

# Fair overlap number of balls (Fair-ONB): an undersampling method for bias reduction based on data morphology

José Daniel Pascual-Triana, Alberto Fernández, Paulo Novais & Francisco Herrera

**To cite this article:** José Daniel Pascual-Triana, Alberto Fernández, Paulo Novais & Francisco Herrera (2025) Fair overlap number of balls (Fair-ONB): an undersampling method for bias reduction based on data morphology, *Statistics*, 59:4, 970-988, DOI: [10.1080/02331888.2025.2476029](https://doi.org/10.1080/02331888.2025.2476029)

**To link to this article:** <https://doi.org/10.1080/02331888.2025.2476029>



Published online: 12 Mar 2025.



Submit your article to this journal [↗](#)



Article views: 47



View related articles [↗](#)



View Crossmark data [↗](#)



# Fair overlap number of balls (Fair-ONB): an undersampling method for bias reduction based on data morphology

José Daniel Pascual-Triana <sup>a</sup>, Alberto Fernández<sup>a</sup>, Paulo Novais<sup>b</sup> and Francisco Herrera<sup>c</sup>

<sup>a</sup>Andalusian Institute of Data Science and Computational Intelligence (DASCI), University of Granada, Granada, Spain; <sup>b</sup>ALGORITMI CENTRE/LASI, University of Minho, Braga, Portugal; <sup>c</sup>ADIA Lab, Abu Dhabi, United Arab Emirates

## ABSTRACT

To attain Trustworthy Artificial Intelligence, Fairness must be ensured in classification problems when protected (sensitive) features are present. Since classifiers might reflect the biases in training data, impacting humans and contravening current regulations, data quality is key. Preprocessing the dataset can help improve data quality; particularly, undersampling can promote simultaneous learning of protected features and classes. Sampling in class overlap areas, close to the decision boundary, should magnify the impact on the classifier. This work proposes Fair Overlap Number of Balls (Fair-ONB), an undersampling method that employs the morphology of data groups (formed by combining classes and protected feature values) to guide the undersampling in overlap areas. It utilizes balls to cover groups and harnesses some of their characteristics, such as their radii, numbers of covered instances and densities, to select where undersampling would be most effective for bias reduction. The experimental framework shows that the Fair-ONB method increases model Fairness with a low effect on predictive performance.

## ARTICLE HISTORY

Received 29 October 2024  
Accepted 1 March 2025

## KEYWORDS

Trustworthy artificial intelligence; machine learning; sampling; fairness; data morphology

## 1. Introduction

The vast amount of data generated in this digital era [1], ranging from simple sensor measurements to biometric data, makes harnessing relevant information impossible unless some processes are automated. Machine Learning techniques learn the inner patterns in data, making it possible for computers to predict the behavior of unseen data [2].

To learn these patterns and, in particular, the classes in supervised classification problems, the characteristics of each dataset are crucial [3]. Models learned from poor-quality data do not reach high predictive performance and might draw unrealistic conclusions [4]. In particular, class overlap hinders the model's estimation of class boundaries, class imbalance can reduce the importance of some classes in the model (which complicates

**CONTACT** José Daniel Pascual-Triana  [jd Pascual@gmail.com](mailto:jd Pascual@gmail.com); Alberto Fernández  [alfh@ugr.es](mailto:alfh@ugr.es)  Andalusian Institute of Data Science and Computational Intelligence (DASCI), University of Granada, Granada 18071, Spain

This article has been corrected with minor changes. These changes do not impact the academic content of the article.

an equitable predictive performance), and the presence of noisy samples can distort the predictions.

When datasets include personal data or sensitive features that might give rise to discrimination, low data quality or mistreatment of said features can cause models to include discriminatory bias [5]. This critical problem has led to auditability requirements and the emergence of Artificial Intelligence regulations, such as the recently approved European Union Regulation 2024/1689<sup>1</sup>, an evolution of the European Union Artificial Intelligence Act<sup>2</sup>.

Some sensitive features, usually referred to as protected features (for example, race or gender), are prone to causing discrimination amongst groups or classes and cannot be included in the model's decision function. However, even if these features are not directly involved in a model, the latter can still include biases depending on how instances with different values on those protected features are classified (sometimes due to the indirect impact of other highly correlated features to the protected one that the model does use) [6]. Data and model Fairness are achieved when these possible biases amongst groups are minimized or eliminated.

While Artificial Intelligence often focuses on optimizing prediction model parameters, another approach focuses on data quality, as models learned on high-quality data tend to achieve high performance [3]. A widespread approach to reduce model bias is to modify the dataset via sampling [7]. Sampling data of specific groups or classes, either by adding instances of those with fewer data (oversampling) or selecting for usage only some from the over-represented ones (undersampling), can foster an equitable treatment of all groups or classes by the model. When protected features are present, this would mean balancing the model's performance for instances of various values of the protected features. Specifically, this means balancing the groups formed by the combinations of the class values and protected feature values as, when taken into account separately, not only might bias reduction be harder, but counterproductive situations might arise (for example, if more *gender = female* samples were added to balance *gender* but they were added in negative class areas, discrimination against *gender = female* would increase).

Sampling strategies in the literature follow varied strategies. Most are based on sampling until group balance is reached, and while some strategies do it randomly [8], others include restrictions or preferences, such as only balancing groups close to the model's decision boundaries [7] or using particular clusters [9]. However, just balancing a priori ratios of protected feature values and classes does not always yield equal performance of the model (particularly, when the distribution of instances with each protected feature value is different).

The hypothesis of this work is that, to balance model performance on different data groups, the focus must be on the group overlap areas of the problem, as those would be the sections of the data space where sampling would be the most effective at modifying model decisions and, thus, at improving Fairness. To this aim, using data morphology, which has proven to be a remarkable strategy for characterizing groups or classes and detecting overlap amongst them, would be a well-founded approach.

Amongst the strategies that locate and measure overlap in tabular datasets, the use of 'Overlap Number of Balls' (ONB) is proposed [10]. This technique is based on the coverage of data using balls that only include points from a single group. ONB employs morphology to characterize the different groups in data, as well as to estimate group overlap and

classifier performance. Furthermore, the information on the boundaries amongst classes or groups provided by the ball coverage can also be harnessed for model explainability [11].

The aim of this work is to leverage the overlap estimation capabilities of ONB and the characterization of groups (obtained using class and protected feature value combinations) derived from its coverage strategy to improve Fairness in classification models according to protected feature values without a noticeable reduction in model performance. To this end, we propose a preprocessing method via sampling called Fair Overlap Number of Balls (Fair-ONB), which is based on group morphology and data neighborhoods. Since undersampling allows for the reduction of both bias and noise (particularly due to using morphology to locate it) and oversampling would only tackle the former, the proposed strategy will perform undersampling.

The Fair-ONB method involves two stages: selecting the undersampling groups, and then using percentile thresholds on coverage ball attributes, such as ball radius, number of instances covered and density, to decide how to undersample. Unlike random sampling strategies [8,12], it produces stable results due to using empirical measures to decide which samples to eliminate, and the location of balls under the chosen thresholds (and the change in group size that this would entail) can be easily quantified.

The effectiveness of the proposed Fair-ONB method will be tested on several of the most broadly-used real datasets in data Fairness studies. The a priori behavior of the different threshold variables for instance filtering is studied, both in terms of predictive performance and Fairness. The results are compared with FAWOS, a state-of-the-art, neighborhood-based sampling technique [7].

The rest of this paper is organized as follows. Firstly, Section 2 presents the preliminaries on Fairness, discrimination and bias. Then, Section 3 details the proposed Fair-ONB method. Next, Section 4 describes the experimental framework for the evaluation of the Fair-ONB method's behavior. Section 5 shows the experimental results. Lastly, Section 6 presents the concluding remarks from this manuscript and some future work lines.

## 2. Preliminaries on fairness

This Section presents some preliminaries regarding bias in Machine Learning. Section 2.1 introduces important concepts and some of the different notions of Fairness, along with their associated metrics. Section 2.2 shows some of the techniques present in the state of the art for the minimization of bias, with an emphasis on sampling techniques.

### 2.1. On fairness metrics

The existence of discrimination amongst groups of people, that is, the treatment that inherently favors some of them, involves substantial ethical problems and is legally unacceptable<sup>3</sup>. Fairness in Machine Learning means the absence of bias or discrimination amongst groups of data. This Fairness is checked with respect to protected features, such as gender or race, which are variables in datasets whose unequal treatment would expose the existence of discrimination.

Fairness and bias can be measured in different ways according to the situation, as studied in multiple surveys [5,13,14]. For example, it might be preferable that an equal rate of acceptance for a grant between men and women, but when assigning a medical procedure,

the actual need for such treatment (which might be different according to gender) must also be considered.

According to these preferences, there are multiple possible approaches in order to determine if a classifier is biased [14]. For example, in some cases the aim is to maintain demographic parity, also named statistical parity, which involves the positive prediction (beneficial for the person) being independent on the protected feature; that is, given a prediction  $\hat{Y}$ , demographic parity is measured according to how different the chances of obtaining a prediction  $\hat{Y} = 1$  is for each value of protected feature  $P$ . This would create 2 similar metrics:

- *Statistical Parity Difference (SPD)*: it measures the difference between the probabilities of a positive classification between two protected feature values (Equation 1). In the optimum case (equal chances), its value is 0.

$$SPD = Pr(\hat{Y} = 1 | P = 0) - Pr(\hat{Y} = 1 | P = 1) \quad (1)$$

- *Disparate Impact (DI)*: it measures the ratio between the probabilities of a positive classification between two protected feature values (Equation 2). In the optimum case (equal chances), its value is 1. Given its extended use to check biases in the USA, this metric is commonly applied [13].

$$DI = \frac{Pr(\hat{Y} = 1 | P = 0)}{Pr(\hat{Y} = 1 | P = 1)} \quad (2)$$

In some cases, a different version of Disparate Impact is used, named Adapted Disparate Impact (ADI), whose values are between 0 and 1, with 1 being the optimum case. This makes ADI more tractable but overlooks bias direction. Its formula is given in Equation (3).

$$ADI = \begin{cases} ID & : ID \leq 1 \\ \frac{1}{ID} & : ID > 1 \end{cases} \quad (3)$$

Other extended metrics would involve equality of probability (also named equalized odds) and equality of opportunity, whose foci are the true positive and/or false positive rates being the same for different values of the protected feature. Their associated metrics are measured according to the difference in probabilities for each group and are named Equal Probability Difference (EPD) and Equal Opportunity Difference (EOD), whose formulae are given in Equations (4) and (5), respectively.

$$EPD = Pr(\hat{Y} = 1 | Y, P = 0) - Pr(\hat{Y} = 1 | Y, P = 1) \quad (4)$$

$$EOD = Pr(\hat{Y} = 1 | Y = 1, P = 0) - Pr(\hat{Y} = 1 | Y = 1, P = 1) \quad (5)$$

Unless they are based on the same types of probabilities (such as in SPD and DI), optimizing multiple bias metrics simultaneously is generally unfeasible [15]; therefore, studies often select one of them to measure Fairness.

## 2.2. Bias and discrimination reduction strategies

Having indicated multiple metrics that measure Fairness, knowing how to reduce biases is also important. The use of model bias reduction techniques so that they comply with the

current Fairness regulations is a currently thriving research area [6]. Consequently, there are many techniques for bias reduction, from adding Fairness restrictions to the model's objective function [16], to adversarial training [17], adding weights to data with specific protected feature values [18] or modifying the dataset so that it includes less bias before learning the model [7]. Since this study is focused on dataset modifications via sampling, the strategies in that area will now be outlined.

Most sampling strategies are mainly based on the equal treatment of the subgroups of data defined by the combinations of class and protected feature values (for example, the pair of  $Class = 1$  and  $Gender = F$ ). This is due to Fairness metrics taking both into account simultaneously and, thus, seeking only class balance or protected feature balance on their own would not be sufficient.

According to when these techniques that reduce classifier bias are applied, they can be divided into three groups: preprocessing strategies, which work on the base data; in-processing strategies, which work during the learning process of the model; or postprocessing strategies, which correct the model's results after its prediction.

*Preprocessing* strategies work on the training data that will be used to learn the classifier. The aim of these techniques is to detect whether the dataset includes a priori anomalies or biases amongst the different subgroups in order to modify it appropriately [19]. This way, the models would not learn and reflect the biased behaviors stemming from data collection.

This a priori bias reduction so that the posterior model can have similar learning capabilities for every group is usually obtained by balancing the number of samples for each subgroup (according to the combination of classes and protected feature values) in the dataset. To this end, the number of samples of the over-represented groups can be reduced (undersampling), or it can be increased for the under-represented groups (oversampling). This sampling is often done randomly in the appropriate subgroups [8,12], but sometimes specific preferences are introduced (such as randomly sampling on the model's boundaries [7,20] or amongst clusters [9,21,22]) or samples are selected iteratively [23]. Another preprocessing technique consists on the modification of data labels (massaging) of certain bias-inducing samples [24]. Albeit another technique is the elimination of the protected feature so that the model is not directly exposed to bias, this strategy does not always work as intended [12] since other features with high correlations with that protected feature can still be present and, thus, the influence of that same protected feature would be indirectly maintained; this would be the case, for example, of post code and race or wage.

Amongst preprocessing techniques, it would be important to highlight Fairness-Aware OverSampling (FAWOS) [7], a sampling strategy based on instance neighborhoods. FAWOS employs the  $k$  nearest neighbors method to label instances as safe (4-5 of its neighbors share its class), borderline (2-3 of its neighbors share its class), rare (only 1 neighbor shares its class) and outlier (no neighbors share its class). The labels of subgroups with positive class and protected feature values that are discriminated against are used to grant different probabilities for those instances for Synthetic Minority Over-sampling Technique (SMOTE) usage, thus balancing the groups.

Regarding *inprocessing* and *postprocessing* techniques, the learning or modification of the model can allow it to work with biased datasets directly. In these cases, different weights can be assigned to data from different groups so that model learners would favor those that would otherwise be disadvantaged [25,26]. Constraints can also be applied so that the model is learned while considering Fairness [27,28]. Another strategy would

be using multiobjective optimization, including both Fairness and performance at the same time [23], or using adversary learning where certain protected feature values are swapped [26,29]. Other techniques include a posteriori massaging of labels [9] or interaction with experts [23,30].

### 3. The fair overlap number of balls method

This Section describes in detail the functioning of the Fair-ONB method, the novel under-sampling proposal based on data morphology. This method undersamples groups with favorable bias so that protected groups are treated similarly by the classifier. This is done in a guided manner, harnessing the morphology of the group coverage obtained from the Overlap Number of Balls (ONB) algorithm and its ball properties in order to select problematic areas of the data space.

Firstly, the ONB algorithm is explained in Section 3.1. Then, some preliminaries on the choices made during the method's creation are indicated in Section 3.2. Lastly, the structure and inner workings of the method are introduced in Section 3.3.

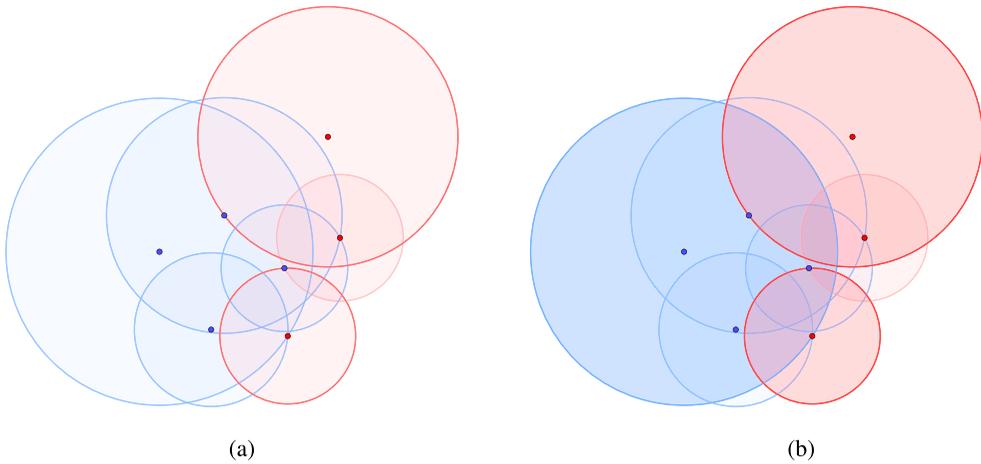
#### 3.1. Overlap number of balls

ONB is a class boundary estimation method that was initially conceived as a complexity metric, that is, a technique that measures how difficult to classify a dataset is according to its inner characteristics [10]. ONB is based on the Pure Class Cover Catch Digraph (P-CCCD) coverage algorithm [31].

Its strategy is simple. Firstly, for each instance in the dataset, the biggest open ball (the interior of the hypersphere of the same radius) that does not cover any points of a different class is generated. Then, for each class, the balls that would cover the most points that are not yet covered by other balls in the coverage set are iteratively chosen as part of said coverage set, until all points of that class are covered. An example of this is shown in Figure 1. ONB has an overall complexity of  $O(kn^3)$ , where  $k$  is the number of classes and  $n$  is the number of instances to cover.

This coverage strategy has very beneficial implications.

- The coverage locates and maps classes in the dataset. This estimation of the class distribution gives information on where it would be normal to find points of each class and where their presence could be considered unusual or noisy.
- The number of balls that are necessary to cover the dataset gives information on both how difficult to classify the dataset is and the degree of overlap amongst classes. This information is very useful towards deciding whether preprocessing is necessary in order to improve classifier performance [10].
- The ball coverage can estimate class boundaries, either in the base dataset or in the classifier that is learnt from it. In the latter case, the morphology of these estimated boundaries can be used to improve the explainability of complex classification models. In particular, Overlap Number of Balls Model-Agnostic Counterfactuals (ONB-MACF) is a counterfactual generator that employs the boundary estimations of ONB to provide counterfactual explanations [11].



**Figure 1.** Example of the P-CCCD ball coverage algorithm that serves as the base for ONB. Firstly, the maximum radii for the balls centered on each instance that do not include points of other classes are calculated (a), and afterwards, the balls that cover the most points of their class are iteratively selected until all points are covered (b).

**Table 1.** Groups obtained from the combinations of protected feature and class values.

Race	Gender	Class	Group
0	0	0	0
0	0	1	1
0	1	0	2
0	1	1	3
1	0	0	4
1	0	1	5
1	1	0	6
1	1	1	7

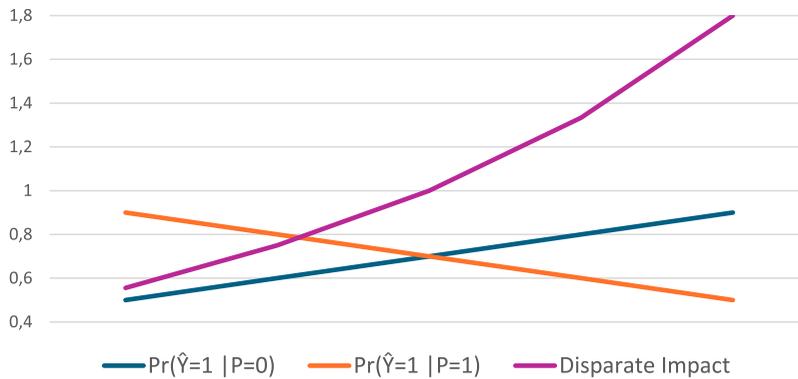
### 3.2. Preliminaries on the Fair-ONB method

Given a dataset whose protected features and class are known, the combination of their values is employed to detect and correct biases, given their paired involvement in Fairness estimation formulae. For the Fair-ONB method, each of these combinations forms a group. For example, a dataset with binary protected features “race” and “gender” and a binary class would have the groups shown in Table 1.

The aim is to see how those groups are distributed since, depending on their overlap and their degrees of imbalance, the classifier learned using the dataset can be biased.

At the same time, before employing the method, the existence of a priori biases in the dataset or model is checked using a Fairness metric. In this study, said Fairness metric is the Disparate Impact (DI, Equation 2), which, as a reminder, is the ratio between the probabilities of obtaining a positive result ( $\hat{Y} = 1$ ) according to the values in the protected feature. This first check allows the selection of which subgroups need to be preprocessed.

As shown in Equation (2), when no bias is present according to a protected feature,  $DI = 1$ ; otherwise, the value is still considered good when it is between 0.8 and 1.25 [13].



**Figure 2.** Disparate impact plot as the preferences of the classifier change.

Figure 2 shows the behavior of DI as the preferences of the model change (as happens when its training data is modified). It includes the neutral DI value (1) and the positive classification probabilities for the different values of a protected feature.

With this in mind, DI can be checked on the base dataset (with the existing ratios, to detect pre-existing bias in the dataset) or on the model (using the predictions given by the classifier, which elucidates whether the model is biased) for each of the protected features. When  $DI > 1$ , the bias favors the group with protected feature value 0; when  $DI < 1$ , it is biased towards the feature value 1.

### 3.3. Structure and functioning of the Fair-ONB method

The Fair-ONB method is based on the characterization of groups derived from the combination of class and protected feature values in the dataset. This is done via the data coverage using balls obtained from ONB (see Section 3.1).

Once the coverage is obtained, the Fair-ONB method undersamples instances of favored groups to foster a fair classification. This is done in a guided way, selecting instances in problematic areas of the data space after harnessing the properties of coverage balls.

The Fair-ONB method selects the balls for undersampling in two stages: first, the groups whose balls can be selected for undersampling are chosen; then, using their characteristics, some balls of those groups (and the instances they contain) are eliminated.

#### 3.3.1. Stage 1: group selection for preprocessing

When Fairness metrics indicate a bias towards a value of a protected feature in the dataset (which would likely translate into that same bias being learnt by the subsequent classifier), the chosen strategy is to perform undersampling of the groups with positive class ( $\hat{Y} = 1$ ) and said value (either 0 or 1) in the protected feature. Since there can be multiple protected features simultaneously, there would be two ways of choosing the groups to preprocess.

- Using the union of the groups with positive class and favored values in the different protected features (OR). Thus, if in the example of Table 1 there are biases

favoring  $race = 1$  and  $gender = 0$ , the groups to be preprocessed would be groups 1, 5 and 7.

- Using the intersection of the groups with positive class and favored values in the different protected features (AND). Likewise, if in the example of Table 1 there are biases favoring  $race = 1$  and  $gender = 0$ , only group 5 would be chosen for preprocessing.

Depending on the dataset characteristics, the chosen undersampling strategy can greatly affect the results. For that reason, both group selection methods will be studied.

### 3.3.2. Stage 2: selection of balls from those groups for elimination

Once the ball coverage has been obtained, some ball attributes that might indicate noisy areas or group overlap that can favor the existence of bias are studied. In particular, the radius, the number of instances they cover and the density of balls will be observed. The hypotheses regarding their possible effects on data bias are now described.

- *Radius*: if it is very small, the ball is in an area of group overlap, be it due to data distribution or to the existence of noise.
- *Number of covered instances*: if it covers very few instances, the ball might not be truly representative in the characterization of its class. It can also be in an area of group overlap or noise.
- *Density*: if it is very low, the ball might have been generated using an isolated point far from the decision boundary or just very far inside a class' area, which might not contribute much to the model's decisions. While some balls with very high density can also be problematic (due to, for example, having a very small radius), these particular cases are already covered by the other premises, so highly dense but useful balls will not be affected.

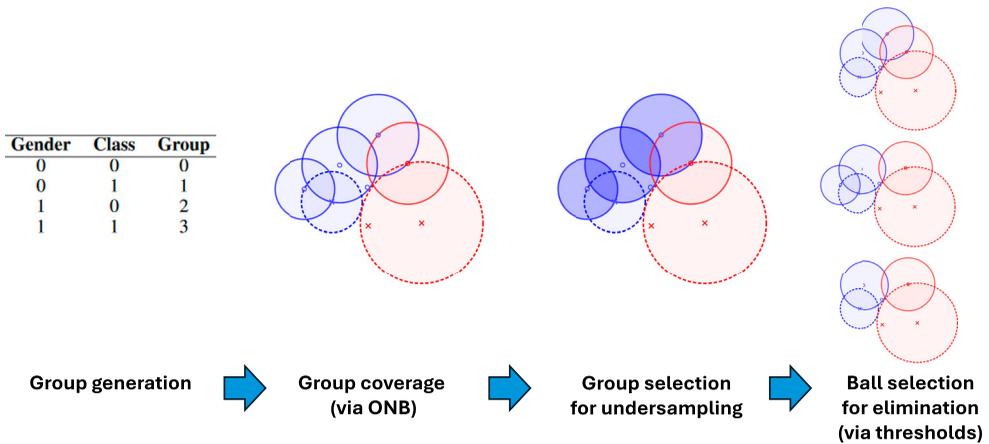
To perform the undersampling, percentiles (values of an attribute for which exactly that percentage of cases have a lower value) of those three characteristics in the coverage are calculated and will serve as thresholds for ball selection. Using percentiles instead of predefined values aims to have a better adaptability to each dataset.

Once obtained, the thresholds of all three ball characteristics are combined to decide different degrees of preprocessing (for example, one test can simultaneously use percentile 5 for radius, 15 for number of covered instances and 10 for density). The balls from the selected preprocessing groups whose radii, number of covered instances or densities are lower than each test's chosen ball characteristic thresholds are eliminated, along with the instances they cover, due to being considered problematic.

The overall modus operandi of the Fair-ONB method is shown in Figure 3.

## 4. Experimental framework

This Section presents the experimental framework that tests the goodness of the proposed Fair-ONB method, with the aim of improving model Fairness while maintaining performance. Section 4.1 presents the datasets that will be used in the experiments. Section 4.2 explains the structure of the experiments and their parameters.



**Figure 3.** The overall strategy of the Fair-ONB method starts with the generation of the groups defined by the pairs of class and protected feature values. The group coverage is then performed using ONB (in the plot, color is used to discern protected feature values, and instance symbols and circumference design to indicate class values). According to a chosen Fairness metric, groups are selected for undersampling. Finally, undersampling is performed, eliminating balls of that group according to thresholds (in the example, radius, density and number of covered instances were employed, respectively, from top to bottom).

#### 4.1. Datasets

This experimental framework includes the use of four datasets. The first three (COMPAS, Adult and German) are the most used in the field of bias reduction [32]; the fourth dataset, Ricci, was included to add more broadness to the study.

The exact versions of the aforementioned datasets have been modified to reduce the computational complexity: categorical variables were transformed into binary ones via one-hot encoding, and in the case of Adult, its size was cut in half (while maintaining its structure via stratified sampling). The dataset characteristics (after those transformations) are as follows.

- The COMPAS dataset predicts whether convicts will re-offend in the subsequent two years, according to their personal data and criminal history. Once transformed, it includes 6172 samples, 7 attributes (3 of them binary, 4 numeric), plus the binary class. The protected features are race and sex.
- The Adult dataset evaluates whether a person will earn over 50.000\$/year according to census data. Once transformed, the reduced dataset includes 24.416 samples, 13 attributes (7 of them binary, 6 numeric), plus the binary class. The protected features are race and sex.
- The German dataset predicts the credit risk of people. Once transformed, it includes 1.000 samples, 24 features (14 of them binary, 10 numeric), plus the binary class. The protected features are gender and age.
- The dataset Ricci evaluates the promotion of a group of firefighters according to exam results. It includes 118 samples, 4 attributes (2 of them binary and 2 numeric), plus the binary class. The protected feature is race.

## 4.2. Experimental structure and parameters

This study has two parts, each of them with a different aim.

- (1) Testing the behavior of the different threshold variables (according to their percentiles) in different situations, to detect whether there are certain thresholds that generally give good results.
- (2) Comparing the bias reduction when using the Fair-ONB method to that of FAWOS [7], a recent neighborhood-based sampling strategy with whose strategy Fair-ONB shares some similarities.

In order to properly evaluate classifier performance, 5-fold cross-validation is used in each experiment for each dataset. Sampling methods are only applied on each experiment's training sets. The predictive performance is evaluated over the test set using two usual metrics: AUC and accuracy. AUC is used in the first part of the study due to being a more conservative metric and its aptitude in case of imbalance; accuracy is used in the second part, once the good behavior of the Fair-ONB method has been proven, and to allow for fair comparison with FAWOS (which samples groups until they are approximately even, a situation where accuracy is a useful metric). Fairness is evaluated using Disparate Impact. Performance and Fairness results are aggregated via the arithmetic mean of the 5 folds'.

For dataset preprocessing using the Fair-ONB method, up to the 5 first different percentiles are used for each threshold feature (radius, number of covered instances, density), starting with 0 and going up in increments of 5. The combination of percentiles for each threshold feature parameterizes the different experiments for each dataset. Moreover, the experiments are performed using both union and intersection for undersampling group selection.

The classifier used is a decision tree with its default parameters in version 0.23.2 of scikit-learn<sup>4</sup> (the necessary version for FAWOS), except for including a fixed seed (*random\_state* = 30). These classifiers are learned using the preprocessed training set of each experiment.

In the second part of the study, the Fair-ONB method is compared to FAWOS [7]. The fact that FAWOS is also a sampling method based on neighborhoods makes it the closest strategy to Fair-ONB. FAWOS employs 5NN to label samples as safe, borderline, rare and outlier according to how many of their neighbors share their class. Then, labels are used to allocate probabilities of those samples being selected for SMOTE.

The parameters used for FAWOS are the same as in the original study [7]. The weights for each label group were 0 for safe, 0.6 for borderline and 0.4 for rare points (outliers always get null weight), and oversampling factors were either 0.8, 1 or 1.2.

## 5. Experimental results

This Section presents the experimental results of the two studies. Firstly, Section 5.1 evaluates the effects of selecting different undersampling thresholds; then, Section 5.2 presents the comparison of the results using Fair-ONB and FAWOS, two bias reduction sampling methods.

### 5.1. Behaviour of threshold features in the Fair-ONB method

This Section presents the results of the experiments on threshold combinations. Disparate Impact is measured according to the different protected variables of each of the four datasets, as well as classifier performance according to AUC, for each preprocessing parametrisation. Thus, the aims are to check whether Fairness can be improved without much effect on AUC, and whether certain thresholding parameters can be widely considered useful. Two sets of experiments are performed on each dataset, depending on whether union or intersection is used for undersampling group selection.

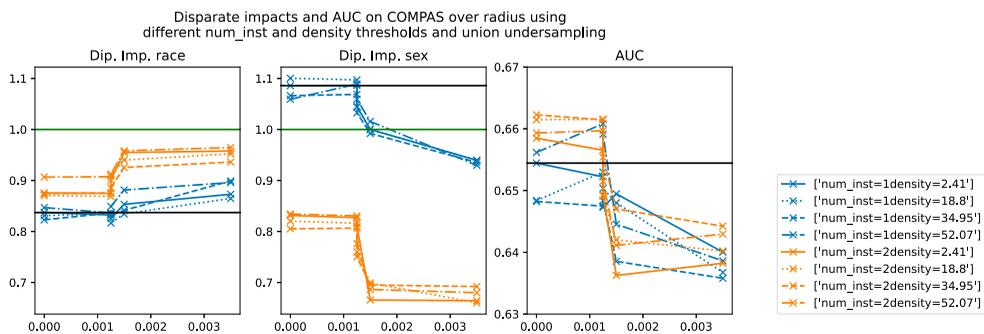
To better exemplify the different observed behaviors of union and intersection undersampling, the results on COMPAS and Adult will be plotted. To facilitate a global comparison, in this type of figure, Disparate Impact values for both protected variables in the dataset and AUC values are shown simultaneously; furthermore, to better reflect whether the method improves the results, the base results (with no preprocessing) are shown as a horizontal black line and the optimum Disparate Impact (1) as a green one. In each case, the radius will be used as the X-axis feature due to providing the most illustrative results.

#### 5.1.1. COMPAS dataset

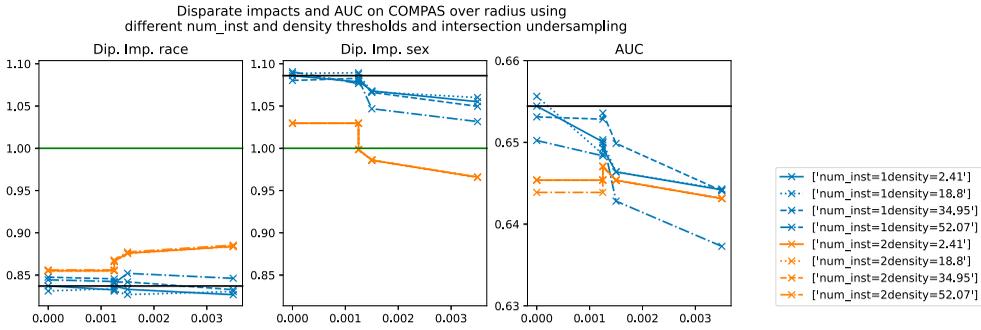
Figures 4 and 5 respectively depict the union and intersection undersampling results on this dataset. As can be observed, the results when using different thresholds for the number of covered instances create result clusters regarding Disparate Impact, with noticeable differences. Meanwhile, an interesting behavior is observed for intermediate values of radius thresholds, where Disparate Impact varies greatly in this dataset. Substantial Disparate Impact improvements can be attained on COMPAS (to the point of reaching the optimum on “sex” with both types of undersampling, although with different parameters). AUC, while slightly reduced, is comparable to the base value.

#### 5.1.2. Adult dataset

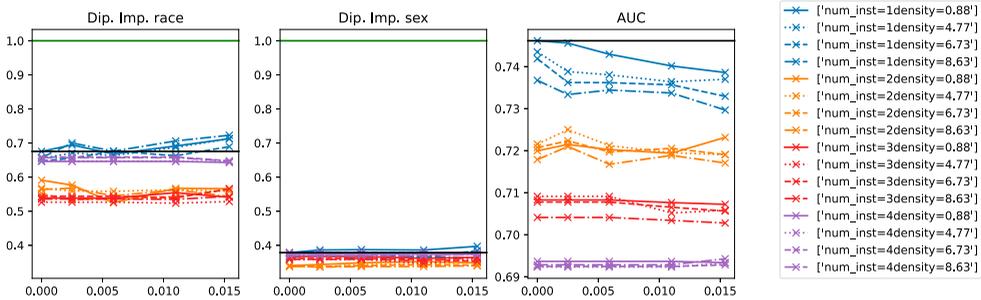
Figures 6 and 7 show the union and intersection undersampling results on this dataset. Here, union undersampling does not obtain noticeable improvements. Only slightly better Disparate Impact results were obtained for both protected variables. However, with intersection undersampling, Disparate Impact improves considerably for both protected



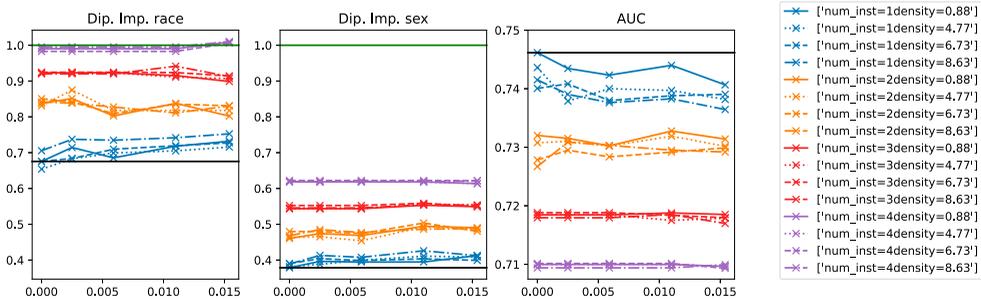
**Figure 4.** Behaviour of Disparate Impact according to the protected features (race and sex) and of AUC when the radius threshold is modified and the other threshold variables are fixed, using union undersampling on COMPAS.



**Figure 5.** Behaviour of Disparate Impact according to the protected features (race and sex) and of AUC when the radius threshold is modified and the other threshold variables are fixed, using intersection undersampling on COMPAS.



**Figure 6.** Behaviour of Disparate Impact according to the different protected features (race and sex) and of AUC when the radius threshold is modified and the other threshold variables are fixed, using union undersampling on Adult.



**Figure 7.** Behaviour of Disparate Impact according to the different protected features (race and sex) and of AUC when the radius threshold is modified and the other threshold variables are fixed, using intersection undersampling on Adult.

variables, reaching the optimum in “race”. In the case of “sex”, much better results than the baseline are also obtained while not reaching the often recommended [0.8,1.25] Disparate Impact interval.

### 5.1.3. Other datasets

For the *German dataset*, only slight improvements were obtained on “gender” since it was already very close to the optimum Disparate Impact, while noticeable improvements were attained according to the protected feature “age”. In this dataset, like in *Adult*, very similar global Disparate Impact results were obtained when using union and intersection undersamplings, and in both cases, it was possible to improve both Fairness and performance simultaneously. As for the *Ricci dataset*, there is only one protected feature (“race”), so it makes no sense to separate union and intersection (since there is only one possible group to undersample). Slight improvements were obtained compared to the base case.

### 5.1.4. Lessons learned

A summary of the results obtained with all datasets, including the best results of each type of undersampling (and the base case for comparison), is presented in Table 2. For each dataset, the best results on each individual metric are presented in bold, while the results of the parametrisation that minimizes DI globally (that is, the one that minimizes the total distance to the optimum disparate impacts amongst the protected features) are presented in bold and italics.

All in all, according to the results shown in Figures 4–7, as well as the overall results in Table 2, the following conclusions can be reached.

- The choice between union and intersection undersampling can greatly affect the minimization of Disparate Impact, as shown by the results on *Adult* (where intersection was clearly the better choice). Given the fact that in the other cases both union and intersection fared similarly to one another, it would seem that using intersection undersampling is the better choice overall.
- Choosing the right number of covered instances threshold is always very important, as this threshold defines the different result clusters in the Disparate Impact plots. One of the extremes amongst the percentiles chosen (being either the lowest or highest value) gave the best results in all datasets, although which one is the right choice depends on the dataset.
- The radius threshold shows the behavior/shape of clusters defined by the number of covered instances threshold, whether it is mostly constant, increasing, decreasing, sigmoidal, etc. Depending on the location of the plotted cluster compared to the optimum Disparate Impact, different radius threshold values will be optimal.
- The density threshold gives different heights in the results cluster, so using different threshold values helps fine-tune Disparate Impact depending on the cluster’s location compared to the optimum Disparate Impact.

## 5.2. Sampling methods comparison

This Section presents the results from the comparative study between Fair-ONB and FAWOS, two bias reduction sampling methods that are based on data neighborhoods.

While the initial intention was to perform the comparison using all 4 datasets indicated in Section 4.1, the particularities of FAWOS would make using *Adult* unfeasible. The computational complexity of FAWOS, which is not prepared to work on datasets that big (even having already reduced its number of samples), is too high, as it uses 5NN for each sample

**Table 2.** Comparison of Disparate Impact and AUC between the base COMPAS, Adult, German and Ricci datasets and the best union and intersection undersamplings, along with the best undersamplings for the optimization of each individual metric.

Dataset	US Type	Aim	US Thresholds			DI by protected feature		Perf.
			N <sub>i</sub>	R	D	Race DI	Sex DI	AUC
COMPAS	Baseline	–	–	–	–	0.837	1.086	0.654
	Union	Global DI	1	0.0015	52.07	0.881	1.016	0.645
		Race DI	2	0.0035	52.07	<b>0.965</b>	0.680	0.643
		Sex DI	1	0.0015	2.41	0.853	<b>1.000</b>	0.649
		AUC	2	0	34.95	0.874	0.805	0.662
	Intersec	Global DI	2	0.0013	52.07	<b>0.868</b>	<b>0.999</b>	<b>0.647</b>
		Race DI	2	0.0035	52.07	0.885	0.966	0.643
		Sex DI	2	0.0013	2.41	0.866	0.999	0.647
		AUC	1	0	2.41	0.831	1.089	<b>0.656</b>
	Adult	Baseline	–	–	–	–	0.675	0.379
Union		Global DI	1	0.0154	0.88	0.713	0.397	0.739
		Race DI	1	0.0154	8.63	0.723	0.363	0.730
		Sex DI	1	0.0154	0.88	0.713	0.397	0.739
		AUC	1	0	0.88	0.675	0.379	<b>0.746</b>
Intersec		Global DI	4	0.0154	6.73	<b>1.006</b>	<b>0.622</b>	<b>0.709</b>
		Race DI	4	0	8.63	<b>0.995</b>	0.619	0.709
		Sex DI	4	0	6.73	0.983	<b>0.622</b>	0.710
		AUC	1	0	0.88	0.675	0.379	<b>0.746</b>
German		Baseline	–	–	–	–	1.028	1.295
	Union	Global DI	1	0.7109	0.49	1.062	<b>1.043</b>	0.646
		Gender DI	1	0.4383	0.22	1.015	1.269	0.643
		Age DI	1	0.7109	0.49	1.062	<b>1.043</b>	0.646
		AUC	1	0.7109	0.22	1.139	1.270	0.599
	Intersec	Global DI	1	0.5810	0.44	<b>0.984</b>	<b>1.080</b>	<b>0.618</b>
		Gender DI	1	0.8257	0.36	<b>1.000</b>	1.419	0.633
		Age DI	1	0.581	0.44	0.984	1.080	0.618
		AUC	2	0.0363	0.44	1.229	1.229	<b>0.649</b>
	Ricci	Baseline	–	–	–	–	0.501	–
Union		Race DI	1	0	5.49	<b>0.513</b>	–	<b>0.916</b>
		AUC	1	0	5.49	<b>0.513</b>	–	<b>0.916</b>

multiple times per experiment. While the sample size could have been reduced again to include the dataset in this study, the drastic reduction necessary would likely distort the dataset characteristics.

In this study, accuracy was used to measure classifier performance, and since FAWOS uses the Adapted Disparate Impact, which is a simple transformation of Disparate Impact (see Equation 3), that was the chosen Fairness metric.

The best results obtained with each method (and their optimal parameters) on the COMPAS, German and Ricci datasets are presented in Tables 3–5, respectively. As a reminder, FAWOS parameters are the weights for safe (S), borderline (B) and rare (R) samples (as outliers receive weight 0), as well as the oversampling factor (OF); meanwhile, the

**Table 3.** Comparison, according to Adapted Disparate Impact and Accuracy, between the best FAWOS and Fair-ONB results on the COMPAS dataset.

Method	Best Parameters	Race ADI	Sex ADI	Tot. Dist. Opti.	Accuracy
FAWOS	Weights: $S = 0$ , $B = 0.6$ , $R = 0.4$ , $OF = 0.8$	0.474	0.758	0.768	<b>0.780</b>
Fair-ONB Intersec.	Thresholds: $N_i = 3$ , $R = 0.00$ , $D = 2.47$	<b>0.886</b>	0.972	0.142	0.771
Fair-ONB Union	Thresholds: $N_i = 1$ , $R = 0.00$ , $D = 34.37$	0.873	<b>0.999</b>	<b>0.128</b>	0.747

**Table 4.** Comparison, according to Adapted Disparate Impact and Accuracy, between the best FAWOS and Fair-ONB results on the German dataset.

Method	Best Parameters	Age ADI	Gender ADI	Tot. Dist. Opti.	Accuracy
FAWOS	Weights: $S = 0$ , $B = 0.6$ , $R = 0.4$ , $OF = 1$	<b>0.820</b>	0.910	0.270	<b>0.686</b>
Fair-ONB Intersec.	Thresholds: $N_i = 1$ , $R = 0.57$ , $D = 0.36$	0.796	<b>0.942</b>	<b>0.262</b>	0.672
Fair-ONB Union	Thresholds: $N_i = 1$ , $R = 0.80$ , $D = 0.22$	0.789	0.925	0.286	0.682

**Table 5.** Comparison, according to Adapted Disparate Impact and Accuracy, between the best FAWOS and Fair-ONB results on the Ricci dataset.

Method	Best Parameters	Race ADI	Accuracy
FAWOS	Weights: $S = 0$ , $B = 0.6$ , $R = 0.4$ , $OF = 1.2$	0.282	0.910
Fair-ONB	Thresholds: $N_i = 1$ , $R = 0.04$ , $D = 5.52$	<b>0.467</b>	<b>0.917</b>

parameters of the Fair-ONB method are the number of covered instances ( $N_i$ ), radius ( $R$ ) and density ( $D$ ) thresholds.

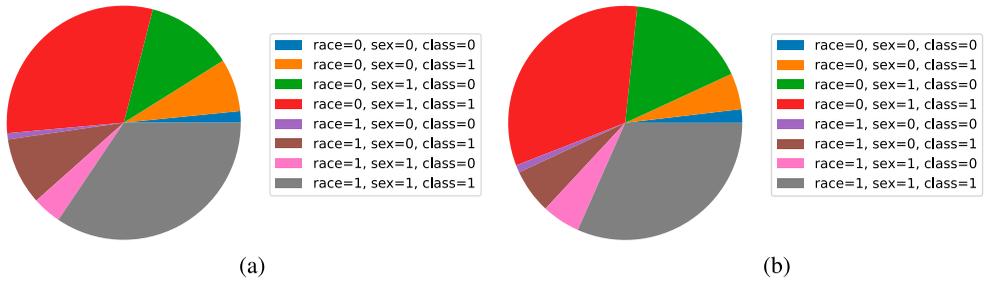
As can be observed, the Fair-ONB method obtains the best results according to Global Adapted Disparate Impact (as indicated by the lower total distance to the optimum Adapted Disparate Impact amongst protected features) and, very often, also the best Adapted Disparate Impact on the individual protected features. Given the usual inverse relationship between predictive performance and bias mitigation, it makes sense that FAWOS obtains better accuracy results.

As a possible reason for this, Figure 8 shows two pie charts with the group distributions in the base COMPAS dataset and when the aforementioned best-performing union undersampling is used. As can be observed, the best result is obtained without group balance, which is what most other methods (including FAWOS) would pursue, potentially giving Fair-ONB the edge. This also exemplifies the importance of correctly harnessing the morphology of protected groups towards modifying their proportions, as sampling in the most relevant parts of the dataset outweighed random sampling in group balancing (even if, in FAWOS, it was informed by neighbors according to 5NN).

## 6. Concluding remarks

This paper presents Fair-ONB, a new sampling method that employs guided undersampling in areas that are close to the decision boundaries, where models usually have problems classifying the different protected groups. The Fair-ONB method harnesses the morphology of groups to select samples for elimination, and unlike other random sampling techniques, its logic is more justifiable due to its empirical approach.

As shown in the experimental study, the Fair-ONB method improves model Fairness while maintaining or improving classification performance, reflecting the usefulness of



**Figure 8.** Example of COMPAS group proportions in the base case (a) and with the best-performing union undersampling using the Fair-ONB method (b), showing that achieving group balance was not the optimal strategy.

using subgroup ball coverage for undersampling instance selection. Moreover, the Fair-ONB method also achieves better bias reduction results than FAWOS, its neighborhood-based competitor in the state of the art.

Therefore, this paper has shown the usefulness of morphology and neighborhood-based sampling approaches towards classifier Fairness, a niche that has not been thoroughly explored. In particular, the use of threshold variables and group coverage strategies, which give information on multiple morphological properties of the groups in a dataset, have been useful for the proper characterization of problematic overlap areas. Thus, the hypothesis that using data morphology for undersampling instance selection can be a very useful strategy is confirmed.

Regarding future work lines, combining morphology-based strategies with other preprocessing techniques could lead to attaining higher fairness levels. Extending the experiments and employing other data complexity metrics to evaluate them from other points of view might also provide more information on the right thresholds for undersampling experiments, reducing the necessary computations. Using a bigger scope and finding synergies could allow for the implementation of more in-depth preprocessing techniques.

## Notes

1. Regulation (EU) 2024/1689 of the European Parliament and of the Council <http://data.europa.eu/eli/reg/2024/1689/oj>
2. Artificial Intelligence Act, <https://artificialintelligenceact.eu/>
3. Regulation (EU) 2024/1689 of the European Parliament and of the Council <http://data.europa.eu/eli/reg/2024/1689/oj>
4. <https://scikit-learn.org/stable/modules/generated/sklearn.tree.DecisionTreeClassifier.html>

## Disclosure statement

No potential conflict of interest was reported by the author(s).

## Funding

This work has been partially supported by Knowledge Generation Projects, funded by the Spanish Ministry of Science, Innovation, and Universities of Spain under the project PID2023-150070NB-I00. It also received funding from the I+D+i project granted by C-ING-250-UGR23 co-funded by

“Consejería de Universidad, Investigación e Innovación” and the European Union related to FEDER Andalucía Program 2021–27. Moreover, the work was partially funded by the Growing Resilient, INclusive and Sustainable (GRINS PE00000018) project, which is part of NextGenerationEU, from the European Union. The views and opinions expressed are solely those of the authors and do not necessarily reflect those of the European Union, nor can the European Union be held responsible for them.

## ORCID

José Daniel Pascual-Triana  <http://orcid.org/0000-0003-4106-3182>

## References

- [1] Hassani H, MacFeely S. Driving excellence in official statistics: unleashing the potential of comprehensive digital data governance. *Big Data Cogn Comput.* 2023;7(3):134. doi: [10.3390/bdcc7030134](https://doi.org/10.3390/bdcc7030134)
- [2] Aggarwal CC. *Artificial intelligence: a textbook*. Cham: Springer International Publishing; 2021.
- [3] Luengo J, García-Gil D, Ramírez-Gallego S, et al. *Big data preprocessing: enabling smart data*. Cham: Springer; 2020.
- [4] Triguero I, García-Gil D, Mailló J, et al. Transforming big data into smart data: an insight on the use of the k-nearest neighbors algorithm to obtain quality data. *WIREs Data Min Knowl Discov.* 2019;9(2):e1289. doi: [10.1002/widm.2019.9.issue-2](https://doi.org/10.1002/widm.2019.9.issue-2)
- [5] Mehrabi N, Morstatter F, Saxena N, et al. A survey on bias and fairness in machine learning. *ACM Comput Surv.* 2021;54(6):35.
- [6] González-Sendino R, Serrano E, Bajo J, et al. A review of bias and fairness in artificial intelligence. *Int J Interact Multimed Artif Intell* 2023;9:13. doi: [10.9781/ijimai.2023.11.001](https://doi.org/10.9781/ijimai.2023.11.001).
- [7] Salazar T, Santos MS, Araújo H, et al. FAWOS: fairness-aware oversampling algorithm based on distributions of sensitive attributes. *IEEE Access.* 2021;9:81370–81379. doi: [10.1109/ACCESS.2021.3084121](https://doi.org/10.1109/ACCESS.2021.3084121)
- [8] Rančić S, Radovanović S, Delibašić B. Investigating oversampling techniques for fair machine learning models. In: Jayawickrama U, Delias P, Escobar MT, Papathanasiou J, editors. *Decision support systems XI: decision support systems, analytics and technologies in response to global crisis management*. Cham: Springer International Publishing; 2021. p. 110–123.
- [9] He F, Wu X, Zhang W, et al. A novel bias-alleviated hybrid ensemble model based on over-sampling and post-processing for fair classification. *Connect Sci.* 2023;35:2184310. doi: [10.1080/09540091.2023.2184310](https://doi.org/10.1080/09540091.2023.2184310)
- [10] Pascual-Triana JD, Charre D, Andrés Arroyo M, et al. Revisiting data complexity metrics based on morphology for overlap and imbalance: snapshot, new overlap number of balls metrics and singular problems prospect. *Knowl Inf Syst.* 2021;63(7):1961–1989. doi: [10.1007/s10115-021-01577-1](https://doi.org/10.1007/s10115-021-01577-1)
- [11] Pascual-Triana JD, Fernández A, Del Ser J, et al. Overlap number of balls model-agnostic counterfactuals (ONB-MACF): a data-morphology-based counterfactual generation method for trustworthy artificial intelligence. 2024. arXiv:2405.12326 [cs].
- [12] Valentim I, Lourenço N, Antunes N. The impact of data preparation on the fairness of software systems. In: *2019 IEEE 30th International Symposium on Software Reliability Engineering (ISSRE)*, Berlin, Germany; 2019. p. 391–401. ISSN: 2332–6549.
- [13] Caton S, Haas C. Fairness in machine learning: a survey. *ACM Comput Surv.* 2024;56(7):166:1–166:38. doi: [10.1145/3616865](https://doi.org/10.1145/3616865)
- [14] Pagano TP, Loureiro RB, Lisboa FVN, et al. Bias and unfairness in machine learning models: a systematic review on datasets, tools, fairness metrics, and identification and mitigation methods. *Big Data Cogn Comput.* 2023;7(1):15. doi: [10.3390/bdcc7010015](https://doi.org/10.3390/bdcc7010015)

- [15] Kleinberg J, Mullainathan S, Raghavan M. Inherent trade-offs in the fair determination of risk scores. In: *LIPICs, ITCS 2017*, Berkeley, CA, USA, Vol. 67, p. 43:1–43:23. Schloss Dagstuhl – Leibniz-Zentrum für Informatik (2017).
- [16] Wu H, Ma C, Mitra B, et al. A multi-objective optimization framework for multi-stakeholder fairness-aware recommendation. *ACM Trans Inf Syst.* 2022;41(2):29.
- [17] Reimers C, Bodesheim P, Runge J, et al. . Conditional adversarial debiasing: towards learning unbiased classifiers from biased data. *Lecture Notes in Computer Science* Vol. 13024. Cham: Springer International Publishing; 2021. p. 48–62 doi: [10.1007/978-3-030-92659-5\\_4](https://doi.org/10.1007/978-3-030-92659-5_4)
- [18] Yu Z, Chakraborty J, Menzies T. FairBalance: how to achieve equalized odds with data pre-processing. *IEEE Trans Softw Eng.* 2024;50:1–18. doi: [10.1109/tse.2024.3431445](https://doi.org/10.1109/tse.2024.3431445).
- [19] Kamiran F, Calders T. Data preprocessing techniques for classification without discrimination. *Knowl Inf Syst.* 2012;33(1):1–33. doi: [10.1007/s10115-011-0463-8](https://doi.org/10.1007/s10115-011-0463-8)
- [20] Sun Y, Haghighat F, Fung BCM. Trade-off between accuracy and fairness of data-driven building and indoor environment models: A comparative study of pre-processing methods. *Energy.* 2022;239:122273. doi: [10.1016/j.energy.2021.122273](https://doi.org/10.1016/j.energy.2021.122273)
- [21] Chakraborty I. Bayesian sampler: fairness in sampling. In: *2020 19th IEEE International Conference on Machine Learning and Applications (ICMLA)*, Miami, FL, USA; 2020. p. 246–251.
- [22] Sonoda R. Fair oversampling technique using heterogeneous clusters. *Inf Sci.* 2023;640:119059. doi: [10.1016/j.ins.2023.119059](https://doi.org/10.1016/j.ins.2023.119059)
- [23] Anahideh H, Asudeh A, Thirumuruganathan S. Fair active learning. *Expert Syst Appl.* 2022;199:116981. doi: [10.1016/j.eswa.2022.116981](https://doi.org/10.1016/j.eswa.2022.116981)
- [24] Sharma S, Zhang Y, Ríos Aliaga JM, et al. Data augmentation for discrimination prevention and bias disambiguation. In: *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society, AIES '20*; New York (NY): Association for Computing Machinery; 2020. p. 358–364.
- [25] Krasanakis E, Spyromitros-Xioufis E, Papadopoulos S, et al. Adaptive sensitive reweighting to mitigate bias in fairness-aware classification. In: *Proceedings of the 2018 World Wide Web Conference, WWW '18, International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, Lyon, France; 2018*, p. 853–862.
- [26] Zeng H, Yue Z, Kou Z, et al. Fairness-aware training of face attribute classifiers via adversarial robustness. *Knowl-Based Syst.* 2023;264:110356. doi: [10.1016/j.knosys.2023.110356](https://doi.org/10.1016/j.knosys.2023.110356)
- [27] Ferry J, Aïvodji U, Gambs S, et al. Improving fairness generalization through a sample-robust optimization method. *Mach Learn.* 2023;112(6):2131–2192. doi: [10.1007/s10994-022-06191-y](https://doi.org/10.1007/s10994-022-06191-y)
- [28] Liu X, Tong X, Liu Q. Sampling with trustworthy constraints: a variational gradient framework. In: *Advances in Neural Information Processing Systems*, Vol. 34; Curran Associates, Inc.; 2021. p. 23557–23568.
- [29] Zhang P, Wang J, Sun J, et al. Automatic fairness testing of neural classifiers through adversarial sampling. *IEEE Trans Softw Eng.* 2022;48(9):3593–3612. doi: [10.1109/TSE.2021.3101478](https://doi.org/10.1109/TSE.2021.3101478)
- [30] Wang H, Mukhopadhyay S, Xiao Y, et al. An interactive approach to bias mitigation in machine learning. In: *2021 IEEE 20th International Conference on Cognitive Informatics & Cognitive Computing (ICCI\*CC)*, Banff, AB, Canada; 2021. p. 199–205.
- [31] Manukyan A, Ceyhan E. Classification of imbalanced data with a geometric digraph family. *J Mach Learn Res.* 2016;17(1):6504–6543.
- [32] Fabris A, Messina S, Silvello G, et al. Algorithmic fairness datasets: the story so far. *Data Min Knowl Discov.* 2022;36(6):2074–2152. doi: [10.1007/s10618-022-00854-z](https://doi.org/10.1007/s10618-022-00854-z)