

# PrivFair: a Library for Privacy-Preserving Fairness Auditing

Sikha Pentyla,<sup>1</sup> David Melanson,<sup>2</sup> Martine De Cock,<sup>2,3</sup> Golnoosh Farnadi<sup>1,4</sup>

<sup>1</sup> Mila – Quebec AI Institute

<sup>2</sup> School of Engineering and Technology, University of Washington Tacoma

<sup>3</sup> Department of Applied Mathematics, Computer Science and Statistics, Ghent University

<sup>4</sup> Department of Decision Sciences, HEC Montréal

## Abstract

Machine learning (ML) has become prominent in applications that directly affect people’s quality of life, including in healthcare, justice, and finance. ML models have been found to exhibit discrimination based on sensitive attributes such as gender, race, or disability. Assessing if an ML model is free of bias remains challenging to date, and by definition has to be done with sensitive user characteristics that are subject of anti-discrimination and data protection law. Existing libraries for fairness auditing of ML models offer no mechanism to protect the privacy of the audit data. We present PRIVFAIR, a library for privacy-preserving fairness audits of ML models. Through the use of Secure Multiparty Computation (MPC), PRIVFAIR protects the confidentiality of the model under audit and the sensitive data used for the audit, hence it supports scenarios in which a proprietary classifier owned by a company is audited using sensitive audit data from an external investigator. We demonstrate the use of PRIVFAIR for group fairness auditing with tabular data or image data, without requiring the investigator to disclose their data to anyone in an unencrypted manner, or the model owner to reveal their model parameters to anyone in plaintext.

## Introduction


Algorithmic decision making, driven by machine learning (ML), has become very prominent in applications that directly affect people’s quality of life, including in healthcare, justice, and finance. ML models have made discriminatory inferences in recidivism prediction (Angwin et al. 2016), credit card approval (Vigdor 2019), advertising (Ali et al. 2019) and job matching (Wall and Schellmann 2021), among others. Concerns over fairness in ML have prompted research into the establishment of fairness metrics and techniques to mitigate bias; see e.g. (Calders and Verwer 2010; Galhotra, Brun, and Meliou 2017; Kearns et al. 2018; Pleiss et al. 2017; Verma and Rubin 2018) and references therein. In state-of-the-art approaches based on *group fairness measures*, audits are done by comparing the ML model predictions for different demographic subgroups in the audit data set. For instance, a classifier satisfies the definition of *demographic parity* (Dwork et al. 2012) if the subjects in the protected group and the unprotected group have equal probability of being assigned to the positive predicted class, e.g. if credit card approval is equally probable for both females and males.

In real world scenarios, a proprietary model  $\mathcal{M}$  held by a company *Alice* may need to be audited by an external investigator *Bob* using sensitive audit data  $\mathcal{D}$  (see Figure 1a). For example, a bank or hospital, represented by Bob, wants to purchase use of a predictive model and needs to investigate whether the model performs fairly on their data. Alice does not want to disclose her trained model parameters as this could assist rival companies to benefit from her technology. Furthermore, ML models can memorize specific examples from the training data (Carlini et al. 2019), hence disclosing  $\mathcal{M}$ , or even giving just black box access through an API interface, can leak very specific information about the training data, which might be sensitive in itself. Likewise, Bob does not want to disclose the audit data  $\mathcal{D}$  to Alice, because it contains sensitive attributes that are on one hand needed for the fairness audit, while on the other hand may be subject of anti-discrimination and data protection law.

While substantial research progress has been made in the area of fairness in AI and, separately, in the area of privacy-preserving ML, there is a gap in literature that addresses fairness and privacy simultaneously. Existing algorithms for fairness auditing, and their implementations in libraries such as Fairlearn (Bird et al. 2020) and AI Fairness 360 (Bellamy et al. 2019), operate without concern for privacy, assuming that the entity performing the fairness audit has unrestricted access to the model  $\mathcal{M}$  as well as the audit data  $\mathcal{D}$ . This assumes that either Alice is willing to disclose  $\mathcal{M}$  to Bob, or that Bob is willing to disclose  $\mathcal{D}$  to Alice, or that they are both willing to disclose  $\mathcal{M}$  and  $\mathcal{D}$  to some trusted third party. None of these scenarios is realistic for proprietary models and sensitive audit data.

We present PRIVFAIR, a library for privacy-preserving evaluation of ML models that does not require disclosure of trained model parameters nor of audit data. PRIVFAIR is based on Secure Multiparty Computation (MPC) (Cramer, Damgard, and Nielsen 2015), a cryptographic approach that allows two or more parties to jointly compute a specified output from their private information in a distributed fashion, without actually revealing the private information.

The price paid to preserve privacy during fairness auditing is an increase in runtime, which stems mostly from the computational and communication cost of the MPC protocols used for labeling the audit instances in a privacy-preserving manner. We argue that this is a reasonable price to pay when working with sensitive data. Indeed, unlike in inference ap-



```

Encrypting model parameters...
Connecting to servers...
Connected to servers...
Secret sharing parameters with servers...

Servers running MPC protocol for DP
Servers running MPC protocol for EOP

Finished protocol execution

```

(b) Model owner’s terminal during audit

```

Encrypting audit data
Connecting to servers...
Connected to servers...
Secret sharing audit data with servers...

Servers running MPC protocol for DP
Servers running MPC protocol for EOP

Finished protocol execution
Aggregating results...

Demographic parity - male: 0.24
Demographic parity - female: 0.19
Equal opportunity - male: 0.399
Equal opportunity - female: 0.40

```

(c) Investigator’s terminal during audit

Figure 1: Alice and Bob send encrypted shares of the model  $\mathcal{M}$  (Alice) and the audit data  $\mathcal{D}$  (Bob) to 3 servers. The servers subsequently execute MPC protocols to measure demographic parity (DP) and equal opportunity (EOP) in a privacy-preserving manner, i.e. through computations over the encrypted data.

plications – such as real-time speech recognition or video classification for surveillance – where fast responsiveness of the system is of utmost importance, in fairness auditing use cases with sensitive data one can typically afford longer runtimes, justifying the use of the most stringent security settings, even if they come at a higher cost.

### Preliminaries

As is common in MPC, all computations in PRIVFAIR are done over integers in a ring  $\mathbb{Z}_q = \{0, 1, 2, \dots, q - 1\}$ . The parameter values of Alice’s trained model  $\mathcal{M}$ , and the feature values of Bob’s audit data  $\mathcal{D}$ , are natively often real numbers represented in a floating point format. In this case, as is common in MPC, Alice and Bob convert all data in their inputs  $\mathcal{M}$  and  $\mathcal{D}$  into integers using a bijection between the reals (represented by fixed point notation with  $k$  bits in total out of which  $a$  bits of decimal precision), and the integers in  $\mathbb{Z}_q$ , making it simple to go from one representation to the other (De Cock et al. 2021).  $q = 2^k$  is chosen large enough so that there is no loss of precision that would affect the utility of the ML models. Next, Alice and Bob encrypt the integers by splitting them into so-called secret shares (see below) and distributing these secret shares across parties (servers) who subsequently perform computations over the shares (see Figure 1). While the original values

of the inputs  $\mathcal{M}$  and  $\mathcal{D}$  can be trivially revealed by combining all shares, the secret-sharing based MPC schemes ensure that nothing about the inputs is revealed to any subset of the servers that can be corrupted by an adversary. This means, in particular, that none of the servers by themselves learns anything about the actual values of the inputs.

PRIVFAIR incorporates a variety of MPC schemes, designed for different numbers of parties (servers) and offering various levels of security that correspond to different threat models. Regarding threat models, the fairness auditing protocols in PRIVFAIR provide protection against *passive* as well as *active* adversaries. While parties corrupted by passive adversaries follow the protocol instructions correctly but try to obtain additional information, parties corrupted by active adversaries can deviate from the protocol instructions. Regarding the number of parties, while fairness auditing is inherently a two-party (2PC) scenario between Alice and Bob, they may not each have a server available on their end and prefer to outsource computations to the cloud. PRIVFAIR accommodates two scenarios:

**Scenario 1** Alice and Bob have the required computational setup to perform MPC computations themselves and proceed by secret-sharing the data with each other only (2PC),

**Scenario 2** Alice and Bob need to perform computation in-

tensive tasks and outsource the MPC based computations, thus secret-sharing the data across a set of untrusted servers (see Figure 1a for an illustration of 3PC).

The 2PC setting is a *dishonest-majority* setting in which each server can only trust itself. In the 3PC setting, we consider an *honest-majority* setting where at most one of the servers is corrupted by an adversary.

The fairness auditing protocols in PRIVFAIR are sufficiently generic to be used in dishonest-majority as well as honest-majority settings, with passive or active adversaries. This is achieved by changing the underlying MPC scheme to align with the desired security setting. We refer to the caption of Table 2 for a listing of the various MPC schemes that we used to obtain the results reported in this paper.

In a replicated secret sharing scheme with 3 servers (3PC) and passive adversaries for example (Araki et al. 2016), a value  $x$  in  $\mathbb{Z}_q$  is secret shared among the servers (parties)  $S_1, S_2,$  and  $S_3$  by picking uniformly random numbers (shares)  $x_1, x_2, x_3 \in \mathbb{Z}_q$  such that  $x_1 + x_2 + x_3 = x \pmod q$ , and distributing  $(x_1, x_2)$  to  $S_1$ ,  $(x_2, x_3)$  to  $S_2$ , and  $(x_3, x_1)$  to  $S_3$ . Note that, while the secret-shared information can be trivially revealed by combining all shares, no single server can obtain any information about  $x$  given its shares. In the remainder of this paper, we use  $\llbracket x \rrbracket$  as a shorthand for a secret sharing of  $x$ , regardless of which secret-sharing scheme is used. In the replicated secret-sharing scheme sketched above,  $\llbracket x \rrbracket = ((x_1, x_2), (x_2, x_3), (x_3, x_1))$ .

The MPC schemes used in PRIVFAIR provide a mechanism for the servers to perform primitive operations over the secret shares, namely addition of a constant, multiplication by a constant, and addition and multiplication of secret-shared values. Building on these cryptographic primitives, MPC protocols for other operations have been developed in the literature. As we explain below, PRIVFAIR’s MPC protocols for fairness auditing rely on established MPC subprotocols for multiplication  $\pi_{\text{MUL}}$ , division  $\pi_{\text{DIV}}$ , and equality testing  $\pi_{\text{EQ}}$  of secret-shared values; see e.g. (Catrina and Saxena 2010; Dalskov, Escudero, and Keller 2020, 2021) and references therein for descriptions of these subprotocols.

## Secure Fairness Auditing Protocols

Below we present MPC protocols for fairness auditing of ML models. By fairness, we mean algorithmic fairness as commonly understood in the ML literature, namely that a person’s experience with an information system should not irrelevantly depend on their personal characteristics such as race, gender, sexual orientation, ethnicity, religion, or age (Ekstrand, Joshaghani, and Mehrpouyan 2018). Fairness has been operationalized through various measures of group fairness, expressing that different groups of users should receive similar statistical treatment. Many different notions of group fairness have been proposed in the literature, some of them even mathematically incompatible (Kleinberg, Mullainathan, and Raghavan 2016); the discussion on the most appropriate metrics is still ongoing and out of scope for this paper.

Many well adopted statistical notions of fairness rely on measuring the number of true positives (TP), false positives (FP), true negatives (TN), and false negatives (FN) when ap-

plying the ML model under investigation to an audit data set. Privacy-preserving computations of these numbers form the basic building blocks of our proposed protocols. In the description of the protocols below, the audit data  $\mathcal{D}$  contains samples of the form  $(\mathbf{x}, y, a)$  in which  $\mathbf{x}$  is the instance that needs to be classified (e.g. an image, or a row of tabular data),  $y$  is the ground truth class label, and  $a$  indicates whether the instance belongs to an unprivileged demographic subgroup.  $a$  is a value of a binary variable  $A$  where  $A = 0$  designates the unprotected group and  $A = 1$  designates the protected (or sensitive) group.  $y$  is a value of a variable  $Y$  that represents the actual outcome. Similarly, we use  $\hat{Y}$  to denote a variable that represents the predicted outcome according to the model  $\mathcal{M}$ . For ease of explanation, we first assume that  $Y$  and  $\hat{Y}$  are binary, i.e. that we are solving a binary classification task. Next, we explain how the protocols can be used for auditing of multi-class classifiers as well.

**Protocol  $\pi_{\text{EOD}}$  for Equalized Odds** By definition, equalized odds is satisfied if  $\hat{Y}$  and  $A$  are independent conditional on  $Y$  (Eq. (1)) (Hardt, Price, and Srebro 2016), i.e.

$$\hat{Y} \perp\!\!\!\perp A \mid Y \quad (1)$$

For a binary classifier, this is equivalent to Eq. (2), implying that the true positive rate  $\text{TPR} = \text{TP}/(\text{TP}+\text{FN})$  and the false positive rate  $\text{FPR} = \text{FP}/(\text{FP}+\text{TN})$  are the same for both the unprotected and the protected groups. An exact equality as in Eq. (2) may be hard to achieve, so it is common to compute the left and the right hand sides of Eq. (2) separately and inspect their differences, sometimes reported through metrics such as the equalized odds difference.

$$\begin{aligned} \text{P}(\hat{Y} = 1 \mid Y = y, A = 0) &= \text{P}(\hat{Y} = 1 \mid Y = y, A = 1), \\ \forall y \in \{0, 1\} \end{aligned} \quad (2)$$

We extend Eq. (2) to multi-class classification using the *one-vs-rest* approach. For a classification task with a label set  $\mathcal{L} = \{0, 1, \dots, C - 1\}$ , for each class label  $c \in \mathcal{L}$ , we define the equalized odds for class label  $c$  as per Eq. (3). In this case, for each class label  $c$ , we consider two cases with  $\hat{Y} = c$  and  $\hat{Y} = \neg c$  (any other class). Adapting Eq. (2) to this gives us Eq. (3) which implies that the TPR and FPR for each class should be same for both the protected and unprotected groups.

$$\begin{aligned} \text{P}(\hat{Y} = c \mid Y = y, A = 0) &= \text{P}(\hat{Y} = c \mid Y = y, A = 1), \\ \forall y \in \{c, \neg c\}, \forall c \in \mathcal{L} \end{aligned} \quad (3)$$

In PRIVFAIR, the left and right hand sides of Eq. (3) are computed in a privacy-preserving manner with protocol  $\pi_{\text{EOD}}$ , presented in pseudocode here as Protocol 1. At the beginning of this protocol, the servers (parties) have secret shares of the parameters of the model  $\mathcal{M}$ , as received from the model owner Alice (see Figure 1). Similarly, from the investigator Bob, the parties have received secret shares of an audit data set  $\mathcal{D}$  with  $N$  instances, including a secret-shared vector  $Y$  of length  $N$  with the ground truth labels and a secret-shared vector  $A$  of length  $N$  with the sensitive attribute. Note that we abuse the notations  $Y$  and  $A$  here to denote random variables as in Eq. (3) as well as vectors with values for those random variables as in Protocol 1.

Protocol 1: Protocol  $\pi_{\text{EOD}}$  for computing equalized odds for multi-class classification

**Input:** The parties have a secret sharing of trained model parameters  $\llbracket \mathcal{M} \rrbracket$ , and a secret sharing of a data set  $\llbracket \mathcal{D} \rrbracket$  with  $N$  instances and secret sharings  $\llbracket Y \rrbracket$  and  $\llbracket A \rrbracket$  of the corresponding ground truth labels (drawn from a set  $\mathcal{L}$  with  $C$  labels) and a binary sensitive attribute.

**Output:** A secret sharing of the equalized odds metrics for each class

```

1:  $\llbracket Y_{\text{pred}} \rrbracket \leftarrow \pi_{\text{INFER}}(\llbracket \mathcal{M} \rrbracket, \llbracket \mathcal{D} \rrbracket)$ 
2: for  $c = 0$  to  $C - 1$  do
3:   for  $i = 1$  to  $N$  do
4:      $\llbracket \text{grnd} \rrbracket \leftarrow \pi_{\text{EQ}}(\llbracket Y[i] \rrbracket, c)$ 
5:      $\llbracket \text{pred} \rrbracket \leftarrow \pi_{\text{EQ}}(\llbracket Y_{\text{pred}}[i] \rrbracket, c)$ 
6:      $\llbracket a \rrbracket \leftarrow \llbracket A[i] \rrbracket$ 
7:      $\llbracket \text{tp} \rrbracket \leftarrow \pi_{\text{MUL}}(\llbracket \text{grnd} \rrbracket, \llbracket \text{pred} \rrbracket)$ 
8:      $\llbracket \text{ta} \rrbracket \leftarrow \pi_{\text{MUL}}(\llbracket \text{grnd} \rrbracket, \llbracket a \rrbracket)$ 
9:      $\llbracket \text{pa} \rrbracket \leftarrow \pi_{\text{MUL}}(\llbracket \text{pred} \rrbracket, \llbracket a \rrbracket)$ 
10:     $\llbracket \text{tpa} \rrbracket \leftarrow \pi_{\text{MUL}}(\llbracket \text{tp} \rrbracket, \llbracket a \rrbracket)$ 
11:     $\llbracket \text{TP}_{A=1}[c] \rrbracket \leftarrow \llbracket \text{TP}_{A=1}[c] \rrbracket + \llbracket \text{tpa} \rrbracket$ 
12:     $\llbracket \text{FN}_{A=1}[c] \rrbracket \leftarrow \llbracket \text{FN}_{A=1}[c] \rrbracket + (\llbracket \text{ta} \rrbracket - \llbracket \text{tpa} \rrbracket)$ 
13:     $\llbracket \text{FP}_{A=1}[c] \rrbracket \leftarrow \llbracket \text{FP}_{A=1}[c] \rrbracket + (\llbracket \text{pa} \rrbracket - \llbracket \text{tpa} \rrbracket)$ 
14:     $\llbracket \text{TN}_{A=1}[c] \rrbracket \leftarrow \llbracket \text{TN}_{A=1}[c] \rrbracket$ 
         $+ (\llbracket a \rrbracket - \llbracket \text{ta} \rrbracket - \llbracket \text{pa} \rrbracket + \llbracket \text{tpa} \rrbracket)$ 
15:     $\llbracket \text{TP}_{A=0}[c] \rrbracket \leftarrow \llbracket \text{TP}_{A=0}[c] \rrbracket + (\llbracket \text{tp} \rrbracket - \llbracket \text{tpa} \rrbracket)$ 
16:     $\llbracket \text{FN}_{A=0}[c] \rrbracket \leftarrow \llbracket \text{FN}_{A=0}[c] \rrbracket$ 
         $+ (\llbracket \text{grnd} \rrbracket - \llbracket \text{ta} \rrbracket - \llbracket \text{tp} \rrbracket + \llbracket \text{tpa} \rrbracket)$ 
17:     $\llbracket \text{FP}_{A=0}[c] \rrbracket \leftarrow \llbracket \text{FP}_{A=1}[c] \rrbracket$ 
         $+ (\llbracket \text{pred} \rrbracket - \llbracket \text{pa} \rrbracket - \llbracket \text{tp} \rrbracket + \llbracket \text{tpa} \rrbracket)$ 
18:     $\llbracket \text{TN}_{A=0}[c] \rrbracket \leftarrow \llbracket \text{TN}_{A=1}[c] \rrbracket$ 
         $+ (1 - \llbracket \text{grnd} \rrbracket - \llbracket \text{pred} \rrbracket - \llbracket a \rrbracket$ 
         $+ \llbracket \text{tp} \rrbracket + \llbracket \text{ta} \rrbracket + \llbracket \text{pa} \rrbracket - \llbracket \text{tpa} \rrbracket)$ 
19:   end for
20:    $\llbracket \text{TPR}_{A=1}[c] \rrbracket \leftarrow \pi_{\text{DIV}}(\llbracket \text{TP}_{A=1}[c] \rrbracket,$ 
         $\llbracket \text{TP}_{A=1}[c] \rrbracket + \llbracket \text{FN}_{A=1}[c] \rrbracket)$ 
21:    $\llbracket \text{TPR}_{A=0}[c] \rrbracket \leftarrow \pi_{\text{DIV}}(\llbracket \text{TP}_{A=0}[c] \rrbracket,$ 
         $\llbracket \text{TP}_{A=0}[c] \rrbracket + \llbracket \text{FN}_{A=0}[c] \rrbracket)$ 
22:    $\llbracket \text{FPR}_{A=1}[c] \rrbracket \leftarrow \pi_{\text{DIV}}(\llbracket \text{FP}_{A=1}[c] \rrbracket,$ 
         $\llbracket \text{FP}_{A=1}[c] \rrbracket + \llbracket \text{TN}_{A=1}[c] \rrbracket)$ 
23:    $\llbracket \text{FPR}_{A=0}[c] \rrbracket \leftarrow \pi_{\text{DIV}}(\llbracket \text{FP}_{A=0}[c] \rrbracket,$ 
         $\llbracket \text{FP}_{A=0}[c] \rrbracket + \llbracket \text{TN}_{A=0}[c] \rrbracket)$ 
24: end for
25: return  $\llbracket \text{TPR}_{A|0,1} \rrbracket, \llbracket \text{FPR}_{A|0,1} \rrbracket$ 

```

On Line 1, the servers perform privacy-preserving labeling of the audit data instances in  $\mathcal{D}$  with the model  $\mathcal{M}$  using an MPC protocol  $\pi_{\text{INFER}}$  for secure inference. Such protocols have been developed by us and others for logistic regression, neural networks, decision tree ensembles etc. (De Cock et al. 2019; Fritchman et al. 2018; Agrawal et al. 2019; Reich et al. 2019; Dalskov, Escudero, and Keller 2020; Pentylala, Dowsley, and De Cock 2021) and can be used in combination with the protocols in PRIVFAIR. After Line 1, the servers have a secret-shared vector  $\llbracket Y_{\text{pred}} \rrbracket$  with predicted class labels for each of the samples in the audit data.

Next, on Line 2–19, for each of the classes, the servers compute the number of true positives (TP), number of false

grnd	pred	Metric	Logic
1	1	TP	$\text{grnd} * \text{pred}$
1	0	FN	$\text{grnd} * (1 - \text{pred})$
0	1	FP	$(1 - \text{grnd}) * \text{pred}$
0	0	TN	$(1 - \text{grnd}) * (1 - \text{pred})$

Table 1: Logic for evaluating TP, FN, FP and TN

positives (FP), number of true negatives (TN) and false negatives (FN) for the (un)protected group. To this end, on Line 4, the servers compute a secret-shared binary variable  $\llbracket \text{grnd} \rrbracket$  representing if instance  $i$  belongs to the class  $c$  that is being inspected. To obtain this, the servers run an MPC protocol  $\pi_{\text{EQ}}$  for equality testing that takes as input the secret-shared ground truth class label  $\llbracket Y[i] \rrbracket$  and the class  $c$  itself, and returns a secret-sharing of 1 if the equality test was positive, and a secret-sharing of 0 otherwise. Next, on Line 5, the servers compute in a similar way a secret-shared binary variable  $\llbracket \text{pred} \rrbracket$  representing if the predicted outcome for instance  $i$  is the class  $c$  that is being inspected. On Line 6, the servers compute a secret-shared binary value  $\llbracket a \rrbracket$  representing if the current sample in the audit data belongs to the protected or unprotected group.

The underlying logic for efficiently computing TP, FN, FP, and TN, is derived from the truth table shown in Table 1. Multiplying column 4 in Table 1 with  $a$  and with  $1 - a$  respectively gives us the contribution of instance  $i$  to the TP, FN, FP, and TN metrics for the protected and unprotected group respectively. For example, the total number of instances of the protected group that get incorrectly classified as belonging to class  $c$  is

$$\text{FP}_{A=1}[c] = \sum_{i=1}^n (1 - \text{grnd}^{(i)}) * \text{pred}^{(i)} * a^{(i)} \quad (4)$$

in which  $\text{grnd}^{(i)}$ ,  $\text{pred}^{(i)}$ , and  $a^{(i)}$  are the secret-shared binary variables computed for instance  $i$  on Line 4–5.

This logic is modified to suit the MPC computations on Line 7–18. To this end, we rewrote the expressions in the summations as in the right hand side of Eq. (4) to reduce the overall number of multiplications as much as possible. As a result,  $\pi_{\text{EOD}}$  requires only 4 multiplications per instance and per class (see Lines 7–10). On Lines 11–14 and Lines 15–18, the servers use these precomputed secret-shared products to obtain secret shares of TP, FN, FP, and TN respectively for the protected group and for the unprotected group.

Finally, on Line 20–23 the servers execute a protocol  $\pi_{\text{DIV}}$  for division with secret-shared values, to obtain secret shares of the TPR and FPR for each class for the protected and unprotected groups. At the end of the protocol, each server sends its secret shares to the investigator Bob who can then combine the shares to learn the TPR and FPR values. The servers themselves do not learn anything about the real values of the TPR and the FPR, nor of the TP, FN, FP, and TN metrics computed along the way, and not even the number of instances of each class in the audit data.

While Protocol 1 is presented for multi-class classification, it can be trivially used for binary classification as well by removing Line 2 and 19, and running the inner for-loop

---

Protocol 2: Protocol  $\pi_{\text{GACC}}$  for computing sub-group accuracy for multi-class classification

---

**Input:** The parties have a secret sharing of trained model parameters  $\llbracket \mathcal{M} \rrbracket$ , and a secret sharing of a data set  $\llbracket \mathcal{D} \rrbracket$  with  $N$  instances and secret sharings  $\llbracket Y \rrbracket$  and  $\llbracket A \rrbracket$  of the corresponding ground truth labels (drawn from a set  $\mathcal{L}$  with  $C$  labels) and a binary sensitive attribute.

**Output:** A secret sharing of the sub-group accuracy metrics

```

1:  $\llbracket Y_{\text{pred}} \rrbracket \leftarrow \pi_{\text{INFER}}(\llbracket \mathcal{M} \rrbracket, \llbracket \mathcal{D} \rrbracket)$ 
2: for  $i = 1$  to  $N$  do
3:    $\llbracket \text{count}_{A=1} \rrbracket \leftarrow \llbracket \text{count}_{A=1} \rrbracket + \llbracket A[i] \rrbracket$ 
4: end for
5:  $\llbracket \text{count}_{A=0} \rrbracket \leftarrow N - \llbracket \text{count}_{A=1} \rrbracket$ 
6:  $\llbracket \text{correct}_{A=1} \rrbracket \leftarrow 0$ 
7:  $\llbracket \text{correct}_{A=0} \rrbracket \leftarrow 0$ 
8: for  $i = 1$  to  $N$  do
9:    $\llbracket \text{corr} \rrbracket \leftarrow \llbracket Y[i] \rrbracket - \llbracket Y_{\text{pred}}[i] \rrbracket$ 
10:   $\llbracket \text{iscorr} \rrbracket \leftarrow \pi_{\text{EQ}}(\llbracket \text{corr} \rrbracket, 0)$ 
11:   $\llbracket \text{iscorr}_{A=1} \rrbracket \leftarrow \pi_{\text{MUL}}(\llbracket \text{iscorr} \rrbracket, \llbracket A[i] \rrbracket)$ 
12:   $\llbracket \text{correct}_{A=1} \rrbracket \leftarrow \llbracket \text{correct}_{A=1} \rrbracket + \llbracket \text{iscorr}_{A=1} \rrbracket$ 
13:   $\llbracket \text{correct}_{A=0} \rrbracket \leftarrow \llbracket \text{correct}_{A=0} \rrbracket$ 
    +  $(\llbracket \text{iscorr} \rrbracket - \llbracket \text{iscorr}_{A=1} \rrbracket)$ 
14: end for
15:  $\llbracket \text{ACC}_{A=1} \rrbracket \leftarrow \pi_{\text{DIV}}(\llbracket \text{correct}_{A=1} \rrbracket, \llbracket \text{count}_{A=1} \rrbracket)$ 
16:  $\llbracket \text{ACC}_{A=0} \rrbracket \leftarrow \pi_{\text{DIV}}(\llbracket \text{correct}_{A=0} \rrbracket, \llbracket \text{count}_{A=0} \rrbracket)$ 
17:  $\llbracket \text{ACC} \rrbracket \leftarrow \pi_{\text{MUL}}(\llbracket \text{correct}_{A=1} \rrbracket + \llbracket \text{correct}_{A=0} \rrbracket, 1/N)$ 
18: return  $\llbracket \text{ACC}_{A \in \{0,1\}} \rrbracket, \llbracket \text{ACC} \rrbracket$ 

```

---

only once with a fixed value of  $c = 1$ .

**Protocol  $\pi_{\text{EOP}}$  for Equal Opportunity** By definition, equal opportunity for a binary predictor is satisfied if (Hardt, Price, and Srebro 2016)

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

implying that the TPR is the same for both the unprotected and the protected group. This fairness notion is a relaxation of equalized odds in a sense that it focuses only the positive or advantaged outcome.

Protocol 1 can be modified directly to compute Eq. (5) resulting in a new protocol  $\pi_{\text{EOP}}$ . This can be done by running the inner for-loop only once, with a fixed value of  $c = 1$  (binary classification) and by removing several lines, i.e. Line 9; the lines that compute FP and TN (Lines 13–14, 17–18); and the lines that compute FPR (Lines 22–23). Through executing the resulting protocol  $\pi_{\text{EOP}}$ , the servers will only compute secret shares of TPR (Line 25) for  $c = 1$ .

**Protocol  $\pi_{\text{GACC}}$  for Sub-Group Accuracy** This notion of fairness, similar to overall accuracy equality (Berk et al. 2021), is satisfied if the classifier is equally accurate for both the protected and unprotected group.

To report sub-group accuracy, PRIVFAIR’s protocol  $\pi_{\text{GACC}}$  (see Protocol 2) computes the accuracy for the protected group, the accuracy for the unprotected group, and the overall accuracy. These notions are defined as follows, for a classifier  $\mathcal{M}$  and an audit data set  $\mathcal{D}$  with instances of the

form  $(\mathbf{x}, y, a)$ :

$$\text{ACC}_{A=1} = \frac{|\{(\mathbf{x}, y, 1) \in \mathcal{D} \mid y = \mathcal{M}(\mathbf{x})\}|}{|\{(\mathbf{x}, y, 1) \in \mathcal{D}\}|} \quad (6)$$

$$\text{ACC}_{A=0} = \frac{|\{(\mathbf{x}, y, 0) \in \mathcal{D} \mid y = \mathcal{M}(\mathbf{x})\}|}{|\{(\mathbf{x}, y, 0) \in \mathcal{D}\}|} \quad (7)$$

$$\text{ACC} = \frac{|\{(\mathbf{x}, y, a) \in \mathcal{D} \mid y = \mathcal{M}(\mathbf{x})\}|}{|\{(\mathbf{x}, y, a) \in \mathcal{D}\}|} \quad (8)$$

Eq. (6)–(8) cover both binary and multi-class classifiers.

On Line 1 in Protocol 2, the servers perform privacy-preserving labeling of the audit data  $\mathcal{D}$  with the model  $\mathcal{M}$ . On Line 2-4, the servers “count” how many instances in the audit data belong to the protected group, i.e. the denominator of Eq. (6). Note that each of the servers obtains only a secret share of the resulting  $\text{count}_{A=1}$ , i.e. no server finds out how many instances in the audit data belong to the protected class. Since the total number of instances  $N$  is public information, the servers can straightforwardly retrieve a secret-sharing of the number of instances belonging to the unprotected group by computing  $N - \llbracket \text{count}_{A=1} \rrbracket$ . This is leveraged in Line 5.

Through Lines 8–14, the servers count how many predictions are correct. To this end, for each instance, on Line 9 the servers compute the difference between the actual outcome and predicted outcome. If this difference is equal to zero as per Line 10, it means it was a correct prediction. Lines 11–13 decide whether this correct prediction is to be accounted for in the total of the protected or the unprotected group. The right hand side of Line 13 for the unprotected group is written to take advantage of the multiplication that was already done for the protected group on Line 11.

On Line 15–16 the servers compute accuracy for the protected and unprotected group respectively. Line 17 computes overall accuracy by performing secure addition of number of correct instances for each subgroup followed by secure multiplication with a constant  $1/N$ .

**Protocol  $\pi_{\text{DP}}$  for Demographic Parity** Demographic parity, also known as statistical parity, is one of the best known and most widely accepted notions of group fairness. A classifier satisfies demographic parity if (Dwork et al. 2012)

$$P(\hat{Y} = 1 \mid A = 0) = P(\hat{Y} = 1 \mid A = 1) \quad (9)$$

implying that both protected and unprotected groups have equal probability of receiving positive outcomes (Verma and Rubin 2018). This is equivalent to the ratio of the number of positive outcomes for the group to the total number of instances belonging to the group. Note that the number of positive outcomes for a group in the audit data is the sum of the number of true positives and the number of false positives for that group, an idea that is leveraged in PRIVFAIR’s MPC protocol  $\pi_{\text{DP}}$  for demographic parity.

Protocol  $\pi_{\text{DP}}$  can be derived by replacing Line 2 in Protocol 1 with  $c = 1$  (for binary classification) and retaining only the lines required for computing TP and FP for each subgroup (Lines 3–10, 11, 13, 15, 17). Furthermore, Lines 20–23 are replaced by lines that compute the sum of TP and FP:

Task	#Samples (owned by Bob)	#Model Parameters (owned by Alice)	Passive		Active		
			DP ( $\pi_{DP}$ )	EOP ( $\pi_{EOP}$ )	DP ( $\pi_{DP}$ )	EOP ( $\pi_{EOP}$ )	
Binary classification	200	47 (LR)	2PC	3.34 sec	3.36 sec	239.24 sec	238.69 sec
			3PC	1.37 sec	1.67 sec	6.41 sec	6.43 sec
				EOD ( $\pi_{EOD}$ )	Grp. Acc ( $\pi_{GACC}$ )	EOD ( $\pi_{EOD}$ )	Grp. Acc ( $\pi_{GACC}$ )
Multi-class classification	56	1.48M (ConvNet)	2PC	5,866.47 sec	5,866.46 sec	42,745.09 sec	42,742.00 sec
			3PC	29.92 sec	29.84 sec	199.02 sec	198.99 sec
				EOD ( $\pi_{EOD}$ )	Grp. Acc ( $\pi_{GACC}$ )	EOD ( $\pi_{EOD}$ )	Grp. Acc ( $\pi_{GACC}$ )

Table 2: Time taken to execute individual MPC protocols in PRIVFAIR – includes time for making inferences and fairness evaluation. Binary classification corresponds to credit score classification on the German credit score data set. Multi-class classification corresponds to detecting one of 7 emotions in an image using the RAVDESS data set. The times do not include compile time and are an average over 5 runs. DP: Demographic Parity, EOP: Equal Opportunity, EOD: Equalized Odds, Grp. Acc: Sub-group accuracy. 2PC-Passive: (Cramer et al. 2018), 2PC-Active: (Damgård et al. 2019) , 3PC-Passive: (Araki et al. 2016) , 3PC-Active: (Dalskov, Escudero, and Keller 2021).

$$\begin{aligned} \llbracket \text{POS}_{A=1} \rrbracket &\leftarrow \llbracket \text{TP}_{A=1} \rrbracket + \llbracket \text{FP}_{A=1} \rrbracket \\ \llbracket \text{POS}_{A=0} \rrbracket &\leftarrow \llbracket \text{TP}_{A=0} \rrbracket + \llbracket \text{FP}_{A=0} \rrbracket \end{aligned}$$

Additionally to compute the ratio, we adopt computations similar to Protocol 2, by including lines for computing counts of the instances belonging to each group (Lines 2–5), followed by the lines to compute the ratio

$$\begin{aligned} \llbracket \text{DP}_{A=1} \rrbracket &\leftarrow \pi_{\text{DIV}}(\llbracket \text{POS}_{A=1} \rrbracket, \llbracket \text{count}_{A=1} \rrbracket) \\ \llbracket \text{DP}_{A=0} \rrbracket &\leftarrow \pi_{\text{DIV}}(\llbracket \text{POS}_{A=0} \rrbracket, \llbracket \text{count}_{A=0} \rrbracket) \end{aligned}$$

At the end, Line 25 in Protocol 1 is replaced with

**return**  $\llbracket \text{DP}_{A=1} \rrbracket, \llbracket \text{DP}_{A=0} \rrbracket$

## Results

PRIVFAIR is implemented<sup>1</sup> on top of MP-SPDZ (Keller 2020) and is sufficiently generic to take advantage of the variety of underlying MPC schemes in MP-SPDZ.

We perform all our experiments on virtual machines (host) on Google Cloud Platform (GCP) with 8 vCPUs, 32 GB RAM and egress bandwidth limited to 16 Gbps. We use mixed computations (Mohassel and Rindal 2018; Araki et al. 2018; Demmler, Schneider, and Zohner 2015; Escudero et al. 2020) that switch between arithmetic ( $\mathbb{Z}_{2^k}$  with  $k = 64$ ) and binary ( $\mathbb{Z}_2$ ) computations for efficiency. All integer additions and multiplications are performed over the arithmetic domain and any non-linear functions such as comparisons are computed over the binary domain.

We report the time taken to execute the MPC protocols in PRIVFAIR under different security settings in Table 2, namely for 2PC and 3PC with passive or with active adversaries. To perform experiments for *binary classification*, we audit a logistic regression (LR) model for credit score classification on the benchmark German credit score data set (Dua and Graff 2017). The servers run MPC protocols  $\pi_{DP}$  and  $\pi_{EOP}$  from PRIVFAIR to compute demographic parity and equality of opportunity, using gender as the sensitive attribute. For *multi-class classification*, we demonstrate

PRIVFAIR to audit a ConvNet ( $\sim 1.48\text{M}$  parameters) model for emotion recognition from images (Pentyala, Dowsley, and De Cock 2021). As audit data, we use 56 images from the RAVDESS data set (Livingstone and Russo 2018) corresponding to different emotions, namely *neutral, happy, sad, angry, fearful, disgust, and surprised*. Using gender as the sensitive attribute, the servers run MPC protocols  $\pi_{EOD}$  and  $\pi_{GACC}$  from PRIVFAIR to compute equalized odds and sub-group accuracy.

Regarding utility we note that the LR and ConvNet models were trained in the clear, and that the  $\pi_{\text{INFER}}$  protocol for secure inference with these models infer the same labels as one would obtain without encryption, i.e. there is no loss of utility (accuracy). The results of the secure fairness auditing protocols are therefore also the same as one would obtain without encryption. In the remainder of this section, we therefore focus on the runtimes.

The times reported in Table 2 correspond to the entire runtime of the MPC protocols, i.e. both the so-called offline and online phases. The offline phase includes any preprocessing required to begin executing the MPC protocol (such as the generation of the correlated randomness that is needed for the secure multiplication  $\pi_{\text{MUL}}$ ) and is independent of the specific input values; the online phase is where the MPC protocol executes on the specific inputs.

Furthermore, the times reported in Table 2 encompass the time needed to infer class labels for all the instances in the audit data (such as on Line 1 in Protocol 1 and Protocol 2) and the time needed to evaluate the individual applicable fairness notions. The substantial differences in runtime for auditing the logistic regression model (binary classification) vs. the ConvNet model (multi-class classification) stem from large differences in runtime for the inference step. For example, auditing 56 images belonging to 7 classes with a ConvNet of 1.48M parameters to report EOD takes 42.745.09 sec in the active 2PC setting. Out of this, 42515.82 sec were taken to classify the 56 images in the audit data in a privacy-preserving manner (Line 1 of Protocol 1), while execution of the rest of Protocol 1 to evaluate the fairness metric itself

<sup>1</sup><https://bitbucket.org/uwtpmml/privfair>

took 229.27 sec. Nearly all the time taken by PRIVFAIR’s protocols for auditing the image classifier are spent on inference, while the time taken to evaluate only the fairness metrics given the labels is near real-time and practical.

Our runtime results for the time-consuming inference (the first step in the fairness auditing protocols) align with the observations in the MPC literature (Dalskov, Escudero, and Keller 2020, 2021; Pentylala, Dowsley, and De Cock 2021) where there is also a substantial difference between protocols for passive and active adversarial settings. Providing security in the presence of active or “malicious” adversaries, i.e. ensuring that no such adversarial attack can succeed, comes at a much higher computational cost than in the passive case. The same holds for a dishonest-majority 2PC setting where each party (server) only trusts itself, versus a much faster 3PC honest-majority setting. We note that unlike in privacy-preserving inference applications, where fast responsiveness of the system can be of utmost importance for practical applications (Dalskov, Escudero, and Keller 2020; Pentylala, Dowsley, and De Cock 2021), in fairness auditing use cases such as the ones we consider in this paper, one can typically afford longer runtimes, justifying the use of the most stringent security settings.

All the above results carry over to any similar, same sized audit data sets and models. With increase in the size of the data set or the number of parameters of the model, the times for execution of the MPC protocols will increase linearly (Pentylala, Dowsley, and De Cock 2021).

## Conclusion & Future Directions

In this paper, we presented PRIVFAIR, a first-of-its-kind library for privacy-preserving fairness audits of ML models. The protocols in PRIVFAIR allow an investigator Bob with audit data to evaluate the fairness of a model owner Alice’s ML model without requiring Bob or Alice to disclose their inputs to anyone in an unencrypted manner. The computations required for the fairness audits are performed by servers (parties) in a privacy-preserving manner. To protect both the trained model parameters and the audit data, we use Secure Multi-Party Computation (MPC) which allows for parties to jointly compute functions over encrypted shares of data. PRIVFAIR enables 2PC scenarios in which Alice and Bob act as the 2 parties themselves, as well as 3PC scenarios in which Alice and Bob outsource the fairness audit to untrusted servers in the cloud. The latter service can be adapted to be used not only for synchronous but also asynchronous MPC computations, where servers each store the corresponding secret shares of the models that are to be investigated and an external investigator can choose the desired model to investigate.

The current open-source implementation of PRIVFAIR includes MPC protocols for the well established fairness metrics of equalized odds, equal opportunity, subgroup accuracy, and demographic parity. The fairness auditing protocols can be selected based on their suitability to the application where fairness is to be guaranteed. PRIVFAIR can be easily used and extended to report other statistical notions of fairness (Verma and Rubin 2018) such as treatment equality (Berk et al. 2021) and predictive equality (Corbett-Davies

et al. 2017). Furthermore, while we have demonstrated the use of PRIVFAIR for auditing of logistic regression and ConvNet models, the fairness auditing protocols can be used for any kind of model architecture for which an MPC protocol for secure inference is available, including decision tree models, random forests, support vector machines, and other kinds of neural networks (De Cock et al. 2019; Fritchman et al. 2018; Agrawal et al. 2019; Reich et al. 2019; Dalskov, Escudero, and Keller 2020).

The applicability and usefulness of PRIVFAIR stretches well beyond the applications considered in our experiments. AI services that deal with sensitive data such as in healthcare, banking, predictive policing etc. not only need to protect data but also ensure that unbiased AI services are available to all. PRIVFAIR is a first step where AI developers and AI end users can collaborate to aim for unbiased AI services, while protecting their data. As such, the impact of using PRIVFAIR can be significant in all fields where ML models appear for automated decision making, e.g., education, housing, law-enforcement, healthcare, and banking, as well as new application domains yet to be discovered.

## Acknowledgements

Funding support for project activities has been partially provided by The University of Washington Tacoma’s Founders Endowment fund, Canada CIFAR AI Chair, Facebook Research Award for Privacy Enhancing Technologies, and the Google Cloud Research Credits Program.

## References

- Agrawal, N.; Shahin Shamsabadi, A.; Kusner, M.; and Gascón, A. 2019. QUOTIENT: two-party secure neural network training and prediction. In *ACM SIGSAC Conference on Computer and Communications Security*, 1231–1247.
- Ali, M.; Sapiezynski, P.; Bogen, M.; Korolova, A.; Mislove, A.; and Rieke, A. 2019. Discrimination through optimization: How Facebook’s Ad delivery can lead to biased outcomes. *Proc. of the ACM on Human-Computer Interaction*, 3: 1–30.
- Angwin, J.; Larson, J.; Mattu, S.; and Kirchner, L. 2016. Machine bias. *ProPublica*, May 23.
- Araki, T.; Barak, A.; Furukawa, J.; Keller, M.; Lindell, Y.; Ohara, K.; and Tsuchida, H. 2018. Generalizing the SPDZ Compiler For Other Protocols. *Cryptology ePrint Archive*, Report 2018/762. <https://eprint.iacr.org/2018/762>.
- Araki, T.; Furukawa, J.; Lindell, Y.; Nof, A.; and Ohara, K. 2016. High-throughput semi-honest secure three-party computation with an honest majority. In *ACM SIGSAC Conference on Computer and Communications Security*, 805–817.
- Bellamy, R. K.; Dey, K.; Hind, M.; Hoffman, S. C.; Houde, S.; Kannan, K.; Lohia, P.; Martino, J.; Mehta, S.; Mojsilović, A.; et al. 2019. AI Fairness 360: An extensible toolkit for detecting and mitigating algorithmic bias. *IBM Journal of Research and Development*, 63(4/5): 4–1.
- Berk, R.; Heidari, H.; Jabbari, S.; Kearns, M.; and Roth, A. 2021. Fairness in criminal justice risk assessments: The state of the art. *Sociological Methods & Research*, 50(1): 3–44.

- Bird, S.; Dudík, M.; Edgar, R.; Horn, B.; Lutz, R.; Milan, V.; Sameki, M.; Wallach, H.; and Walker, K. 2020. Fairlearn: A toolkit for assessing and improving fairness in AI. Technical Report MSR-TR-2020-32, Microsoft.
- Calders, T.; and Verwer, S. 2010. Three naive Bayes approaches for discrimination-free classification. *Data Mining and Knowledge Discovery*, 21(2): 277–292.
- Carlini, N.; Liu, C.; Erlingsson, Ú.; Kos, J.; and Song, D. 2019. The secret sharer: Evaluating and testing unintended memorization in neural networks. In *USENIX 2019*, 267–284.
- Catrina, O.; and Saxena, A. 2010. Secure computation with fixed-point numbers. In *International Conference on Financial Cryptography and Data Security*, 35–50. Springer.
- Corbett-Davies, S.; Pierson, E.; Feller, A.; Goel, S.; and Huq, A. 2017. Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd acm sigkdd international conference on knowledge discovery and data mining*, 797–806.
- Cramer, R.; Damgård, I.; Escudero, D.; Scholl, P.; and Xing, C. 2018. SPD $\mathbb{Z}_{2^k}$ : Efficient MPC mod  $2^k$  for Dishonest Majority. In *Annual International Cryptology Conference*, 769–798. Springer.
- Cramer, R.; Damgard, I.; and Nielsen, J. 2015. *Secure Multiparty Computation and Secret Sharing*. New York: Cambridge University Press Print.
- Dalskov, A.; Escudero, D.; and Keller, M. 2020. Secure evaluation of quantized neural networks. *Proceedings on Privacy Enhancing Technologies*, 2020(4): 355–375.
- Dalskov, A.; Escudero, D.; and Keller, M. 2021. Fantastic four: Honest-majority four-party secure computation with malicious security. In *USENIX 2021*, 2183–2200.
- Damgård, I.; Escudero, D.; Frederiksen, T.; Keller, M.; Scholl, P.; and Volgushev, N. 2019. New primitives for actively-secure MPC over rings with applications to private machine learning. In *IEEE Symposium on Security and Privacy (SP)*, 1102–1120.
- De Cock, M.; Dowsley, R.; Horst, C.; Katti, R.; Nascimento, A.; Poon, W.-S.; and Truex, S. 2019. Efficient and Private Scoring of Decision Trees, Support Vector Machines and Logistic Regression Models based on Pre-Computation. *IEEE Trans Dependable Secure Comput*, 16(2): 217–230.
- De Cock, M.; Dowsley, R.; Nascimento, A. C. A.; Railsback, D.; Shen, J.; and Todoki, A. 2021. High performance logistic regression for privacy-preserving genome analysis. *BMC Medical Genomics*, 14(1): 23.
- Demmler, D.; Schneider, T.; and Zohner, M. 2015. ABY: A framework for efficient mixed-protocol secure two-party computation. In *NDSS*.
- Dua, D.; and Graff, C. 2017. UCI Machine Learning Repository.
- Dwork, C.; Hardt, M.; Pitassi, T.; Reingold, O.; and Zemel, R. 2012. Fairness through awareness. In *Proc. of the 3rd Innovations in Theoretical Comp. Science Conference*, 214–226.
- Ekstrand, M. D.; Joshaghani, R.; and Mehrpouyan, H. 2018. Privacy for all: Ensuring fair and equitable privacy protections. In *Conf. on Fairness, Accountability and Transparency*.
- Escudero, D.; Ghosh, S.; Keller, M.; Rachuri, R.; and Scholl, P. 2020. Improved Primitives for MPC over Mixed Arithmetic-Binary Circuits. Cryptology ePrint Archive, Report 2020/338. <https://eprint.iacr.org/2020/338>.
- Fritchman, K.; Saminathan, K.; Dowsley, R.; Hughes, T.; De Cock, M.; Nascimento, A.; and Teredesai, A. 2018. Privacy-Preserving Scoring of Tree Ensembles: A Novel Framework for AI in Healthcare. In *Proc. of 2018 IEEE BigData*, 2412–2421.
- Galhotra, S.; Brun, Y.; and Meliou, A. 2017. Fairness testing: testing software for discrimination. In *Proc. of the 11th Joint Meeting on Foundations of Software Engineering*, 498–510.
- Hardt, M.; Price, E.; and Srebro, N. 2016. Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems*, 3315–3323.
- Kearns, M.; Neel, S.; Roth, A.; and Wu, Z. S. 2018. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In *Proc. of ICML*, 2564–2572.
- Keller, M. 2020. MP-SPDZ: A versatile framework for multi-party computation. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, 1575–1590.
- Kleinberg, J.; Mullainathan, S.; and Raghavan, M. 2016. Inherent trade-offs in the fair determination of risk scores. *arXiv preprint arXiv:1609.05807*.
- Livingstone, S. R.; and Russo, F. A. 2018. The Ryerson Audio-Visual Database of Emotional Speech and Song (RAVDESS): A dynamic, multimodal set of facial and vocal expressions in North American English. *PLoS one*, 13(5).
- Mohassel, P.; and Rindal, P. 2018. ABY3: A Mixed Protocol Framework for Machine Learning. Cryptology ePrint Archive, Report 2018/403. <https://eprint.iacr.org/2018/403>.
- Pentyala, S.; Dowsley, R.; and De Cock, M. 2021. Privacy-Preserving Video Classification with Convolutional Neural Networks. In *38th International Conference on Machine Learning*, volume 139 of *PMLR*.
- Pleiss, G.; Raghavan, M.; Wu, F.; Kleinberg, J.; and Weinberger, K. Q. 2017. On fairness and calibration. In *Advances in Neural Information Processing Systems*, 5680–5689.
- Reich, D.; Todoki, A.; Dowsley, R.; De Cock, M.; and Nascimento, A. 2019. Privacy-Preserving Classification of Personal Text Messages with Secure Multi-Party Computation. In *Advances in Neural Information Processing Systems*, 3752–3764.
- Verma, S.; and Rubin, J. 2018. Fairness definitions explained. In *2018 IEEE/ACM International Workshop on Software Fairness (FairWare)*.
- Vigdor, N. 2019. Apple card investigated after gender discrimination complaints. *The New York Times*, Nov 10.
- Wall, S.; and Schellmann, H. 2021. LinkedIn’s job-matching AI was biased. The company’s solution? More AI. *MIT Technology Review*, June 23.