. This under-representation may occur in different ways. It may be at random, i.e. the subgroup is less numerous than it should be, but without any particular skewness in the other characteristics: in this case this single mechanism is not sufficient to create disparities, but it may exacerbate existing ones. Alternatively, the under-represented subgroup might contain individuals with disproportionate characteristics with respect to their corresponding world population, e.g. only low-income individuals, or only low-education individuals. In this last case, representation bias may be enough to create disparities in decision making processes based on those data. The mechanism underlying the representation disparities should be analyzed on the basis of the assumed worldview and/or chosen EO framework: e.g. if the data has an under-represented ethnic minority one should investigate *why* it is so. If the target population is itself different from world population (i.e. it is not merely a poor data collection), then one should consider the reasons by which this ethnic minority is under-represented in the target population, e.g. in the Substantive EO framework one should understand if these reasons have to be regarded as *circumstance* or as consequence of individual *effort*. In the latter case, the representation disparities are not to be considered as “unfair” *per se*.

Representation bias is strictly connected to *sampling bias*, in that it embodies problems arising during data collection, e.g. by collecting disproportionately less observations from one subgroup, possibly skewed with respect to some characteristics. Like measurement bias, this is a form of

bias going *from data to algorithm*.

Omitted variable bias may occur when the collected dataset omits a variable relevant to some specific goal or task. In this case, if the variables that are present in the dataset have some dependence on sensitive characteristics of individuals, then a machine learning model trained on such dataset will learn such dependencies, thus producing outcomes with spurious dependence on sensitive attributes. Assuming the Formal EO framework, sensitive features are omitted by default. While this may appear fairer because the decision is made solely on the basis of the relevant attributes, on the other hand it becomes arduous to mitigate on structural biases that affect achievements. Depending on worldview assumptions or on the chosen EO framework, the mechanism through which the residual variables happen to depend on sensitive individual characteristics should be as well analyzed to understand/decide whether this dependence is legitimate or if they are themselves a consequence of some bias at work.

The above list of biases should be seen as the set of the most important mechanisms through which “unfair” disparities happen to result in data-driven decision making. Notice, however, that in terms of consequences on the data, it may well be that different types of bias result in very similar effects. E.g. representation bias may create in the dataset spurious correlations among sensitive characteristics of individuals and other characteristics relevant to the problem at hand, a situation very similar to the correlations present as a consequence of historical bias. This reminds us that in reality we are not aware of the type of bias (or biases) affecting the data and that their interpretation depends on former assumptions.

4. Dataset Generation

We propose a simple modeling framework able to simulate the bias-generating mechanisms described in section 3.

The rationale behind the model is that of being at the same time sufficiently flexible to accommodate all the main forms of bias generation *while* maintaining a structure as simple and intuitive as possible to facilitate *human readability* and ensure *compactness* avoiding unnecessary complexities that might hide the relevant patterns.

As noted in section 3, following [3] we can distinguish between biases from *users to data* and from *data to algorithm*. Namely, between biases that impact the phenomenon to be studied and thus the dataset, and biases that impact directly the dataset but not the phenomenon itself. Formally, we model the relevant quantities describing a phenomenon as random variables, in particular we label Y the *target* variable, namely the quantity to be estimated or predicted on the basis of other *feature* variables, that we collectively call X . As usual, we assume that the underlying phenomenon is described by the formula

$$Y = f(X) + \epsilon, \quad (1)$$

where f represents the actual relationship between features and target variables, modulated by some idiosyncratic noise ϵ .

A data-driven decision maker infers from a (training) set of samples $\{(\tilde{x}_i, y_i)\}_{i=1, \dots, N}$, an estimate for f that we label \hat{f} , thus producing its best estimate for Y , namely

$$\hat{Y} = \hat{f}(\tilde{X}). \quad (2)$$

The use of \tilde{X} rather than X describes the fact that the set of variables employed to make inferences about a phenomenon may not coincide with the actual variables that play a role in that phenomenon. This is precisely what happens in some forms of biases such as measurement bias or omitted variable bias.

Notice that *users to-data* types of bias impact directly equation (1), while *data-to-algorithm* biases have a role at the level of equation (2).

It is possible to make a schematic representation of building blocks of biases as discussed in section 3 via Directed Acyclic Graphs (DAGs), representing causal impacts among variables (see e.g. [29, 30, 31]). In general, in order to provide an intuitive grasp on interesting mechanisms and patterns, we shall make reference to the following situation:

- we label with R variables representing *resources* of individuals —be them economic resources, or personal talents and skills— which are relevant for the problem, i.e. they directly impact the target Y ;
- we label with A variables indicating sensitive attributes, such as ethnicity, gender, etc.;
- we label with P_R proxy variables that we have access to instead of the original variable R . We may think of R as soft skills, or talent, or intelligence, i.e. something that we cannot measure or sometimes even quantify directly;
- we label with Q additional variables, that may or may not be relevant for the problem (i.e. impacting Y) and that may or may not be impacted either by R or A , e.g. the neighborhood one lives in.

With these examples in mind, we can e.g. think of Y —i.e. the target of our decision making process— as the repayment of a debt or the work performance.

In particular, figure 1 shows 3 minimal graph representations of historical, omitted variable and measurement biases. Historical bias occurs when the relevant variable R is somehow impacted by sensitive feature A . Omitted variable bias occurs when, for some reasons, we omit the relevant variable R from our dataset and we employ another variable which happens to be impacted by A . Measurement bias occurs when the relevant variable R is, in general, free of bias, but we cannot access it, thus we employ a proxy P_R which is impacted by sensitive characteristic A . The measurement bias could occur also on the target variable Y , when we can access only on a proxy P_Y of the phenomenon that we want to predict.

The following system of equations formalizes the relationships between variables that are used to simulate specific form of biases:

$$A = B_A, \quad B_A \sim \mathcal{Ber}(p_A); \quad (3a)$$

$$R = -\beta_h^R A + N_R, \quad N_R \sim \mathcal{N}(\mu_R, \sigma_R^2); \quad (3b)$$

$$Q = B_Q, \quad B_Q | (R, A) \sim \mathcal{Bin}(K, p_Q(R, A)), \quad (3c)$$

$$p_Q(R, A) = \text{sigmoid} \left(-(\alpha_{RQ} R - \beta_h^Q A) \right);$$

$$S = \alpha_R R - \alpha_Q Q - \beta_h^Y A + N_S, \quad N_S \sim \mathcal{N}(0, \sigma_S^2); \quad (3d)$$

$$Y = \mathbf{1}_{\{S > \bar{s}\}}. \quad (3e)$$

$$(3f)$$

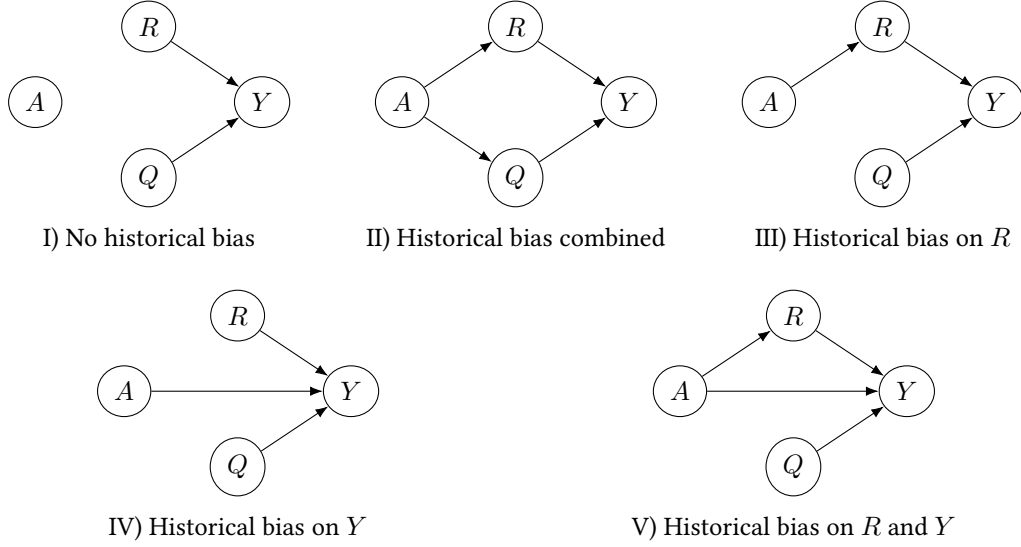


Figure 2: The 5 considered scenarios with *users-to-data* biases.

When simulating measurement bias, either on resources R or on target Y , we are going to use the following *proxies* as noisy (and biased) substitutes of the actual variables:

$$P_R = R - \beta_m^R A + N_{P_R}; \quad N_{P_R} \sim \mathcal{N}(0, \sigma_{P_R}^2); \quad (4a)$$

$$P_S = S - \beta_m^Y A + N_{P_S}; \quad N_{P_S} \sim \mathcal{N}(0, \sigma_{P_S}^2); \quad (4b)$$

$$P_Y = \mathbf{1}_{\{P_S > \bar{s}\}} \quad (4c)$$

By varying the values of the parameters, we are able to generate different aspect of biases as follows:

- β_h^j determines the presence and amplitude of the *historical bias* on the variable $j \in \{R, Q, Y\}$;
- β_m^j , when the proxy P_j is used instead of the original variable j , governs the intensity of *measurement bias* on $j \in \{R, Y\}$;
- α_R, α_Q control the linear impact on (S and thus) Y of R and Q , respectively; α_{RQ} represents the intensity of the dependence of Q on R .

Additionally, in order to account for *representation bias*, we undersample the $A = 1$ group. The amount of undersampling is governed by the parameter p_u defined as the proportion of the under-represented group $A = 1$ with respect to the majority group $A = 0$. We draw the undersampling *conditioned* on R by selecting the $A = 1$ individuals with lower R . Finally, simulating *omission bias* is as simple as dropping the variable R from the set of features the model uses to estimate Y .

5. Experiments

We employ our modeling framework to generate a total of twenty-five different datasets by combining a set of five *users-to-data* biases with a set of four *data-to-algorithm* biases, plus a configuration without additional biases. The aim of the experiments is to investigate the effects of these bias combinations on machine learning algorithms.

As clarified in section 3, *users-to-data* biases impact the phenomenon to be studied and thus the dataset. The five scenarios that we have chosen to explore, represented in figure 2, are defined as follows:

I) **No historical bias:**

$$Y = f(R, Q) + \epsilon, \quad R, Q \perp\!\!\!\perp A. \quad (5)$$

II) **Historical bias with effect compensation:**

$$Y = f(R, Q) + \epsilon, \quad R = R(A), Q = Q(A), Y \perp\!\!\!\perp A. \quad (6)$$

III) **Historical bias on R :**

$$Y = f(R, Q) + \epsilon, \quad R = R(A), Q \perp\!\!\!\perp A. \quad (7)$$

IV) **Historical bias on Y :**

$$Y = f(R, Q, A) + \epsilon, \quad R, Q \perp\!\!\!\perp A. \quad (8)$$

V) **Historical bias on R and Y :**

$$Y = f(R, Q, A) + \epsilon, \quad R = R(A), Q \perp\!\!\!\perp A. \quad (9)$$

We combine each scenario with the following *data-to-algorithm* biases, having an effect on the variables that we assume we can access in order to train the models:

- *no additional bias*;
- *measurement bias* on R ;
- *omission bias* on R ;
- *representation bias* on $A = 1$;
- *measurement bias* on Y .

We thus generate twenty-five synthetic datasets, each of which characterized by a specific choice of the set of parameters discussed in section 4. Details on the values of the parameters and the complete python code to reproduce the experiments is publicly available at github.com/rcrupiISP/ISParity.

On each generated dataset, we train and test the following three algorithms: (i) a *Random Forest* (RF) [32] that aims to maximize performance by utilizing all available variables; (ii) *Blinded Random Forest* (BRF): a random forest that does not use the sensitive variable A in order to avoid creating decision-making patterns based on it; (iii) *Equalized Random Forest* (ERF): the

Table 1

Experiments results. Results regarding 25 experiments: 5 *users-to-data* scenarios with varying types of biases (rows) combined with 5 *data-to-algorithm* types of bias. Metrics employed are Accuracy (Acc), Demographic Parity and Accuracy difference between $A = 0$ and $A = 1$ groups (ΔDP and ΔAcc). ΔDP for the “true” target variable Y is also provided for each scenario (values in round brackets). RF is a random forest trained with all the available information, BRF is a RF blinded with respect to A , ERF is a RF with a post-processing fairness mitigation to impose Demographic Parity. Values are expressed as percentage points.

users-to-data bias		data-to-algorithm bias														
		no additional bias			+ meas. bias on R (P_R substitutes R)			+ omission bias (R is omitted)			+ representation bias ($A = 1$ randomly undersampled)			+ meas. bias on Y (P_Y substitutes Y)		
Scenario	metrics	RF	BRF	ERF	RF	BRF	ERF	RF	BRF	ERF	RF	BRF	ERF	RF	BRF	ERF
I) No historical bias	Acc	86.4	86.4	86.4	81.2	81.1	81.2	62.8	62.8	62.8	86.3	86.3	86.3	83.3	84.5	84.3
	ΔDP (0)	0.1	0.1	0.0	0.8	7.4	0.5	1.2	1.2	1.0	0.7	0.6	0.4	11.2	0.1	0.0
	ΔAcc	0.1	0.1	0.1	0.5	0.7	0.5	0.4	0.4	0.4	2.0	1.8	2.0	4.0	0.1	0.0
II) Historical bias compensation	Acc	86.3	86.4	86.4	81.1	81.1	81.1	62.7	61.6	62.4	86.3	86.3	86.4	83.7	84.2	84.6
	ΔDP (0.9)	0.1	0.8	0.3	2.2	5.6	0.9	9.8	27.1	0.9	1.6	0.7	0.6	11.5	3.9	0.4
	ΔAcc	0.0	0.1	0.0	0.4	0.7	0.3	0.7	1.6	0.0	1.4	1.5	1.2	3.4	0.7	0.5
III) Historical bias on R	Acc	86.7	86.6	85.5	81.5	81.4	80.6	64.5	62.6	62.6	86.4	86.3	86.2	84.0	84.9	83.7
	ΔDP (5)	12.6	12.9	0.0	11.8	19.1	0.7	36.7	1.2	1.0	11.2	10.3	0.0	25.1	13.1	0.1
	ΔAcc	2.2	2.1	2.3	2.5	2.2	2.3	1.2	2.6	2.6	0.9	1.0	1.3	1.7	1.5	4.9
IV) Historical bias on Y	Acc	86.5	85.5	85.5	81.5	81.3	80.6	64.5	62.6	62.6	86.4	86.1	86.2	84.3	83.5	83.9
	ΔDP (12)	11.2	0.1	0.0	10.9	7.4	0.6	36.7	1.2	1.0	10.7	0.5	0.0	23.1	0.2	0.1
	ΔAcc	1.7	0.0	2.4	2.7	2.9	2.2	1.3	2.6	2.6	1.0	1.6	1.3	1.1	5.7	4.7
V) Historical bias on R and Y	Acc	86.9	86.0	83.1	82.0	81.6	78.7	66.4	62.1	62.0	86.5	86.3	85.6	84.6	84.3	81.2
	ΔDP (7)	25.1	13.1	0.1	27.0	18.8	0.6	36.7	1.2	1.0	19.1	10.3	1.3	37.4	13.3	0.1
	ΔAcc	4.0	3.8	3.7	5.1	5.0	4.6	4.7	4.1	4.1	4.2	1.4	6.5	0.1	6.6	11.3

same model in (i) with a post-processing fairness mitigation block [16] to impose the same rate of $\hat{Y} = 1$ between the different classes of A (i.e. imposing Demographic Parity [33]).

Fairness of outcomes is evaluated through *Demographic Parity* difference (ΔDP) [33], i.e. the difference between the rate of $\hat{Y} = 1$ among the two sensitive groups $A = 0$ and $A = 1$, namely $\Delta DP = P(\hat{Y} = 1 | A = 0) - P(\hat{Y} = 1 | A = 1)$. ΔDP can also be computed for the “true” target variable Y , thus capturing the dependence of the target variable on the sensitive attribute A . Performance is evaluated by overall *Accuracy* (Acc). As an additional metric both of fairness and of performance, we compute the difference in accuracy between $A = 0$ and $A = 1$ groups (ΔAcc).

6. Discussion

Table 1 summarizes the results obtained by the twenty-five experiments on generated datasets, on the basis of which we draw the following non-exhaustive list of insights:

On the effect of equalization: as expected, the adoption of a model to enforce a group equalization of the outcomes results in lower performance in the majority of scenarios considered. This gap is proportional to the level of disparity present in the data. On the contrary, when there is representation bias, the effects on performance are less visible, since enforcing equalization impacts fewer instances. This suggests that observing significant reductions in ΔDP with negligible performance deterioration is not *per se* sufficient to conclude that a mitigation strategy has no effects on performance —as expressed by numerous studies in the field of fairness— since the protected class happens often to be a minority. Observing the difference in accuracy (ΔAcc), on the other hand, is a much more robust indicator because it is unaffected by the different numerosity of the protected subgroups.

On the effect of blindness: excluding sensitive attributes when training machine learning models —as suggested by the Formal EO framework— has a negative impact on performance only when Y is directly influenced by A . When the factors that determine Y contain historical bias, this technique has no benefit in terms of demographic parity. It’s worth noting that in the presence of measurement bias in scenarios I, II, and III, the effects on ΔDP are quite worse than when A is included in the model. This occurs because the machine learning model actually needs A to remove the excess of dependence induced by the inclusion of the proxy.

On the performance-bias trade-off in moral frameworks: one crucial aspect to keep in mind when developing a decision system is that what is and what is not a *bias* depends on the moral framework and worldview chosen in the first place. In this regard, the employment of a modeling strategy over another (e.g. using BRF instead of ERF) is in response to the chosen worldview. Thus, *given a specific moral view*, there are circumstances where it is justifiable to prefer performance over group equality and vice versa. Notice that the choice of a framework needs often be complemented with additional (crucial) decisions, e.g. choosing Substantive EO *per se* is not enough to understand whether some phenomenon underlying data is or not a bias, since one needs to agree on which characteristics are to be considered as “individual effort” (as opposed to “circumstance”), and thus legitimately employed as rationale of decisions.

On the interpretation of results: in this work, we simulate the presence of one or more biases and observe the corresponding effects on data. In real-world scenarios, data are typically observed without knowing the exact bias-generation mechanism underlying them. It is the assumption on the data generation mechanism that shapes the interpretation of biases present in the context under consideration.

This concept becomes more evident when the target variable in the data serves as proxy for the true target. In our experiments with measurement bias on Y , we trained the models with the proxy and evaluated the results with the real variable. This approach, even if unlikely in real-world situations, demonstrates that reducing bias can also lead to improved performances. These set of experiments shows that under a specific worldview, mitigation strategies that result in lower performance on the target variable in the observed space, may lead to higher performance on the “true” target.

On the absence of bias: when operating in an ideal bias-free environment, i.e. when any

possible relevant characteristic is independent of sensitive features, all machine learning strategies provide the same performance and fairness results. Furthermore, in this ideal situation, all worldviews would agree, even if they are mutually incompatible in the general case. Regardless of how trivial this conclusion may be, it is essential to mention it in order to promote any legal, social, or strategic initiative with the goal of eliminating society biases.

7. Conclusion

In this work we have investigated how different types of bias impact fairness and performance metrics of machine learning models by introducing of a unified framework for generating synthetic data. Regardless of whether the test scenarios used to draw conclusions are synthetic, we have discussed moral frameworks and worldviews that can be applied in real-world situations.

We release the presented model framework as open-source at github.com/rcrupiISP/ISParity in order to encourage further research into the challenge of bias in data. For example, one crucial aspect to be further analyzed is the sensitiveness of results with varying magnitude of parameters; another is to control the effect of random fluctuations by averaging results of many simulations for each choice of the parameters set.

References

- [1] J. Angwin, J. Larson, S. Mattu, L. Kirchner, Machine bias: There's software used across the country to predict future criminals, and it's biased against blacks, ProPublica (2016).
- [2] E. Ntoutsi, P. Fafalios, U. Gadiraju, V. Iosifidis, W. Nejdl, M.-E. Vidal, S. Ruggieri, F. Turini, S. Papadopoulos, E. Krasanakis, et al., Bias in data-driven artificial intelligence systems—an introductory survey, *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 10 (2020) e1356.
- [3] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, A. Galstyan, A survey on bias and fairness in machine learning, *ACM Computing Surveys (CSUR)* 54 (2021) 1–35.
- [4] A. Castelnovo, L. Malandri, F. Mercorio, M. Mezzanzanica, A. Cosentini, Towards fairness through time, in: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, Springer, 2021, pp. 647–663.
- [5] A. Castelnovo, R. Crupi, G. Del Gamba, G. Greco, A. Naseer, D. Regoli, B. S. M. Gonzalez, Befair: Addressing fairness in the banking sector, in: *2020 IEEE International Conference on Big Data (Big Data)*, IEEE, 2020, pp. 3652–3661.
- [6] S. A. Friedler, C. Scheidegger, S. Venkatasubramanian, The (im) possibility of fairness: Different value systems require different mechanisms for fair decision making, *Communications of the ACM* 64 (2021) 136–143.
- [7] M. A. Madaio, L. Stark, J. Wortman Vaughan, H. Wallach, Co-designing checklists to understand organizational challenges and opportunities around fairness in ai, in: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 2020, pp. 1–14.
- [8] F. Kamiran, T. Calders, Classifying without discriminating, in: *2009 2nd International Conference on Computer, Control and Communication*, IEEE, 2009, pp. 1–6.

- [9] A. Agarwal, A. Beygelzimer, M. Dudík, J. Langford, H. Wallach, A reductions approach to fair classification, in: International Conference on Machine Learning, PMLR, 2018, pp. 60–69.
- [10] M. Hardt, E. Price, N. Srebro, Equality of opportunity in supervised learning, in: Advances in neural information processing systems, 2016, pp. 3315–3323.
- [11] P. K. Lohia, K. N. Ramamurthy, M. Bhide, D. Saha, K. R. Varshney, R. Puri, Bias mitigation post-processing for individual and group fairness, in: Iccasp 2019-2019 ieee international conference on acoustics, speech and signal processing (icassp), IEEE, 2019, pp. 2847–2851.
- [12] T. Le Quy, A. Roy, V. Iosifidis, W. Zhang, E. Ntoutsi, A survey on datasets for fairness-aware machine learning, Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery (2022) e1452.
- [13] T. E. Raghunathan, Synthetic data, Annual Review of Statistics and Its Application 8 (2021) 129–140.
- [14] C. Hertweck, C. Heitz, M. Loi, On the moral justification of statistical parity, in: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, 2021, pp. 747–757.
- [15] H. Heidari, M. Loi, K. P. Gummadi, A. Krause, A moral framework for understanding fair ml through economic models of equality of opportunity, in: Proceedings of the conference on fairness, accountability, and transparency, 2019, pp. 181–190.
- [16] S. Bird, M. Dudík, R. Edgar, B. Horn, R. Lutz, V. Milan, M. Sameki, H. Wallach, K. Walker, Fairlearn: A toolkit for assessing and improving fairness in AI, Technical Report MSR-TR-2020-32, Microsoft, 2020. URL: <https://www.microsoft.com/en-us/research/publication/fairlearn-a-toolkit-for-assessing-and-improving-fairness-in-ai/>.
- [17] S. A. Assefa, D. Dervovic, M. Mahfouz, R. E. Tillman, P. Reddy, M. Veloso, Generating synthetic data in finance: opportunities, challenges and pitfalls, in: Proceedings of the First ACM International Conference on AI in Finance, 2020, pp. 1–8.
- [18] A. Majeed, S. Lee, Anonymization techniques for privacy preserving data publishing: A comprehensive survey, IEEE access 9 (2020) 8512–8545.
- [19] H. Surendra, H. Mohan, A review of synthetic data generation methods for privacy preserving data publishing, International Journal of Scientific & Technology Research 6 (2017) 95–101.
- [20] C. C. Aggarwal, S. Y. Philip, Privacy-preserving data mining: models and algorithms, Springer Science & Business Media, 2008.
- [21] A. D’Amour, H. Srinivasan, J. Atwood, P. Baljekar, D. Sculley, Y. Halpern, Fairness is not static: deeper understanding of long term fairness via simulation studies, in: Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, 2020, pp. 525–534.
- [22] W.-Y. Loh, L. Cao, P. Zhou, Subgroup identification for precision medicine: A comparative review of 13 methods, Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery 9 (2019) e1326.
- [23] C. Reddy, D. Sharma, S. Mehri, A. Romero-Soriano, S. Shabanian, S. Honari, Benchmarking bias mitigation algorithms in representation learning through fairness metrics, in: Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1), 2021.
- [24] A. Castelnovo, R. Crupi, G. Greco, D. Regoli, I. G. Penco, A. C. Cosentini, A clarification of

- the nuances in the fairness metrics landscape, *Scientific Reports* 12 (2022) 1–21.
- [25] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, R. Zemel, Fairness through awareness, in: *Proceedings of the 3rd innovations in theoretical computer science conference*, 2012, pp. 214–226.
 - [26] M. Fleurbaey, *Fairness, responsibility, and welfare*, OUP Oxford, 2008.
 - [27] J. E. Roemer, A. Trannoy, Equality of opportunity, in: *Handbook of income distribution*, volume 2, Elsevier, 2015, pp. 217–300.
 - [28] J. Rawls, *A theory of justice*, in: *Ethics*, Routledge, 2004, pp. 229–234.
 - [29] J. Pearl, *Causality*, Cambridge university press, 2009.
 - [30] J. Pearl, D. Mackenzie, *The book of why: the new science of cause and effect*, Basic Books, 2018.
 - [31] J. Peters, J. M. Mooij, D. Janzing, B. Schölkopf, Causal discovery with continuous additive noise models, *Journal of Machine Learning Research* 15 (2014).
 - [32] G. Biau, E. Scornet, A random forest guided tour, *Test* 25 (2016) 197–227.
 - [33] S. Barocas, M. Hardt, A. Narayanan, *Fairness and Machine Learning*, fairmlbook.org, 2019. <http://www.fairmlbook.org>.