

# Fairness and bias correction in machine learning for depression prediction: results from four different study populations

Vien Ngoc Dang<sup>\*,a</sup>, Anna Cascarano<sup>a</sup>, Rosa H. Mulder<sup>b,c</sup>, Charlotte Cecil<sup>d,e,f</sup>, Maria A. Zuluaga<sup>g</sup>, Jerónimo Hernández-González<sup>h</sup>, Karim Lekadir<sup>a</sup>

<sup>a</sup>Artificial Intelligence in Medicine Lab, Facultat de Matemàtiques i Informàtica, Universitat de Barcelona, Spain

<sup>b</sup>Department of Pediatrics, Erasmus University Rotterdam, the Netherlands

<sup>c</sup>The Generation R Study Group, Erasmus MC, University Medical Center Rotterdam, Rotterdam, the Netherlands

<sup>d</sup>Department of Child and Adolescent Psychiatry/Psychology, ErasmusMC-Sophia, Rotterdam, the Netherlands

<sup>e</sup>Department of Epidemiology, Erasmus MC, Rotterdam, The Netherlands

<sup>f</sup>Molecular Epidemiology, Department of Biomedical Data Sciences, Leiden University Medical Center, Leiden, The Netherlands

<sup>g</sup>Data Science Department, EURECOM, Sophia Antipolis, France

<sup>h</sup>Departament de Matemàtiques i Informàtica, Facultat de Matemàtiques i Informàtica, Universitat de Barcelona (UB), Spain

\*Corresponding author. E-mail address: dangn@ub.edu

A significant level of stigma and inequality exists in mental healthcare, especially in under-served populations, which spreads through collected data. When not properly accounted for, machine learning (ML) models learned from data can reinforce the structural biases already present in society. Here, we present a systematic study of bias in ML models designed to predict depression in four different case studies covering different countries and populations. We find that standard ML approaches show regularly biased behaviors. However, we show that standard mitigation techniques, and our own post-hoc method, can be effective in reducing the level of unfair bias. We provide practical recommendations to develop ML models for depression risk prediction with increased fairness and trust in the real world. No single best ML model for depression prediction provides equality of outcomes. This emphasizes the importance of analyzing fairness during model selection and transparent reporting about the impact of debiasing interventions.

## 1 Introduction

Depression is a leading cause of disability worldwide, a major risk factor for the global burden of disease, and can even lead to suicide [1-2]. Taking into account that the global prevalence of depression increased by 25% during the COVID-19 outbreak [3]. Being able to identify those individuals at risk would be of great value to enable the application of personalized preventive measures. To this end, it is necessary to characterize the factors leading up to the development of depression in order to include appropriate predictors in the model. Research to date points to the importance of both genetic and environmental factors

(as well as their interaction) in the etiology of depression [4-5]. Furthermore, environmental factors have been shown to co-occur, exerting cumulative effects on depression risk. The totality of these environmental influences is often referred to as the exposome and includes environmental and lifestyle factors, as well as traumatic life events [6]. In addition to being necessary to get the full picture, exposome data are also relatively inexpensive and easy to acquire, typically through questionnaires [7], which effectively move ML models from paper to practice.

Motivated by the successful application of ML in different contexts of the medical domain, there is a spike in the use of ML for the detection, diagnosis, and treatment of depression. Specifically, supervised ML methods are commonly used to learn predictive models from historical data, which are then applied to predict possible illness development in new cases and patients. Recently, concerns have been raised about algorithmic bias [8] and the undesirable ability of ML models of amplifying unfair behaviors masked in past practice, that is, in the data used for model learning. The term “*algorithmic bias*” usually refers to differences in the predictive power of models when applied to different subgroups of the population. These differences are particularly pertinent if they are found when the subgroups are obtained according to some *protected attribute* such as ethnicity, socioeconomic status, sex, age, or co-morbidities. Within these subgroups, the subgroups that are adversely impacted by the bias are known as the *unprivileged groups*, and those that are unfairly benefited by the bias are the *privileged groups*. Former groups should be protected from the malicious effects of bias. Despite this undesirable behavior, it is not a common practice to assess the bias of ML models (and try to mitigate it) in healthcare applications. This focus is even more scarce in the mental health literature. In [8], an ML algorithm on psychiatric notes to predict 30-day psychiatric readmission regarding sex, ethnicity, and insurance type is examined without addressing algorithmic bias. In [9], they reduce bias for clinical prediction models of postpartum depression only associated with one protected attribute - binarized ethnicity (Black individuals and White individuals).

In mental health, several unintentional discriminative behaviors have been detected, which could potentially be reproduced by ML models if they are reflected in the training data. Specifically, a lack of representation of the patient subgroups has been reported; for example, some racial and ethnic groups do not use mental health services as much as others due to cultural stigma surrounding mental illness [10-11]. Prior research shows that the prevalence or incidence of depression differs across subgroups; for example, women are about twice as likely as men to develop depression during their lifetime [12]. In addition, manifest discrepancies in relevant factors such as lifestyle or nutritional habits between

subgroups have also been reported [13]. These, and possibly other factors, can be rooted in the data which is used to learn ML models.

In this paper, motivated by these known facts about depression, we investigate how ML algorithms could perpetuate or reduce structural inequality and unfair bias learned from the data. We present a systematic analysis of unfair bias in ML models designed to predict the presence or absence of depression from environmental and lifestyle data, using four public datasets: LONGSCAN [14], FUUS [15], NHANES [16], and the UK Biobank (UKB) [17]. We study unfair bias on protected attributes including demographic factors (sex, ethnicity, nationality), socioeconomic status (age, income, academic qualifications), and co-morbidities (cardiovascular disease (CVD), diabetes), and evaluate the interplay of model accuracy and fairness. We analyze the ability of standard bias mitigation techniques to reduce the discrimination level of the models learned for our four case studies. Discrimination is measured as the performance difference, in terms of fairness and standard ML metrics, before and after performing bias mitigation. We have found unfair biases in the behavior of the models learnt with standard ML techniques regarding several protected attributes in all the case studies. We also found, however, that mitigation techniques are effective in reducing discrimination levels. Our results demonstrate that bias monitoring is essential in the evaluation of ML-based predictive models in mental health and currently available bias mitigation techniques provide a powerful toolset to mitigate any systematic bias against protected groups.

## **2 Methods**

In this section, we introduce and perform an initial descriptive analysis of the available data of our case studies (Sec. 2.1). In the following, we describe the predictive models tested (Sec. 2.2) and the five strategies we considered to mitigate bias (Section 2.3). We also explain both standard ML and fairness evaluation metrics used in our experiments (Sec. 2.4). Finally, we explain several implement details for the sake of reproducibility (Sec. 2.5).

### **2.1 Datasets**

This study uses four public datasets: LONGSCAN, FUUS, NHANES, and UK Biobank. All the participants and/or their carers/guardians provided written informed consent in LONGSCAN, NHANES, and UKB studies. This was not required in the case of FUUS according to the laws that regulate “non-interventional clinical research” in France [15]. While LONGSCAN and FUUS are datasets of late adolescents, NHANES and UKB have most subjects between 40 and 80 years of age. Table 1 describes the protected attributes of these

datasets. Note that not all the datasets report the same protected attributes. The LONGSCAN, FUUS, NHANES, and UKB datasets have relatively equal proportions of male and female subjects. A higher prevalence of depression in women versus men is evident in the LONGSCAN, NHANES, and UKB datasets. No sex effect also is found with depression symptoms among college freshmen in the FUUS dataset, which is consistent with previous studies [18-20]. Other protected attributes have a highly skewed distribution.

Table 1. Protected attributes and other relevant features of the used datasets

Subgroup	Attribute	LONGSCAN	FUUS	NHANES	UKB
	No. of subjects	911	4,184	36,259	461,033
	No. of features	47	62	86	143
Sex, % (Depression rate, %)	Male	44.2 (19.4)	42.6 (12.2)	49.1 (6.5)	46.6 (3.2)
	Female	55.8 (36.0)	57.4 (12.9)	50.9 (10.9)	53.4 (4.5)
Age (%); Depression rate (%)	0–20 years	-	-	7.5 (7.1)	-
	20–40 years	-	-	31.7 (8.2)	1.1 (4.2)
	40–60 years	-	-	31.1 (10.8)	59.9 (4.0)
	60–80 years	-	-	24.1 (8.0)	39.0 (3.8)
	>80 years	-	-	5.7 (5.7)	-
Ethnicity (%); Depression rate (%)	White	24.7 (35.1)	-	41.9 (8.6)	93.9 (3.9)
	Mexican	-	-	16.2 (8.6)	-
	Other Hispanic	-	-	9.5 (11.8)	-
	Black	55.9 (24.4)	-	21.9 (8.9)	1.7 (3.2)
	Asian	-	-	-	2.4 (3.1)
	Other/Multiracial	19.4 (32.8)	-	10.6 (6.2)	1.5 (4.7)
Nationality (%); Depression rate (%)	French	-	92.4 (12.3)	-	-
	Foreigner	-	7.6 (16.2)	-	-
Income (%); Depression rate (%)	Low-income	-	-	36.6 (13.3)	39.9 (4.9)
	High-income	-	-	63.4 (6.1)	44.7 (2.8)

Qualifications (%) Depression rate (%)	Level 0	-	-	5.9 (7.0)	18.8 (5.4)
	Level 1	-	-	9.5 (12.9)	20.9 (3.9)
	Level 2	-	-	13.3 (13.2)	5.3 (5.0)
	Level 3	-	-	21.5 (9.3)	11.0 (3.6)
	Level 4	-	-	21.8 (8.9)	6.5 (4.4)
	Level 5	-	-	28.0 (3.9)	32.3 (2.9)
	Level 6	-	-	-	5.1 (4.0)
Diabetes (%) Depression rate (%)	Absence	-	-	-	94.4 (3.9)
	Presence	-	-	-	5.1 (6.3)
CVD (%) Depression rate (%)	Absence	-	-	-	76.8 (3.4)
	Presence	-	-	-	23.2 (5.6)

Qualification levels have different meaning in the NHANES and UKB datasets due to the difference in the United States and British school systems, more detailed in the Supplementary Table 7; The sum of all the percentages might not be equal to 100 as a result of missing data. Note that participants are all in the same age group in the LONGSCAN and FUUS datasets. The income attribute in the LONGSCAN is a time-varying variable. The FUUS dataset does not report ethnicity and income attributes.

*Participants and features.* Participants in LONSCAN were enrolled from five sites, located in different regions of the United States (South, East, Midwest, Northwest, and Southwest), with different selection criteria, representing varying levels of risk or exposure to maltreatment during the period spanning from 1991 to 2012. The LONGSCAN interview and questionnaire data were collected when target children were at ages 4, 6, 8, 12, 14, 16, and 18. Out of 1354 total participants, we kept the 67.3% of children who completed an interview at age 18 years including depression outcomes, leading to 911 samples available for our study. Among these 911 individuals, there were 363 cases with depression at the age of 18, and 548 controls. The study design is shown in Supplementary Figure 1: data from three different stages (early childhood, late childhood, teen) were collected to predict depression at age of 18. Up to 23 descriptive variables grouped as demographic variables, lifestyles variables, and adverse exposures variables, both time-invariant and repeatedly measured along these stages, were considered (see Supplementary Table 3). Note that the use of data for this study was approved by the National Data Archive on Child Abuse and Neglect (NDACAN). Participants in FUUS were undergraduate students who underwent a

compulsory medical visit at the university medical service in Nice (France) between September 2012 and June 2013. Among 4184 total participants, there are 528 cases with depression and 3656 controls. A total of 62 biomedical and demographic features were used including binary, ordinal and continuous variables (see Supplementary Table 4). Participants in NHANES provided data for this study between 2005 and 2018, and were selected by random sampling of the American population. Among 36259 total participants, there were 3168 cases with depression and 33091 controls. 86 potential predictors were used including demographic data, socioeconomic status, medical history, prescription medications, and lifestyle characteristics (see Supplementary Table 5). Our ML-based models for the FUUS and NHANES datasets are a screening tool for depression. Note that the FUUS and NHANES dataset were publicly available. Participants in UKB were enrolled from 22 assessment centers, located in England, Wales, and Scotland, during the period spanning between 2006 and 2010. Among the 461,033 participants initially without depression, 18112 cases (3.93%) developed depression. Our ML-based models predict a depression event within 10 years. Up to 143 descriptive variables were used, including demographic data, socioeconomic status, medical history, lifestyle characteristics, early life factors, and traumatic events (see Supplementary Table 6). Note that the use of data for this study was approved by UKB, under the project title “Association between Early-Life-Stress and Psycho-Cardio-Metabolic Multi-Morbidity: The EarlyCause H2020 Project” (application number 65769).

*Building the ground truth - depression outcome.* We acquired ground truth label information –whether a participant has depression or not– using dataset-specific information. In LONGSCAN, depression was assessed using a self-reported questionnaire at age 18, which includes a specific question regarding having depression. In NHANES, the Patient Health Questionnaire-9 (PHQ-9) was used. This screening questionnaire consists of 9 items (scored 0-3) and has a specificity and sensitivity of 88% for major depressive disorder (MDD) at a threshold score of 10 or more [21]. Therefore, we chose the threshold at PHQ-9 score 10. Additionally, we carefully excluded participants' feelings and expressions, as well as lifestyle characteristics (physical activity, diet, sleep habits) from the set of descriptive features, which could cause label leakage. In FUUS, the depression outcome was evaluated in a two-stage process. If the result of an initial 4-item screening questionnaire indicated possible presence of MDD (at least two of the four symptoms present), the participants were assessed by a medical provider for the full Diagnostic and Statistical Manual of Mental Disorders Fourth Edition (DSM IV) criteria [22]. In UKB, the depression outcome was defined as an occurrence of a depressive episode (ICD10 code F32 and F33) after the date of assessment, which was drawn from hospital inpatient diagnoses or conditions self-reported.

## 2.2 Predictive models

Two popular types of ML models have been considered: (i) logistic regression (LR) [23], a linear classifier, and (ii) extreme gradient boosting (XGB) [24], a boosting ensemble method. In this paper, we do not focus on the algorithmic aspects of the ML methods considered, but rather on their clinical application and the fairness assessment of their predictions. These techniques work as explained next:

## 2.3 Bias mitigation approaches

Among all the available bias-mitigation techniques, we consider four standard methods in this study: Suppression (SUP), Reweighting (RW) [25], Disparate impact remover (DIR) [26], Calibrated equalized odds post-processing (CPP) [27]. Moreover, we propose a novel post-hoc disparity mitigation named Subgroup-specific score threshold (3ST). Three different pre-processing mitigation techniques (SUP, RW, and DIR) operate over training data, and two different post-processing mitigation techniques (CPP and 3ST) operate over the model's predictions. This is illustrated in Figure 1.

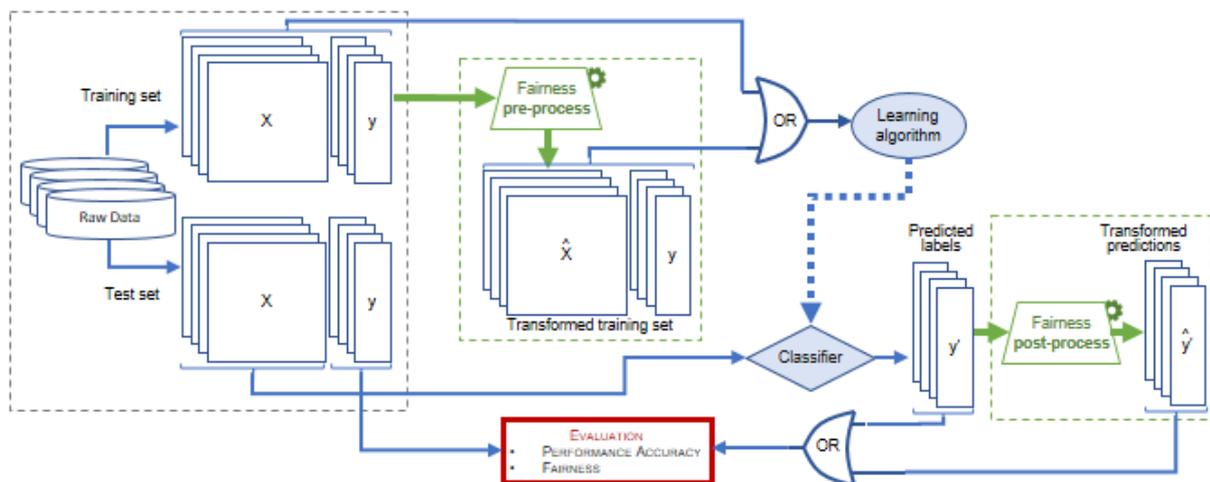


Figure 1. The fairness ML pipeline used in this study.

*Suppression (SUP)*. Protected attributes are directly removed from the training dataset. First, the protected attribute is removed from the dataset. A new ML model is learnt from this new version of the dataset.

*Reweighting (RW)*. This method weights the samples in each (group, label) combination differently to make the protected attribute and outcome statistically independent of each other before model learning. For example, the weight for a sample with positive outcome from a privileged group would be (Number of cases in the privileged group  $\times$  Number of

cases with positive outcome)/(Total sample size × Number of cases of the privileged group with positive outcome). Weights are computed similarly for other samples from the remaining groups and classes. The model is learnt from this new training dataset with weighted samples.

*Disparate impact remover (DIR)*. Given a dataset  $D = (X, Y, C)$ , with protected attribute  $X$ , remaining attributes  $Y$ , and class variable or outcome  $C$ , this repair process attempts to remove bias in the remaining features in  $Y$ . New values are assigned to all the cases and variables in  $Y$ . The new values ensure that all the groups follow the same distribution over every variable, making adjustments based on percentiles and quantile functions. The predictive model is learnt from the new training dataset.

In the three previous bias-mitigation techniques, we act on the training data. Once a newly prepared version of the training dataset is obtained, an ML model is learned from it. Bias mitigation is assessed on the outputs of this model. The following two techniques perform differently. The model is learned from the original data. The outputs of the model are modified under some criteria. Bias mitigation is assessed on the modified outputs of this model.

*Calibrated equalized odds postprocessing (CPP)*. The method calibrates the predicted probability so that the false-positive rate or the false-negative rate of privileged and unprivileged groups are on average equal. It modifies the score outputs of the model for the different subgroups so that the output labels meet an equalized odds objective. In our clinical application, we focus on identifying individuals at risk of depression and desire equitable outcomes; therefore, we consider recall more important than precision, which leads us to set a cost constraint objective: the equal false-negative rates between the subgroups.

*Subgroup-specific score threshold (3ST)*. We propose this novel post-processing bias mitigation method. It is inspired by the idea that, as class balance discrepancies between groups in the training dataset can lead most of the ML models for the classification problem to learn that one group has a higher probability of being part of one class or another, a single default decision threshold of 0.5 for all groups can result in outcome disparities. Therefore, we choose subgroup-specific thresholds for unprivileged groups that are adversely impacted by the unfair bias so that their sensitivity is similar to that of the overall training set. A fixed threshold of 0.5 is used for the remaining groups. This procedure is described in Algorithm 1. Note that [28] uses the name 3ST for their method, where they determine thresholds that balance this metric across groups using a different procedure on the test set.

**Algorithm 1:** Subgroup-specific score threshold (**3ST**) tunes the decision threshold of an unprivileged group towards the objective of equal sensitivity within the overall training set

**Input:** A training set  $D = \{(c_i, p_i, g_i)\}_{i=1}^n$ , where  $c_i$  is the class label,  $p_i$  is the vector of predicted probabilities, and  $g_i$  is the group membership

**Result:** A vector of threshold values  $\theta$ , one per group

1. Initialize threshold vector to  $\theta_g = 0.5$  for all the groups,  $g$
2.  $\text{sensitivity\_overall} :=$  Calculate sensitivity in the overall training set with decision threshold at 0.5
3. For each unprivileged group,  $g$ :
  - a. Consider all the possible thresholds and calculate the sensitivity at each threshold of the unprivileged group  $g$

Find the candidate threshold  $\vartheta$  such that the absolute difference between its associated sensitivity and  $\text{sensitivity\_overall}$  is minimum

$$\text{diff} := \text{abs}(\text{sensitivity}_g, \text{sensitivity\_overall})$$

- b. if  $\text{diff} < \text{mindiff}$ , then adopt threshold  $\vartheta$  for group  $g$ ,  $\theta_g = \vartheta$
4. Use the vector of thresholds  $\theta$  to predict the class of new data

## 2.4 Evaluation metrics

*Performance metrics.* To assess predictive performance, we consider two standard ML performance metrics, namely, the area under the receiver operating characteristic curve (AUC-ROC) that is a plot of false positive rate versus true positive rate, and the balanced accuracy (BAcc), which is the arithmetic mean of sensitivity and specificity, using the presence and absence of depression as the positive and negative class, respectively. When discussing the fairness-accuracy trade-off, we report on empirical accuracy measured by the latter metric because, in practice, classification will be performed at a fixed threshold [9]. However, we report AUC-ROC performance after applying the most effective bias mitigation algorithms in Supplementary Table 2.

*Fairness metrics.* There are three kinds of fairness metrics depending on the fairness concept that they account for: metrics that account for individual fairness, for group fairness, and for both [29]. In this study, we focus exclusively on group fairness metrics for binary classification tasks, which means that subgroups defined by a protected attribute should be received similar model outcome measured by some statistical metrics. In this clinical application, identifying all individuals at risk of depression is desirable; therefore, we

consider sensitivity over specificity which leads us to compare true-positive rates across subgroups. Specifically, we choose the equal opportunity difference (EOD) metric, which states that a binary classifier is fair if its true-positive rates (TPR) are equal across groups (i.e., a value of 0 indicates complete fairness). Let  $D = (X, Y, C)$  be the dataset, with the protected variable  $X$ , regular descriptive variables  $Y$ , and the class variable  $C$ . Predictions provided by a ML model are denoted  $\hat{C}$ . Let us define  $\Omega_X$  as the set of possible values of variable  $X$ . For example, for the protected attribute “sex”,  $\Omega_X = \{\text{male}, \text{female}\}$ . A subgroup of population is formally defined as all the samples in the dataset  $D$  with the same value  $x \in \Omega_X$  assigned to the protected attribute  $X$ . We define the TPR of a specific subgroup  $x \in \Omega_X$  as

$$TPR_x(\hat{C}) := E_Y[\hat{C} \mid C = 1, X = x]$$

Exact equality,  $TPR_{x_0}(\hat{C}) = TPR_{x_1}(\hat{C})$ , for  $x_0, x_1 \in \Omega_X$ , is often hard to verify or enforce in practice. More generally, we use differences among the TPRs of the different subgroups  $x \in \Omega_X$  to measure the level of discrimination  $\Gamma$  [30]:

$$\Gamma^{TPR}(\hat{C}) := \left| TPR_{x_0}(\hat{C}) - TPR_{x_1}(\hat{C}) \right|$$

Specifically, EOD measures the difference of TPR between the unprivileged and privileged groups:

$$EOD = TPR_{x=\text{unprivileged}} - TPR_{x=\text{privileged}}$$

## 2.5 Implementation details

We use Python’s Scikit-learn [31] and XGBoost libraries [24] to implement the LR and the XGB model, correspondingly. RW, DIR, and CPP bias-mitigation techniques are implemented using the AI Fairness 360 (AIF360) library [29]. We perform k-fold cross-validation for model performance measurements and nested cross-validation to carry out hyper-parameter optimization (see details in Supplementary Table 8), with  $k = 5$  for the LONGSCAN, FUUS, NHANES datasets, and  $k = 10$  for the UKB dataset. Cost-sensitive learning is applied to every LR and XGB model through the *class\_weight* and *scale\_pos\_weight* parameters respectively as a means to improve accuracy for positive classes due to the problem of imbalanced classification. We impute the missing data using the statistical imputation algorithm. We use Tukey’s range test [32] to find the 95%-significance level for the true-positive rates for each group over k-fold cross-validation.

### 3 Results

In this section, we study the behavior of the different ML models before and after bias mitigation techniques are applied. Firstly, ML models are learned from data and their predictive performance, and unfair bias, are quantified (Sec. 3.1). Secondly, bias mitigation techniques are applied to these models and, to assess their efficiency, we evaluate the interplay between predictive performance and fairness in the adapted models (Sec. 3.2). We measure performance and fairness metrics of predictive models for depression trained in four datasets: LONGSCAN, FUUS, NHANES, and UKB.

#### 3.1 Initial model performance and bias assessment

We present results with LR models below, whereas the XGB models, qualitatively similar to those of LR models, are available in the Supplement. Regarding predictive performance, LR models achieved lower prediction accuracy in the LONGSCAN dataset and the FUUS dataset. The predictive performance of LR models was considerably better in the NHANES dataset and UKB (see Table 2). These results support that large amounts of data give standard ML models much more predictive power. Rankings of feature importance are available for each predictive model in the Supplement.

Table 2. Predictive performance of LR classifiers learned without bias mitigation from the different datasets

Dataset	BAcc*	AUC-ROC*
LONGSCAN	0.621 [0.577-0.664]	0.644 [0.590-0.699]
FUUS	0.615 [0.598-0.632]	0.671 [0.647-0.695]
NHANES	0.719 [0.711-0.727]	0.793 [0.785-0.800]
UKB	0.729 [0.725-0.734]	0.800 [0.794-0.806]

\*Mean [95%-confidence interval]

Figure 2 displays the fairness metric of the LR models by dataset, protected attribute and subgroup. In each plot, the horizontal shift or difference in performance for different subgroups reveals the bias. Let us analyze the bias assessment by (type of) protected attribute: demographic factors, socioeconomic status, and co-morbidities.

*Protected attributes: demographic factors (sex, ethnicity, nationality).* We show consistent sex differences across datasets in the top row of Figure 2, with higher true positive rates for female subjects. The differences in true positive rates between sex are only statistically

significant at the 95% confidence level for the LONGSCAN, NHANES, and UKB datasets. Note that there are no sex differences in rates of depression among subjects in the FUUS dataset (see Table 1). Interestingly, the mean difference between sexes in the UKB and NHANES datasets (0.1121, 0.167,  $p < 0.001$  respectively) is less than in the LONGSCAN dataset (0.4103,  $p < 0.001$ ). Sex, the top importance feature in LR and XGB for the LONGSCAN dataset, has a much lower rank in the feature importance ranking order in the NHANES and UKB datasets (see the feature importance rankings in the Supplement), which means depression outcome is less sensitive to sex bias in the NHANES and UKB datasets with large sample size and features compared to the LONGSCAN dataset. We highlight the benefit obtained from considering a larger number of risk factors in the predictive model and adequate sample size to reduce bias. This evidence is in line with [30], where enhanced data collection is pointed out as a means to lessen discrimination without sacrificing accuracy. The second row of Figure 2 shows differences in true-positive rates between racial groups, which were not statistically significant, with black subjects and "other/multiracial" subjects having the lowest true-positive rates for the LONGSCAN dataset and the NHANES dataset, respectively; except for the case between black subjects and white subjects in the LONGSCAN dataset, their true-positive rates have nonoverlapping confidence intervals. Interestingly, the UKB subjects in the "do not know/prefer not to answer" ('Missing') group have the lowest true-positive rates, compared with others. As shown in Figure 2, true-positive rates for nationality have non-overlapping confidence intervals. Specifically, foreigner subjects have a higher rate than French subjects. We note that foreigner subjects have a higher observed depression rate in the training set.

*Protected attributes: socioeconomic status (age, income, academic qualifications).* We find that the true-positive rates do not differ much across age groups with many overlapping intervals. However, subjects under 20 years old have the lowest true-positive rate. This may be partially due to the fact that small subset sizes (see Table 1) may not reflect accurate depression rates amongst the adolescents and young adults subpopulation, whose symptoms of depression and other mental illnesses have increased significantly over the last decades [33]. Differences in true-positive rates in the NHANES dataset are also observed between qualification groups. As seen in Figure 2, subjects in the "Level 5" group have the lowest true-positive rates, compared with others in both the NHANES and UKB datasets and subjects in the "Level 0" (Refused/Don't know/Missing) group in the NHANES dataset is also in the unprivileged side. As shown in Figure 2, true-positive rates for income have non-overlapping confidence intervals. Specifically, low-income subjects have a much higher rate than high-income subjects. We note that low-income subjects have a higher baseline observed event rate.

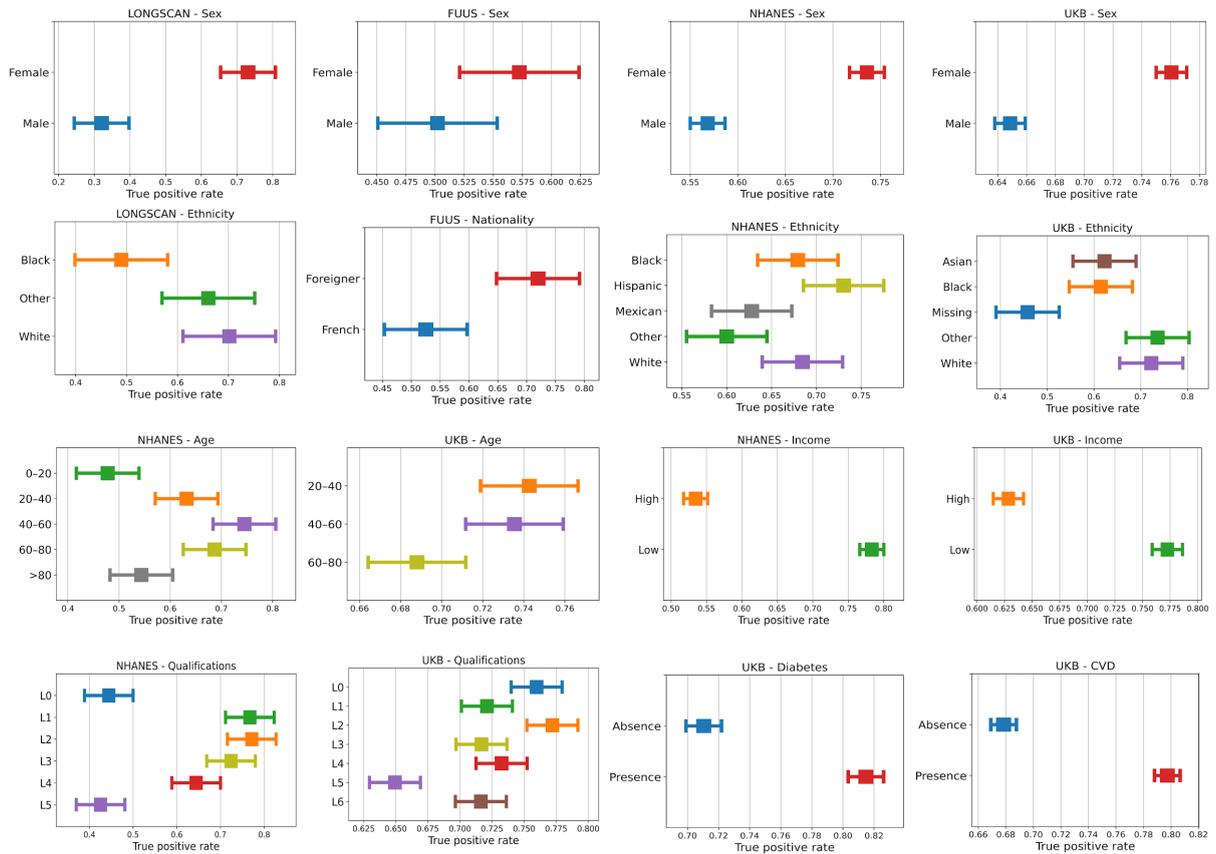


Figure 2. Group-specific true positive rates for LR classifiers learned without bias mitigation from the different datasets. Each plot shows the results per subgroup for a Dataset-Protected attribute pair. Note that the classifiers show regularly biased behaviors due to a lack of representation, varying rates of depression across groups, unequal distribution of features between groups, or a combination of any of these characteristics in the four different study populations.

*Protected attributes: co-morbidities (CVD, diabetes).* As shown in Figure 2, true-positive rates for CVD and diabetes all have non-overlapping confidence intervals. Specifically, individuals experiencing CVD/diabetes have a much higher rate than subjects without CVD/diabetes. These inequitable outcomes support that CVD and diabetes should be considered as an important comorbidity of depression.

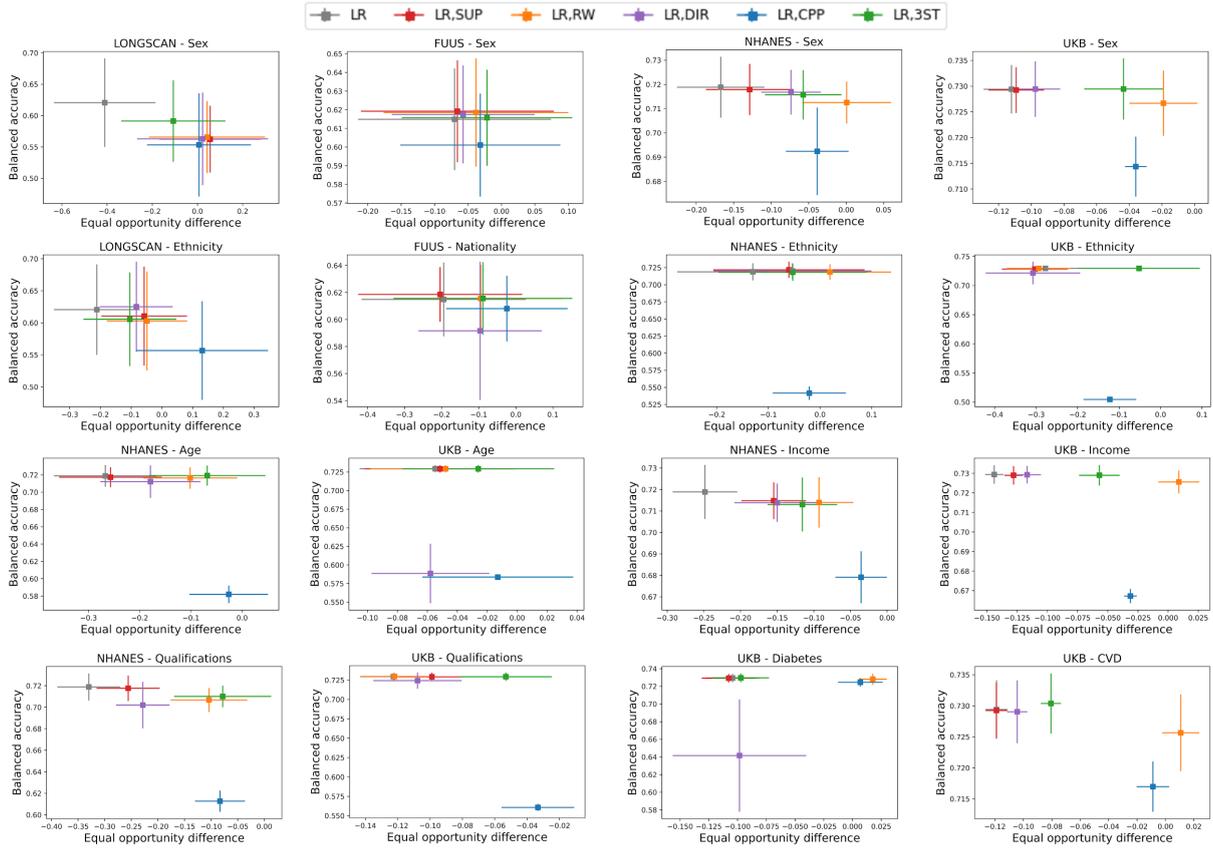


Figure 3. Fairness-accuracy performance in terms of EOD vs. BAcc of the base model and the new classifiers after applying five bias mitigation algorithms to LR classifiers. Each plot shows the results per subgroup for a Dataset-Protected attribute pair.

### 3.2 Model performance after bias mitigation

To assess bias mitigation, we analyze changes in fairness metrics between the previously discussed base models and the new classifiers obtained after bias mitigation. Moreover, given the clinical application, most people would not find it fair to reduce discriminatory outcomes if it identifies fewer actual positives overall. There exists an open discussion in the related literature [25,34] regarding the actual existence of the so-called fairness-accuracy trade-off when bias mitigation is implemented, meaning that predictive performance is reduced if one tries to make the model fairer. A trade-off of accuracy for fairness is often undesirable in healthcare. Thus, we have two objectives: increasing accuracy and non-discrimination. We report the results using the 2D points combining two metrics: (i) a fairness metric, measured by EOD and (ii) a standard ML performance metric, measured by BAcc, on the test set, using 10-fold cross-validation, to determine whether our models for depression prediction are experiencing a fairness-accuracy trade-off. A model can be considered as fair if EOD is between -0.1 and 0.1, its ideal value is 0 and for BAcc, the larger the better. Figure 3 presents the 2D points in the fairness-accuracy space achieved by each

bias mitigation technique. The base model, without bias mitigation, is also shown. The implemented bias methods help to improve fairness, for all protected attribute perspectives. This is a desirable mitigation result. Figure 4 reports the combined best fairness-accuracy results among all bias-mitigation methods for all protected attributes. It is noteworthy that debiasing through RW, DIR, and 3ST substantially improves fairness without compromising model accuracy. On the other hand, the CPP technique preserves precision when calibrating recall, which results in BAcc reduction, according to most of the protected attributes, except for the diabetes attribute. In addition, SUP only partially achieves fairness between groups. This method removes the protected attribute from the training dataset, as they are considered biased features. This result suggests that bias is not only contained in those features but elsewhere. DIR not only excludes them but also adjusts non-protected features that are highly correlated to them. In any case, this repair tool is a good baseline to investigate whether this bias comes from non-protected features.

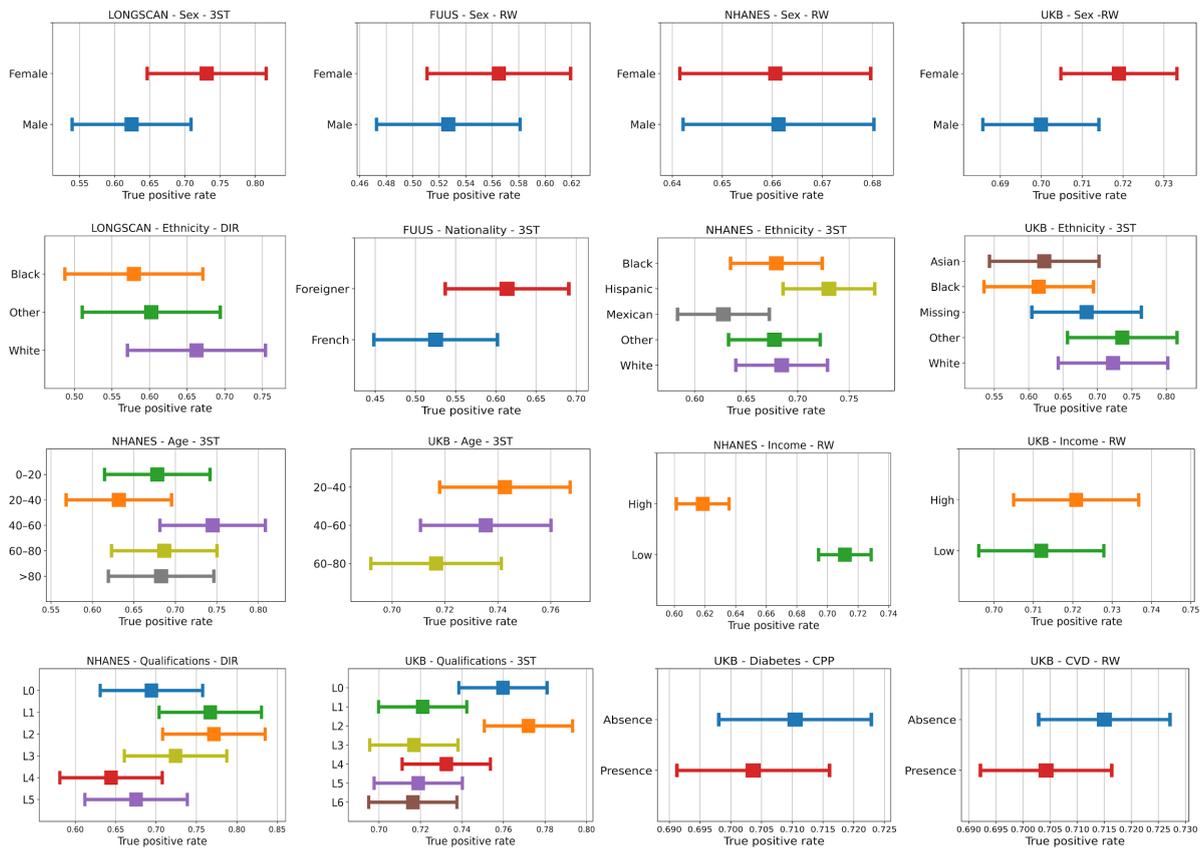


Figure 4. Group differences in true positive rates for an LR classifier after applying bias mitigation algorithms, reporting the best results

Note that true-positive rates for the best mitigation technique (RW) regarding the “income” protected attribute on models learned from the NHANES dataset still have non-overlapping

confidence intervals, although the mean difference between groups was considerably reduced from 0.2485 ( $p < 0.001$ ) to 0.0929 ( $p < 0.001$ ).

#### 4 Discussion

ML algorithms have achieved state-of-the-art performance in many clinical tasks. However, to deploy them in these life-or-death-stakes applications, it must be understood that they can induce biases against unprivileged subgroups and precautionary actions need to be taken in different deployment stages [35]. Here, we leverage our empirical study on four datasets to analyze the need of bias mitigation techniques, as well as their effectiveness, when using ML models to predict mental health issues such as depression.

*Bias found when following the standard ML approach.* Our results indicate that models learned following the standard ML approach show regularly unfair biased behaviors (see Figure 2). We find that, for the classification problem, unequal distribution of classes between groups in the training dataset can lead the predictive model to learn that one group has a higher probability of being part of one class or another, as can be observed in Table 1 and Figure 2. Therefore, ML models trained on the unbalanced dataset of a trial population, and even though the sample in clinical trials is representative of the patient population (not biased sampling), provide potential inequitable outcomes if deployed without integrating bias mitigation techniques. This evidence encourages ML practitioners in healthcare not only to report the model performance on the overall population regardless of the subject membership to subpopulations, but also to audit and address algorithmic bias.

*Bias can be mitigated.* Our results show that the bias mitigation techniques improve fairness compared to the no-intervention base models. All the techniques considered enhance results in terms of the difference in true-positive rates, in different proportions. However, the techniques exhibit differences regarding the effect on the accuracy of the classifiers. Those learned in combination with the RW, DIR, and 3ST mitigation techniques preserve predictive performance, whereas other techniques (SUP, CPP) compromise the level of accuracy, leading them to perform less reliably (see Figure 3). Our findings point out that the RW technique combined with the use of a larger number of risk factors diminishes the impact of the protected attributes (and other possible proxy attributes) on the outcome of the model, which leads to reduced algorithmic bias in NHANES and UKB datasets. This solution is appropriate in our case study as it performs well when it is integrated into predictive model learning while preserving the distribution and values of the original training data, unlike DIR and SUP. In addition, we find that our proposed post-hoc disparity mitigation method (3ST) mitigates bias while preserving predictive performance. However, as other post-doc technical

solutions for imposing fairness, it might be difficult to obtain the optimal threshold for subgroups with small populations underrepresented in the training set. Therefore, RW, DIR, and 3ST are three types of bias-mitigation algorithms that could be tested with clinical predictive models with known differences in disease manifestation in subjects by protected attributes (sex, age, race/ethnicity). We also highlight the importance of having adequate data collection to improve fairness and encourage communities to open health data.

*Fairness-accuracy, a real trade-off?* In fairness literature, the existence of a trade-off between fairness and accuracy is a common assumption, that is, that fairness cannot be improved without sacrificing predictive performance [25]. A few studies also have pointed out that this trade-off may necessitate the application of more complex methods [36-37]. In our case, the standard techniques RW, DIR, and CPP or our proposed method 3ST reduce bias while preserving predictive performance for each specific (dataset, protected attribute) pair. Thus, it is noteworthy that the fairness-accuracy trade-off is not consistently observed in our case studies in depression prediction without requiring complex ML methods. In line with [38], this evidence encourages the ML community to intentionally propose frameworks that maximize both predictive performance enhancement and bias reduction, not accepting an alleged trade-off between them. There is probably no golden bullet, as there is no single best ML model for all prediction problems providing equality of outcomes naturally. ML practitioners need to figure out which combination of type of classifier and bias mitigation algorithm is appropriate for the use case at hand, so that it produces the best results in terms of both accuracy and fairness.

In conclusion, we conduct an empirical study on four exposome datasets to show the ability of bias mitigation techniques to increase fairness of machine learning models obtained to predict depression from environmental and lifestyle data. In addition, our main effort in this work has been directed toward providing empirical evidence to encourage clinical decision-makers to carefully evaluate a proposed framework in terms of both its accuracy and fairness prior to deployment. Experimental results support the idea that it is possible to improve algorithmic fairness without sacrificing their accuracy. We consider that our promising results could enable a wider use of ML techniques in mental healthcare. This should inevitably go hand-in-hand with the assessment of possible biases in the models and the appropriate mitigation techniques, if required.

In the future, we expect to examine simultaneously the effects of having multiple protected attributes such as ethnicity, sex, socio-economic status, geographical location, or co-morbidities (type 1 and type 2 diabetes, or cardiovascular disease), as well as extending

to use protective factors data, such as intelligence, temperament, cognitive appraisal, support from a significant person, which may counteract the negative effects of risk factors for depression, along with environmental and lifestyle data. In addition, we aim to integrate genetic and biological data, which are robust risk factors for depression, in our research.

### **Data availability**

All four datasets are free to download for research purposes from

- LONGSCAN: <https://www.ndacan.acf.hhs.gov/datasets/dataset-details.cfm?ID=170>
- FUUS: <https://datadryad.org/stash/dataset/doi:10.5061/dryad.54qt7>
- NHANES: <https://wwwn.cdc.gov/nchs/nhanes/default.aspx>
- UK Biobank (UKB): <https://www.ukbiobank.ac.uk/enable-your-research/register>

FUUS and NHANES datasets are open access and can be downloaded directly from their links above. LONGSCAN and UKB datasets can be accessed under request to the National Data Archive on Child Abuse and Neglect (NDACAN) and the UK Biobank Access Management Team, respectively. The results for UKB in this study will be returned to the UK Biobank within 6 months since publication, as required.

### **Code availability**

Sample code for all data processing and analysis presented in this work can be found at <https://github.com/ngoc-vien-dang/FairML-Depression>.

### **Declaration of Competing Interest**

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper

### **CRedit authorship contribution statement**

Vien Ngoc Dang: Methodology, Investigation, Data curation, Validation, Software, Visualization, Writing – original draft, Writing – review & editing. Anna Cascarano, Rosa H. Mulder, Charlotte Cecil, and Maria A. Zuluaga: Writing - review & editing. Jerónimo Hernández-González: Conceptualization, Writing - original draft, Writing - review & editing, Co-supervision. Karim Lekadir: Conceptualization, Resources, Writing - review & editing, Supervision.

### **Acknowledgments**

VND, RHM, CC, JHG, and KL have received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement N° 848158, EarlyCause.

## References

- [1] Friedrich, M. J. Depression is the leading cause of disability around the world. *JAMA* **317**, 1517-1517 (2017)
- [2] Bachmann, S. Epidemiology of suicide and the psychiatric perspective. *International Journal of Environmental Research and Public Health* **15**, 1425 (2018)
- [3] Bueno-Notivol, J. et al. Prevalence of depression during the COVID-19 outbreak: A meta-analysis of community-based studies. *International Journal of Clinical and Health Psychology* **21**, 100196 (2021)
- [4] Anttila, V. et al. Analysis of shared heritability in common disorders of the brain. *Science* **360**, eaap8757 (2018)
- [5] Geschwind, D. H. & Flint, J. Genetics and genomics of psychiatric disease. *Science* **349**, 1489-1494 (2015)
- [6] Harald, R. et al. An exposome perspective: early-life events and immune development in a changing world. *The Journal of Allergy and Clinical Immunology* **140**, 24-40 (2017)
- [7] Olesen, J. et al. The economic cost of brain disorders in Europe. *European Journal of Neurology* **19**, 155-162 (2012)
- [8] Chen, I. Y., Szolovits, P., & Ghassemi, M. Can AI help reduce disparities in general medical and mental health care? *AMA journal of ethics* **21**, E167–179 (2019)
- [9] Park, Y. et al. Comparison of Methods to Reduce Bias From Clinical Prediction Models of Postpartum Depression. *JAMA Network Open* **4**, 213909-e213909 (2021)
- [10] Corrigan, P. W., Druss, B. G., & Perlick, D. A. The impact of mental illness stigma on seeking and participating in mental health care. *Psychological Science in the Public Interest* **15**, 37-70 (2014)
- [11] Wong, E. C., Collins, R. L., Cerully, J., Seelam, R., & Roth, B. Differences in mental illness stigma and discrimination among californians experiencing mental health challenges (RAND Corporation, Santa Monica, 2017)
- [12] Albert, R. P. Why is depression more prevalent in women? *Journal of Psychiatry and Neuroscience* **40**, 219--221 (2015)
- [13] Lubin, F., Lusky, A., Chetrit, A., & Dankner, R. Lifestyle and ethnicity play a role in all-cause mortality. *The Journal of Nutrition* **133**, 1180–1185 (2003)
- [14] Runyan, D., Dubowitz, H., English, D., Kotch, J., Litrownik, A., & Thompson, R. Longitudinal studies on child abuse and neglect (LONGSCAN) ages 0-18, Version 1.4 Dataset. National Data Archive on Child Abuse and Neglect (2014). <https://doi.org/10.34681/CXPJ-5M96>
- [15] Tran, A. et al. Health assessment of French university students and risk factors associated with mental health disorders. *PLOS ONE* **12**, e0188187 (2017). <https://doi.org/10.1371/journal.pone.0188187>

- [16] Centers for Disease Control and Prevention (CDC). National Center for Health Statistics (NCHS). National Health and Nutrition Examination Survey Data. Hyattsville, MD: U.S. Department of Health and Human Services, Centers for Disease Control and Prevention. <https://wwwn.cdc.gov/nchs/nhanes/default.aspx> (2005-2018)
- [17] Sudlow, C. et al. UK biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS medicine* **12**, e1001779 (2015). <https://doi.org/10.1371/journal.pmed.1001779>
- [18] Feng, Q., Zhang, Q., Du, Y., Ye, Y., & He, Q. Associations of physical activity, screen time with depression, anxiety and sleep quality among Chinese college freshmen. *PloS One* **9**, e100914 (2014).
- [19] Bayram, N. & Bilgel, N. The prevalence and socio-demographic correlations of depression, anxiety and stress among a group of university students. *Soc Psychiat Epidemiol* **43**, 667–672 (2008)
- [20] Ovuga, E., Boardman, J., & Wasserman, D. Undergraduate student mental health at Makerere University, Uganda. *World Psychiatry* **5**, 51–52 (2006)
- [21] Kroenke, K., Spitzer, R. L., & Williams, J. B. W. The PHQ-9: validity of a brief depression severity measure. *Journal of General Internal Medicine* **16**, 606–613 (2001)
- [22] Nemesure, M.D., Heinz, M.V., Huang, R. et al. Predictive modeling of depression and anxiety using electronic health records and a novel machine learning approach with artificial intelligence. *Scientific Reports* **11**, 1980 (2021)
- [23] Yu, HF., Huang, FL. & Lin, CJ. Dual coordinate descent methods for logistic regression and maximum entropy models. *Mach Learn* **85**, 41–75 (2011)
- [24] Chen, T. & Guestrin, C. XGBoost: A Scalable Tree Boosting System. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794 (2016)
- [25] Calders, T., Kamiran, F., & Pechenizkiy, M. Building classifiers with independency constraints. *IEEE International Conference on Data Mining Workshops*, 13-18 ( 2009)
- [26] Feldman, M., Friedler, S. A., Moeller, J., Scheidegger, C., & Venkatasubramanian, S. Certifying and Removing Disparate Impact. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 259–268 ( 2015)
- [27] Pleiss, G., Raghavan, M., Wu, F., Kleinberg, J., & Weinberger, K. Q. On Fairness and Calibration. *Advances in Neural Information Processing Systems* **30**, (2017)
- [28] Rodolfa K. T. et al. Case study: predictive fairness to reduce misdemeanor recidivism through social service interventions. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 142–153 (2020)

- [29] Bellamy, R. K. E. et al. AI Fairness 360: An Extensible Toolkit for Detecting, Understanding, and Mitigating Unwanted Algorithmic Bias. Preprint at <https://arxiv.org/abs/1810.01943> (2018)
- [30] Chen, I. Y., Johansson, F. D, & Sontag, D. Why is my classifier discriminatory? In Proceedings of the 32nd International Conference on Neural Information Processing Systems, 3543–3554 (2018)
- [31] Pedregosa, F. et al. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* **12**, 2825–2830 (2011)
- [32] Tukey, J. W. Comparing individual means in the analysis of variance. *Biometrics* **5**, 99-114 (1949)
- [33] Twenge, J. M., Cooper, A. B., Joiner, T, Duffy, M., & Binau, S. Age, period, and cohort trends in mood disorder indicators and suicide-related outcomes in a nationally representative dataset, 2005-2017. *Journal of abnormal psychology* **128**, 185-199 (2019).
- [34] Menon, A. K. & Williamson, R. C. The cost of fairness in binary classification. In Proceedings of the 1st Conference on Fairness, Accountability and Transparency **81**, 107–118 (2018)
- [35] de Hond, A.A.H., Leeuwenberg, A.M., Hooft, L. et al. Guidelines and quality criteria for artificial intelligence-based prediction models in healthcare: a scoping review. *npj Digital Medicine* **5**, 2 (2022)
- [36] Friedler, S. A. et al. A comparative study of fairness-enhancing interventions in machine learning. In Proceedings of the Conference on Fairness, Accountability, and Transparency, 329–338 (2019)
- [37] Zafar, M. B., Valera, I., Rogniguez, M. G. & Gummadi, K. P. Fairness constraints: mechanisms for fair classification. In Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, 962–970 (2017).
- [38] Rodolfa, K. T, Lamba, H., & Ghani, R. Empirical observation of negligible fairness–accuracy trade-offs in machine learning for public policy. *Nature Machine Intelligence* **3**, 896–904 (2021)