

## RESEARCH ARTICLE

# The fairness-accuracy Pareto front

Susan Wei<sup>1</sup>  | Marc Niethammer<sup>2</sup>

<sup>1</sup>School of Mathematics and Statistics,  
University of Melbourne, Melbourne,  
Australia

<sup>2</sup>Department of Computer Science,  
University of North Carolina at Chapel  
Hill, Chapel Hill, North Carolina, USA

## Correspondence

Susan Wei, School of Mathematics and  
Statistics, University of Melbourne,  
Victoria, Australia.

Email: susan.wei@unimelb.edu.au

## Funding information

Australian Research Council,  
Grant/Award Number: DE200101253

## Abstract

Algorithmic fairness seeks to identify and correct sources of bias in machine learning algorithms. Confoundingly, ensuring fairness often comes at the cost of accuracy. We provide formal tools in this work for reconciling this fundamental tension in algorithm fairness. Specifically, we put to use the concept of Pareto optimality from multiobjective optimization and seek the fairness-accuracy Pareto front of a neural network classifier. We demonstrate that many existing algorithmic fairness methods are performing the so-called linear scalarization scheme, which has severe limitations in recovering Pareto optimal solutions. We instead apply the Chebyshev scalarization scheme which is provably superior theoretically and no more computationally burdensome at recovering Pareto optimal solutions compared to the linear scheme.

## KEYWORDS

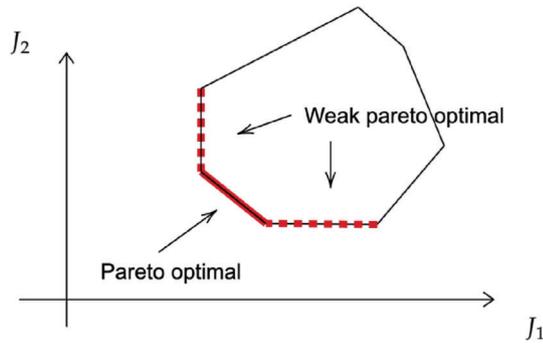
fairness, multiobjective optimization, neural network, Pareto front, Pareto optimality

## 1 | INTRODUCTION

The emerging field of algorithmic fairness is concerned with the understanding and correcting of biases in machine learning algorithms. The scope of algorithmic fairness research is wide, including such goals as understanding post-hoc the cause of bias in both the data collection and the algorithm, correcting pre-hoc the data collection and algorithm design process to mitigate unfairness, defining measurements of fairness, and establishing regulatory guidelines to govern algorithm deployment. In this work, we highlight an underappreciated challenge to ascertaining algorithmic fairness—while we would like an algorithm to be both fair and accurate, these two objectives actually conflict with each other. As an extreme example, a completely randomized classifier is also a perfectly fair classifier by almost all measures of fairness; we would, however, be loathe to accept the low accuracy and arbitrariness of randomization. Thus, recognizing fairness and accuracy as competing objectives, we put to use the concept of Pareto optimality. Specifically, we apply tools from

multiobjective optimization to find the fairness-accuracy Pareto front of a deep neural network classifier.

A typical multiobjective optimization problem is given by  $\arg \min_{\theta} [J_1(\theta), \dots, J_k(\theta)]^T$ , where each  $J_i$  is a loss function from some parameter space  $\Theta$  to  $\mathbb{R}$ . The task is made difficult by the fact that it is rarely possible to find a  $\theta^*$  such that all individual objectives are minimized *simultaneously at  $\theta^*$* . In such circumstances, one can appeal to Pareto optimality, a well-established notion in optimization [1] that is used often in engineering and economics. Pareto optimality is based on the following partial order: for  $a, b \in \mathbb{R}^p$ , we say  $a \leq b$  if and only if every component of  $a$  is less than or equal to the corresponding component of  $b$ . We say  $\theta \in \Theta$  is **Pareto optimal** if and only if it is nondominated, that is, there does not exist any  $\tilde{\theta} \in \Theta$  such that  $(J_1(\tilde{\theta}), \dots, J_k(\tilde{\theta})) \leq (J_1(\theta), \dots, J_k(\theta))$  with at least one strict inequality. Similarly, we say  $\theta$  is **weakly Pareto optimal** if there exists no other  $\tilde{\theta}$  such that  $J_i(\tilde{\theta}) < J_i(\theta)$  for all  $i = 1, \dots, k$ . Note that while a Pareto optimal point is also weakly Pareto optimal, the converse is not true. The **Pareto front** is the set of all Pareto optimal



**FIGURE 1** This diagram illustrates the difference between weakly Pareto optimal and Pareto optimal points. All other points in the feasible region, represented by the polygon, are dominated points. The Pareto front itself, consisting of Pareto optimal points, is highlighted in solid red, while the weakly Pareto optimal points that are not Pareto optimal are highlighted in dashed red

points. Figure 1 displays a toy Pareto front for a bi-objective optimization problem.

Almost all techniques for finding the Pareto front are based on scalarization, that is, turning the vector objective into a scalar one by aggregating the individual objectives. A successful scalarization scheme should satisfy the following two desiderata.

D1. The solutions to the scalarized objective are (weakly) Pareto optimal.

D2. All Pareto optimal points are recoverable from the scalarization scheme.

## 1.1 | Linear scalarization

The simplest scalarization scheme is the linear scalarization Scheme [1]. Here, the vector objective is replaced with a weighted sum, which we will call a **linear scalarization problem (LSP)**,

$$\sum_{i=1}^k \lambda_i J_i, \quad \lambda_i \in \mathbb{R}_{>0}, \quad i = 1, \dots, k. \quad (1)$$

Sometimes, the constraint,  $\sum_{i=1}^k \lambda_i = 1$ , is imposed, but this is not in fact necessary, for example, if the individual

objectives are not normalized. The linear scalarization scheme satisfies desideratum D1. Specifically, no matter the weights  $\lambda_i$  chosen, a solution of the linear scalarization problem is weakly Pareto optimal, see Proposition 8 in [1]. However, the linear scalarization scheme critically fails desideratum D2. In fact, it is only able to recover points on the *convex hull* of the Pareto front [2]. In particular, when the Pareto front of interest is nonconvex, the naive linear scalarization scheme may only find such  $\theta$  that strongly favors one of the objectives [1]. We illustrate this phenomenon in Figure 2.

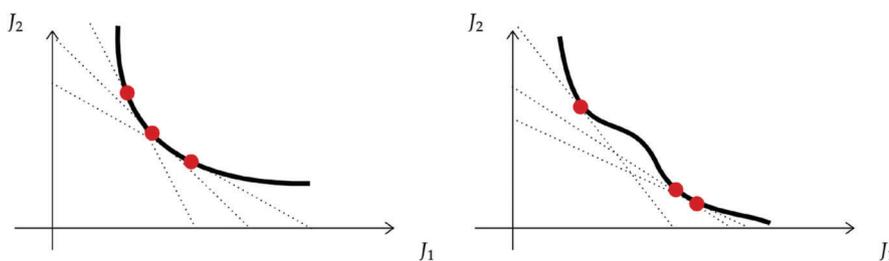
Despite the known drawbacks to the linear scalarization scheme for finding Pareto optimal points, almost all train-time algorithmic fairness methods, to be reviewed in the following section, ascribe to this approach. To see this, we note that many popular algorithmic fairness techniques minimize an objective of the form

$$\text{error}(\theta) + \lambda \times \text{unfairness}(\theta),$$

where  $\theta$  is the model parameters to be learned. That is, they regularize the learning algorithm by employing the fairness criterion as a penalty term during training [3–6]. It is easy to recognize this constrained optimization as the linear scalarization problem in Equation (1). Thus existing algorithmic-fairness methods can be expected to inherit all the limitations of the linear scalarization scheme in regards to finding Pareto optimal points.

## 1.2 | Proposal

In this work, we use instead the Chebyshev scalarization scheme (Section 3) to find the fairness-accuracy Pareto front of a deep neural network classifier. The Chebyshev scheme is no more computationally burdensome than the linear scheme while enjoying theoretically superior performance when the Pareto front is nonconvex (as it may well be). The final deliverable is a set of neural networks spanning the fairness-accuracy space from the high-accuracy-low-fairness corner to the high-fairness-low-accuracy corner. This could then be consulted by an interested party, for example, a relevant



**FIGURE 2** Two hypothetical Pareto fronts are displayed: convex (left) and nonconvex (right). The dashed lines represent different weight combinations in the LSP (1). Pareto optimality occurs only at the tangential points of the dashed lines. It is clearly that the LSP cannot recover Pareto optimal points that reside in the nonconvex region of the Pareto front

regulatory body, to decide on a particular classifier with the acceptable trade-off.

In recognition of the dominant role inhabited by deep learning in predictive tasks, our examples are limited to deep neural network classifiers. We should note, however, that nothing prevents the application of the proposed framework to classifiers besides neural networks, as long as the classifier is amenable to training by backpropagation. This includes models that are more familiar to the statistical audience such as logistic regression and support vector machine. In fact, logistic regression is simply a multilayer perception with no hidden layers and thus a special case of the feedforward neural network we will work with.

## 2 | RELATED WORK

### 2.1 | Multiobjective optimization in machine learning and statistics

Multiobjective optimization techniques are gaining increasing traction in machine learning [7]. Bayesian optimization [8] and reinforcement learning [9] seem to have particularly embraced ideas from multiobjective optimization. The multiple-gradient descent algorithm, a common gradient-based multiobjective optimization technique, has been applied to kernel learning [10] and Bayesian optimization [11]. Although multiple-gradient descent in its original form does not scale up to the high dimensionality of the parameter space of a neural network, Sener and Koltun [12] proposed a workaround that demonstrated good results on standard deep learning benchmarking datasets. Another work in this vein seeks to further articulate user preference for a specific trade-off, by dividing a deep multitask learning problem into different subproblems [13].

### 2.2 | Algorithmic fairness

Let us now briefly review those works in algorithm fairness that are concerned with correcting the bias of machine learning algorithms. These methods are set apart according to the stage in which intervention is taken. The first class of methods attempts to remove bias from the input data itself. These methods rest on the premise that once proper preprocessing is accomplished, any classifier can be used to produce subsequently fair predictions [14–17]. On the other hand, postprocessing techniques directly operate on the classifier output and are, in principle, amenable to any classifier. The technique in Hardt et al. [3] for instance seeks to learn a monotone transformation of the classifier

output to enforce demographic parity or equalized odds, two standard definitions of fairness that we will review in Section 4.

The third type of algorithmic fairness methods directly intervenes during training. Generally speaking, these train-time methods add the fairness criterion as a regularization term to the main objective of minimizing predictive error [5, 6, 18–23]. While early works of this type [6, 21, 24] focused on simple machine learning algorithms such as logistic regression, more recent work [25–28] can handle more complex models such as neural networks.

The proposed work also intervenes at train time. It appears to be the first in the algorithmic fairness literature to advocate for the important role that Pareto optimality should play, considering that fairness and accuracy are competing objectives of interest. Many methods do not consider the trade-off curve or when they do, they do not use proper techniques to recover Pareto optimal points in the fairness-accuracy space.

## 3 | THE FAIRNESS-ACCURACY PARETO FRONT

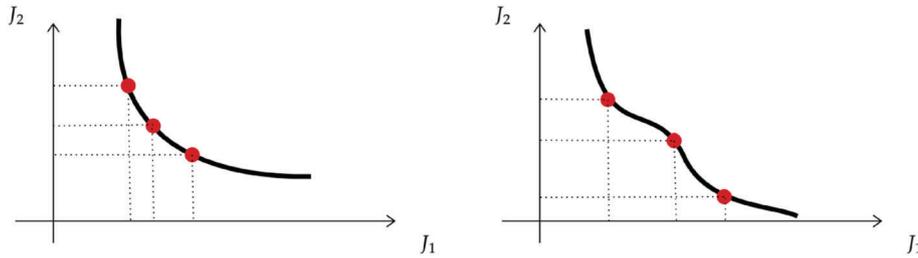
In this section, we apply the Chebyshev scalarization scheme to estimate the fairness-accuracy Pareto front of a neural network binary classifier. Suppose the data consist of input variables  $\mathbf{x} \in \mathbb{X} \subset \mathbb{R}^p$  standardized to mean zero and unit variance, binary response  $\mathbf{y} \in \mathbb{Y}$  indicating class membership, and binary sensitive variable  $\mathbf{a}$ . Let all discrete variables be dummy encoded. A feedforward neural network is a repeated composition of affine transformation followed by nonlinear transformation. Let  $m_0 = p$  and define  $w^{(l)} \in \mathbb{R}^{m_l \times m_{l-1}}$  and  $b^{(l)} \in \mathbb{R}^{m_l}$ ,  $l = 1, \dots, L$  as the parameters in the  $l$ -th layer of a fully connected feedforward neural network with  $L$  layers. Consider the affine transformation  $h^{(l)} : \mathbb{R}^{m_{l-1}} \rightarrow \mathbb{R}^{m_l}$  and nonlinear transformation  $v^{(l)}$  given by

$$\begin{aligned} h^{(l)} &= w^{(l)}v^{(l-1)} + b^{(l)}, \quad l = 1, \dots, L \\ v^{(l)} &= \sigma^{(l)} \circ h^{(l)}, \quad l = 1, \dots, L, \end{aligned}$$

where  $v^{(0)} = id$  is the identity function and  $\sigma^{(l)} : \mathbb{R}^{m_l} \rightarrow \mathbb{R}^{m_l}$  is a nonlinear transformation (more frequently called the activation function) that acts component-wise.

The activation function in the final layer,  $\sigma^{(L)}$ , will be restricted to the sigmoid function,  $(1 + \exp(-x))^{-1}$ , so that the classifier outputs scores between 0 and 1. We use the ReLU activation function,  $\max(0, x)$ , in all other layers for our experiments. Let  $\mathbf{h}_i^{(l)}$  be shorthand for the application of the function  $h^{(l)}$  to input feature  $\mathbf{x}_i$ , that is,

$$\mathbf{h}_i^{(l)} = h^{(l)}(\mathbf{x}_i).$$



**FIGURE 3** We again consider the two hypothetical Pareto fronts from Figure 2. This time the dashed lines represent different weight combinations in the CSP (3). We can see that the CSP, in contrast to the LSP, can recover Pareto optimal points that reside in the nonconvex region of the Pareto front

Collecting all parameters  $(w^{(l)}, b^{(l)}), l = 1, \dots, L$  into the parameter vector  $\theta$ , let  $f_\theta : \mathbb{X} \rightarrow [0, 1]$  be a fully connected feedforward neural network with parameter  $\theta \in \Theta$  constructed as above, that is, the neural network acts as  $f_\theta(x) = v^{(L)}(x)$ .

Let  $P$  be the joint distribution of the data  $(\mathbf{x}, \mathbf{a}, \mathbf{y})$ . The accuracy of  $f_\theta$  will be measured by its risk,

$$R(\theta) = \mathbb{E}_P \mathcal{L}(f_\theta(\mathbf{x}), \mathbf{y}),$$

where  $\mathcal{L}$  is a loss function. Since we are interested in binary classification, we will limit future discussion to the binary cross-entropy loss,  $\mathcal{L} : [0, 1] \times \{0, 1\} \rightarrow \mathbb{R}$  given by

$$\mathcal{L}(\hat{p}, y) = y \log \hat{p} + (1 - y) \log(1 - \hat{p}).$$

Next, let  $U(\theta)$  be some population measure for the fairness of  $f_\theta$ . We shall take the convention that higher values of  $U$  are undesirable, yielding algorithms that are more unfair. Since we wish for the learning algorithm  $f_\theta$  to be both accurate and fair, we have the following *vector* objective function

$$\arg \min_{\theta} \begin{bmatrix} R(\theta) \\ U(\theta) \end{bmatrix}. \quad (2)$$

### 3.1 | Chebyshev scalarization

As an alternative to the linear scalarization scheme in (1), we consider the Chebyshev scalarization Scheme [29, 30] whereby the vector objective,  $[J_1(\theta), \dots, J_k(\theta)]^T$ , is replaced with what we will call the **Chebyshev scalarization problem (CSP)**,

$$\max_{i=1, \dots, k} \{ \lambda_i |J_i(\theta) - z_i^*| \}, \quad \lambda_i \in \mathbb{R}_{>0}, \quad i = 1, \dots, k \quad (3)$$

where  $z_i^* = \inf_{\theta} J_i(\theta)$ . The Chebyshev scalarization satisfies both desiderata D1 and D2. Namely the solutions to (3) are at least weakly Pareto optimal and hence possibly

Pareto optimal but not necessarily. Furthermore, any Pareto optimal solution can be obtained for some configuration of  $\lambda_i$ 's, see Proposition 10 of [1]. In particular, this means the Chebyshev scheme can find Pareto optimal points that reside in the nonconvex area of the Pareto front in contrast to the linear scalarization scheme. Figure 3 contains an illustration of this. However, we caution that this does *not* mean that every CSP solution is Pareto optimal because again the solution may be weakly Pareto optimal.

More advanced scalarization schemes than either the LSP or the CSP certainly exist. However, we focus on the CSP in this article as a first step in applying MOO to algorithmic fairness. CSP is an obvious candidate given its clear superiority to LSP. Furthermore, there is no additional computational overhead to solve the CSP compared to the LSP.

### 3.2 | Estimates of the individual components

Solving (3) first requires estimation of the population unknowns  $R(\theta)$  and  $U(\theta)$ . Since most learning algorithms proceed by empirical risk minimization, we can without loss of generality estimate  $R(\theta)$  with the plug-in estimator. Specifically, let  $R_n(\theta) = \mathbb{P}_n \mathcal{L}(f_\theta(\mathbf{x}), \mathbf{y})$  where  $\mathbb{P}_n = \frac{1}{n} \sum_{i=1}^n \delta_{(\mathbf{x}_i, \mathbf{a}_i, \mathbf{y}_i)}$  is the empirical measure for the training set  $(\mathbf{x}_1, \mathbf{a}_1, \mathbf{y}_1), \dots, (\mathbf{x}_n, \mathbf{a}_n, \mathbf{y}_n) \sim P$ . We will defer discussion of the estimation of  $U(\theta)$  to Section 4 in the context of specific fairness metrics. For now, let  $U_n(\theta)$  denote some estimator of  $U(\theta)$  based on the dataset of sample size  $n$ . (It does not need to be the plug-in estimator.)

### 3.3 | Standardization

The individual estimators  $R_n$  and  $U_n$  should be properly scaled so that the weights in (3) can better articulate user preferences. We will focus on discussing the standardization of  $R_n$  and  $U_n$  in the context where  $f_\theta$  is a neural network. As mini-batches are almost

always employed during training of a neural network, the risk and unfairness measures will be evaluated on mini-batches. Let  $\{(\mathbf{x}_1^*, \mathbf{a}_1^*, \mathbf{y}_1^*), \dots, (\mathbf{x}_b^*, \mathbf{a}_b^*, \mathbf{y}_b^*)\}$  be a mini-batch sample drawn without replacement from  $\mathbb{P}_n$ . Then, define the mini-batch estimate of the risk as  $R_b^*(\theta) = \frac{1}{b} \sum_{i=1}^b \mathcal{L}(f_\theta(\mathbf{x}_i^*), \mathbf{y}_i^*)$ . Similarly, let  $U_b^*(\theta)$  be the mini-batch estimate of the fairness measure.

First we solve (3) for  $\lambda = 0$ . This will not typically be an exact solution. For instance if  $f_\theta$  is a neural network, (3) can only be solved approximately, for example, optimization via stochastic gradient descent. Let  $R_{\min}$  and  $R_{\max}$  denote, respectively, the minimum and maximum value of  $R_b^*$  observed across the mini-batches over all epochs. We can similarly obtain  $U_{\min}$  and  $U_{\max}$  by solving (3) for  $\lambda = 1$ . Then, we standardize the accuracy and fairness components as follows

$$\tilde{R}_n(\theta) = (R_n(\theta) - R_{\min}) / (R_{\max} - R_{\min})$$

and

$$\tilde{U}_n(\theta) = (U_n(\theta) - U_{\min}) / (U_{\max} - U_{\min}).$$

### 3.4 | Pareto front candidates

Using the Chebyshev scalarization scheme leads us to Pareto front candidates given by the set  $\{\hat{\theta}_n^\lambda : \lambda \in \Lambda\}$  where

$$\hat{\theta}_n^\lambda = \arg \min_{\theta} \max \left\{ (1 - \lambda) \tilde{R}_n(\theta), \lambda \tilde{U}_n(\theta) \right\}. \quad (4)$$

and  $\Lambda \subset [0, 1]$  is a finite set of  $\lambda$  values, which includes  $\lambda = 0$  and  $\lambda = 1$ . Note that (4) follows directly from (3) since the standardization step renders  $z_i^* = 0$  for all  $i$ . Ideally, we would then solve (4) for a dense  $\Lambda$ . While intuitive, evenly distributed  $\lambda$ 's in the interval  $[0, 1]$  should be avoided as this often produces solutions that form clumps in the Pareto front, that is, evenly distributed  $\lambda$ 's in  $[0, 1]$  do *not* produce evenly distributed points in the multiobjective space. The experiments in Section 5 employ approximately evenly spaced values of  $\lambda$  on the log scale for the set  $\Lambda$ . Future work might involve more sophisticated techniques from multiobjective optimization, for example, the Normal-Boundary-Interactive [31], which adaptively selects  $\lambda$ .

Finally, to get a robust estimate of the Pareto front, we should average out, so to speak, the randomness associated with the training–testing split. If  $M$  pairs of training and testing splits are considered, then CANDIDATES, the output of Algorithm 1, contains in total  $M |\Lambda|$  Pareto candidates.

---

### Algorithm 1. Fairness-accuracy Pareto front candidates

---

**Input:** multiple splits of the data into training set (sample size  $n$ ) and testing set (sample size  $T$ ) and a finite set  $\Lambda \subset [0, 1]$  which includes  $\lambda = 0$  and  $\lambda = 1$ .  
Initialize CANDIDATES =  $\emptyset$ .  
**for** each training–testing split **to**  
    Solve (4) on the training set to obtain  $\hat{\theta}_n^\lambda$  for each  $\lambda \in \Lambda$   
    Add to CANDIDATES the set  $\{\hat{\theta}_n^\lambda : \lambda \in \Lambda\}$   
**end for**  
**Output:** CANDIDATES

---

### 3.5 | Visualization of results

To visualize the estimated Pareto front, we plot each element of CANDIDATES, the output of Algorithm 1, in the fairness-accuracy space as evaluated on the test set. Specifically, we calculate the out-of-sample risk of  $\theta$  as

$$R(\theta; \hat{P}_{\text{test}}) = \mathbb{E}_{\hat{P}_{\text{test}}} \mathcal{L}(f(\mathbf{x}; \theta), \mathbf{y}) = \frac{1}{T} \sum_{i=1}^T \mathcal{L}(f(\mathbf{x}_i^*; \theta), \mathbf{y}_i^*),$$

where  $\hat{P}_{\text{test}} = \frac{1}{T} \sum_{i=1}^T \delta_{(\mathbf{x}_i^*, \mathbf{a}_i^*, \mathbf{y}_i^*)}$  is the empirical measure of the test set  $\{(\mathbf{x}_i^*, \mathbf{a}_i^*, \mathbf{y}_i^*)\}_{i=1}^T$  drawn independently from the (unknown) distribution  $P$ . Let  $U(\theta; \hat{P}_{\text{test}})$  denote the fairness metric assessed on the test set. Again it need not be a plug-in estimator. Each Pareto front candidate can then be visualized by plotting  $R(\theta; \hat{P}_{\text{test}})$  versus  $U(\theta; \hat{P}_{\text{test}})$  for every  $\theta$  in candidates. We further cull these candidates by keeping only the nondominated points, forming the final Pareto front estimate.

Each point on the estimated Pareto front corresponds to a different neural network classifier  $f_\theta$ . A practitioner can examine the estimated Pareto front to select the neural network with an acceptable trade-off between fairness and accuracy.

## 4 | FAIRNESS MEASURES

We will be examining the fairness-accuracy Pareto front corresponding to a variety of fairness notions. In this section, we begin by reviewing two common notions of fairness—demographic and conditional parity. We then review causal approaches to fairness. Along the way, we introduce new estimators of conditional parity and introduce a new causal fairness measure, each of which may be of independent interest.

## 4.1 | Demographic parity and conditional parity

Many standard fairness measures involve checking the statistical independence between the prediction and the sensitive attribute. The classifier  $f_\theta(\mathbf{x})$  is said to exhibit **demographic parity** with respect to the sensitive attribute  $\mathbf{a}$  if  $f_\theta(\mathbf{x}) \perp\!\!\!\perp \mathbf{a}$ , where  $\perp\!\!\!\perp$  stands for independence. Intuitively, demographic parity assesses whether the predicted score depends on the sensitive variable. For example, a classifier predicting if a convicted criminal will reoffend exhibits demographic parity with respect to race if the distribution of  $f_\theta(\mathbf{x})$  is the same irrespective of race. Despite its intuitiveness, the drawbacks to demographic parity are well-documented [3, 32]. For instance, when the base rates differ across values of the sensitive attribute, satisfying demographic parity can come at the cost of discrimination.

Conditional parity, a term coined in Ritov et al. [33], is another notion of fairness that encompasses several fairness measures. The prediction  $f_\theta(\mathbf{x})$  is said to exhibit **conditional parity** with respect to sensitive attribute  $\mathbf{a}$  conditioned on  $\mathbf{u}$  if  $f_\theta(\mathbf{x}) \perp\!\!\!\perp \mathbf{a} \mid \mathbf{u}$ . Conditional parity in fact unifies several existing fairness definitions. For instance, the notion of equalized odds, introduced in Hardt et al. [3], is recovered by setting  $\mathbf{u}$  to  $\mathbf{y}$ , the true target class membership itself.

That the notions of demographic and conditional parity can strongly differ and may lead to seemingly paradoxical results was strikingly illustrated in Bickel et al. [34] for graduate admissions at UC Berkeley. Consider a classifier predicting whether an applicant should be admitted to graduate school. One may desire admission decisions to be independent of gender (demographic parity) or independent of gender *conditional* on a particular university department (conditional parity). Depending on which fairness measure is employed, radically different conclusions may be reached.

We propose to assess demographic and conditional parity using the mean–variance statistic proposed in Cui et al. [35]. The development of these estimators is not central to the proposed work but may be of separate interest so we defer its presentation to Appendix B.2.

## 4.2 | Causal fairness in the overlap population

Taking the causal approach to defining fairness means replacing the question “Is the learning prediction (conditionally) dependent on the sensitive attribute?” with the question “Does the sensitive attribute have a *causal effect* on the learning predictions?” Causal approaches

to defining fairness [36–39] are motivated by the consideration that selection bias will result in a study population different from the target population. This is precisely the situation in many datasets of interest in algorithmic fairness. For instance, in predicting recidivism, the training data suffers selection bias leading to a training population (reoffenders that were caught) likely different from the target population (all would-be reoffenders).

We propose a new causal fairness measure, which assesses the average effect of  $\mathbf{a}$  on the predicted probabilities  $f_\theta(\mathbf{x})$  in the so-called *overlap population* introduced in Li et al. [40]. The overlap population is of particular policy relevance. It is the subpopulation where the sensitive attribute  $\mathbf{a}$  occurs with equal probability, given the prediction inputs  $\mathbf{x}$ . For instance, in the context of recidivism prediction, the overlap population is the subpopulation with similar inputs, for example, criminal record, demographic characteristics, who could easily be either Caucasian or from a minority group. In contrast, many traditional causal estimands focus on unrealistic target populations comprised of individuals who are atypical for their particular value of  $\mathbf{a}$ .

Adopting the potential outcome framework of Imbens and Rubin [41] and assuming the Stable Unit Treatment Value Assumption, suppose the outcome of interest  $\mathbf{o}$  takes on one of two potential outcomes, either  $\mathbf{o}(0)$  or  $\mathbf{o}(1)$ , depending on whether  $\mathbf{a} = 0$  or  $\mathbf{a} = 1$ . Note that  $\mathbf{o} = \mathbf{o}(1)\mathbf{a} + \mathbf{o}(0)(1 - \mathbf{a})$ , that is, we can only ever observe one of the two potential outcomes. Furthermore, suppose the condition of *unconfoundedness* is satisfied, that is,  $\mathbf{a}$  is independent of  $\{\mathbf{o}(0), \mathbf{o}(1)\}$  conditional on  $\mathbf{x}$ .

The *average treatment effect for the overlap population* (ATO) introduced in Li et al. [40] is given by

$$\tau_{\text{ATO}}(\mathbf{o}; P) = \frac{\mathbb{E}_P[e(\mathbf{x})(1 - e(\mathbf{x}))(\mathbb{E}(\mathbf{o}(1)|\mathbf{x}) - \mathbb{E}(\mathbf{o}(0)|\mathbf{x}))]}{\mathbb{E}_P[e(\mathbf{x})(1 - e(\mathbf{x}))]}$$

where  $e(x) = P(\mathbf{a} = 1 | \mathbf{x} = x)$  is also called the propensity score. Again, the term *overlap* refers to the fact that the ATO articulates the causal effect among the **overlap population** which consists of subjects, given their covariates, who could appear with substantial probability in either value of the sensitive attribute. Consider measuring fairness as the causal effect of  $\mathbf{a}$  on the predicted probability  $f_\theta(\mathbf{x})$  in the overlap population, that is,

$$U(\theta) = \left| \tau_{\text{ATO}}(f_\theta(\mathbf{x}); P) \right|.$$

Its evaluation on the test set is straightforward. Define

$$U(\theta, \hat{P}_{\text{test}}) = \left| \tau_{\text{ATO}}(f_\theta(\mathbf{x}); \hat{P}_{\text{test}}) \right|.$$

Estimation of the ATO causal estimand first requires estimating the propensity score. In our experiments, we used a neural network to do this. We then calibrated the predicted probabilities using the temperature scaling procedure of Guo et al. [42]. Further details follow in the experiments section. For now, let  $\hat{e}(x)$  denote the propensity score estimate. Li et al. [40] proposed estimating  $\tau_{\text{ATO}}$  as follows:

$$\hat{\tau}_{\text{ATO}}(\mathbf{o}) = \frac{\sum_{i=1}^n \mathbf{a}_i \mathbf{o}_i w_i}{\sum_{i=1}^n \mathbf{a}_i w_i} - \frac{\sum_{i=1}^n (1 - \mathbf{a}_i) \mathbf{o}_i w_i}{\sum_{i=1}^n (1 - \mathbf{a}_i) w_i}, \quad (5)$$

where  $w_i$  are the so-called overlap weights given by

$$w_i = \begin{cases} 1 - \hat{e}(\mathbf{x}_i) & \text{if } \mathbf{a}_i = 1 \\ \hat{e}(\mathbf{x}_i) & \text{if } \mathbf{a}_i = 0. \end{cases} \quad (6)$$

Notably, the overlap weights smoothly down-weight subjects in the tails of the propensity score distribution, thereby mitigating the common problem of extreme propensity scores. In contrast, the standard inverse probability weights can suffer from excessive variance and correspond to emphasis on a target population, which may consist of subjects very atypical for their particular value of  $\mathbf{a}$ .

In our experiments, we considered two possibilities for  $U_n(\theta)$ . In the first, we measure the average effect in the overlap population of  $\mathbf{a}$  on the penultimate layer of the neural network,  $\mathbf{h}^{(L-1)}$  leading to

$$\begin{aligned} U_n(\theta) &= \left| \hat{\tau}_{\text{ATO}}(\mathbf{h}^{(L-1)}) \right| \\ &= \frac{\sum_{i=1}^n \mathbf{a}_i \mathbf{h}_i^{(L-1)} w_i}{\sum_{i=1}^n \mathbf{a}_i w_i} - \frac{\sum_{i=1}^n (1 - \mathbf{a}_i) \mathbf{h}_i^{(L-1)} w_i}{\sum_{i=1}^n (1 - \mathbf{a}_i) w_i} \end{aligned} \quad (7)$$

where  $w_i$  are the overlap weights as in (6). In the second option for  $U_n$ , we measure the average effect of  $\mathbf{a}$  on all intermediate layers of the neural network, leading to

$$U_n(\theta) = \sum_{l=1}^{L-1} \left| \hat{\tau}_{\text{ATO}}(\mathbf{h}^{(l)}) \right|. \quad (8)$$

The benefits of learning fair internal representations in a neural network were recognized in Madras et al. [27]. In particular, we may expect doing so can safeguard against bias in further downstream analyses, such as transfer learning.

## 5 | EXPERIMENTS

We apply the proposed Pareto front estimation procedure to two benchmarking datasets in the algorithmic

fairness literature: the ProPublica recidivism dataset and the UCI adult income dataset. The two datasets are summarized in Table 1. We also examine two possibilities for  $U_n$  in Algorithm 1 corresponding to either (7) or (8), that is, calculating the ATO in either the penultimate layer or all intermediate layers of the neural network, respectively.

### 5.1 | Data preprocessing

Missing values were preprocessed according to the accompanying code. In the UCI dataset, we wish to predict whether an individual has income above 50 K USD while remaining fair with respect to *gender*. Separately, we wish to perform the same prediction task in the UCI dataset while remaining fair with respect to *race*. In the recidivism dataset, we wish to predict whether an individual will recommit a crime in 2 years while remaining fair with respect to *race*. In total, we have three datasets—UCI (*gender*), UCI (*race*), and recidivism.

### 5.2 | Comparison to alternatives

For comparison, we implement, as a representative of the regularization approach, the adversarial learning technique proposed in Louppe et al. [43], which is not based on any specific fairness criterion. The idea is intuitive: the classifier and adversarial are engaged in a zero-sum game. Let  $\theta_{\text{clf}}$  and  $\theta_{\text{adv}}$  denote the parameters of the classifier and adversarial network, respectively. The classifier network attempts to make the best possible prediction of the binary output  $\mathbf{y}$  given the input  $\mathbf{x}$ , while ensuring that  $\mathbf{a}$  cannot be predicted well from the predicted score  $\hat{p}$ . The adversary, on the other hand, attempts to make the best possible prediction of the binary sensitive attribute  $\mathbf{a}$  based on the classifier's prediction  $\hat{p}$ . The adversary behaves according to the objective,

$$\min_{\theta_{\text{adv}}} \text{Loss}_{\mathbf{a}}(\theta_{\text{clf}}, \theta_{\text{adv}}),$$

where  $\text{Loss}_{\mathbf{a}}$  is associated to the prediction of the sensitive attribute  $\mathbf{a}$  given  $\hat{p}$ . The classifier behaves according to

$$\arg \min_{\theta_{\text{clf}}} \left[ \text{Loss}_{\mathbf{y}}(\theta_{\text{clf}}) - \beta \text{Loss}_{\mathbf{a}}(\theta_{\text{clf}}, \theta_{\text{adv}}) \right], \quad (9)$$

where  $\beta \in [0, 1]$  and  $\text{Loss}_{\mathbf{y}}$  measure the loss of predicting  $\mathbf{y}$  based on  $\mathbf{x}$ . The binary cross-entropy loss is employed in both objectives.

TABLE 1 Dataset descriptions

Dataset	Dataset features			Training size	Testing size	Minibatch size
	$Dim(\mathbf{x})$	Dinary outcome $y$	Sensitive $\mathbf{a}$			
Recidivism	12	Re-offend in 2 years?	binary race	3086	3086	150
UCI	93	Income above 50 K?	binary race	15,470	15,470	1000
UCI	93	Income above 50 K?	binary gender	15,470	15,470	1000

### 5.3 | Experimental details

For each of the three datasets—UCI (gender), UCI (race), and recidivism—the data are split into 100 training–testing sets, with the split reported in Table 1. Then, for each of the 100 training sets, we apply Algorithm 1 to find Pareto front candidates with  $R_n$  corresponding to the binary cross-entropy loss,  $U_n$  given by either (7) or (8), and  $\Lambda$  containing 15 different  $\lambda$  values in the interval  $[0, 1]$ , approximately evenly spaced on the log scale. Exact values used for  $\Lambda$  can be consulted in the accompanying code.

### 5.4 | Propensity score network

To calculate  $U_n$  in either (7) or (8), propensity scores must first be estimated. We used the same neural network architecture in all three datasets to estimate the propensity scores  $P(\mathbf{a} = 1|\mathbf{x})$ . The network has three fully connected layers, with 32 hidden units each, interspersed with a dropout layer with dropout probability 0.2. The ReLU activation function is used in all intermediate layers while the sigmoid function is used in the output layer.

Since the propensity score network is performing binary classification of  $\mathbf{a}$  based on input  $\mathbf{x}$ , we used the binary cross-entropy loss. The Adam optimization algorithm [44] was used to train the propensity network. The learning rate is set to 0.001. Training took place over 100 epochs. Mini-batch size is reported in Table 1 and was chosen to be around 5% of the training set size. After the propensity network is trained, we calibrate the probability prediction according to the methodology proposed in Guo et al. [42] where we used their GitHub code with no modification.

### 5.5 | Architecture $f_\theta$

The architecture of the neural network  $f_\theta$  in Algorithm 1 is reported in Table 2. The number of fully connected layers and number of hidden nodes in each layer (held constant over the layers) were tuned for each data setting

TABLE 2  $f_\theta$  network architecture

Dataset	Neural network features	
	Layers $L$	Hidden nodes
Recidivism	4	4
UCI	32	10

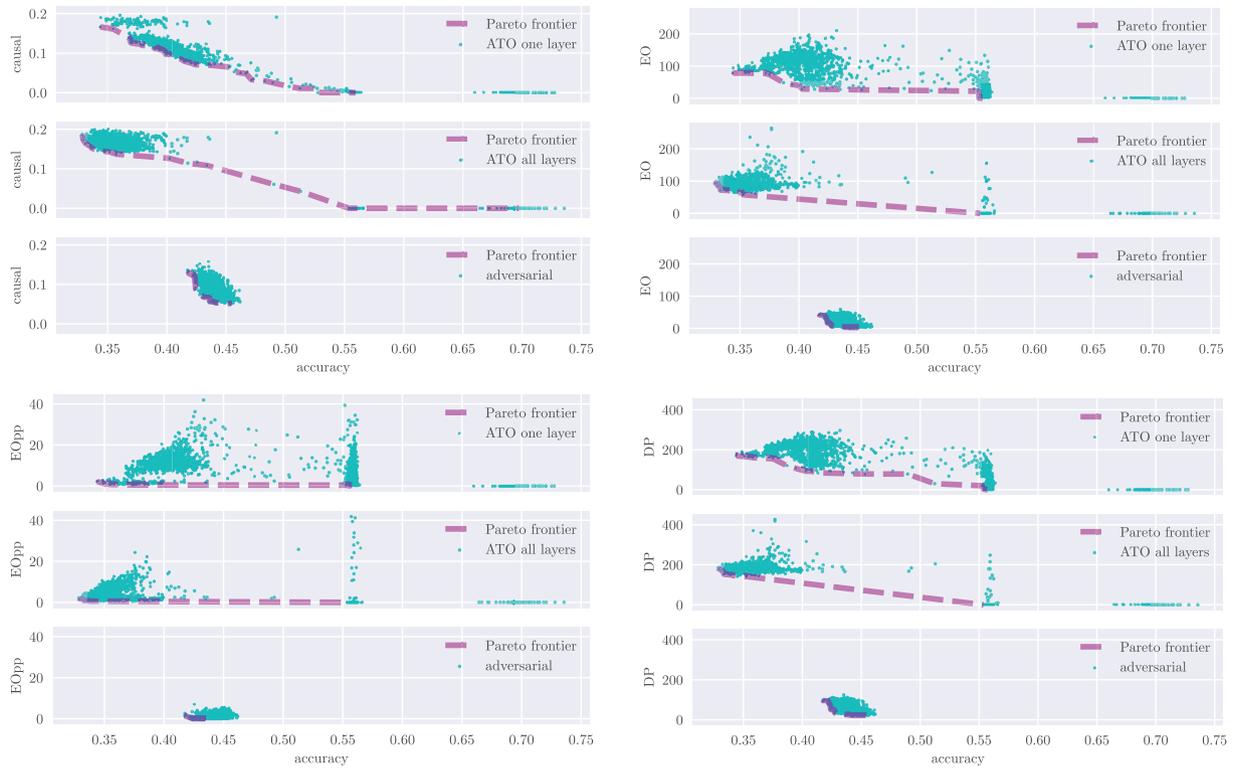
with the goal of not incurring over-fitting in the held-out test set. Each fully connected layer is interspersed with a dropout layer with dropout probability 0.2. The ReLU activation function is used in all intermediate layers, while the sigmoid function is used in the output layer.

To learn the network  $f_\theta$  in Algorithm 1, we again use Adam. The initial learning rate is set to 0.001. We reduce the learning rate when the training loss has stopped decreasing by using the ReduceLRonPlateau scheduler in PyTorch, setting the factor and patience variables to 0.9 and 10, respectively. All training took place over 500 epochs. Mini-batch sizes are as reported in Table 1.

### 5.6 | Adversarial alternative

Our implementation of the adversarial technique of Louppe et al. [43] is based on GoDataDriven's code base. The following steps are alternated over 200 epochs: (1) train the adversarial network for a *single epoch*, holding the classifier network fixed, and (2) train the classifier network on a *single sampled mini batch*, holding the adversarial network fixed.

The adversarial network has 4 hidden layers with 32 hidden units in each. The ReLU activation was used throughout except in the final layer where the sigmoid function is used. The adversarial network was pretrained for 5 epochs. For the classifier network, we employed the same network as that of  $f_\theta$  in Algorithm 1 and kept all training choices, such as the optimisation algorithm and mini-batch size, the same. The classifier was pretrained for 2 epochs. Further implementation details for the adversarial approach can be found in the accompanying code.



**FIGURE 4** Pareto front estimation in the UCI (gender) test set. Each three by one block corresponds to a different fairness evaluation in Table 3. Within each block, the row corresponds to different Pareto front estimation procedures—“ATO one layer” refers to Algorithm 1 employing the ATO causal estimand in the penultimate layer as  $U_n$ , “ATO all layers” refers to Algorithm 1 with the ATO causal estimand calculated over all intermediate layers and lastly the adversarial approach. In each subplot, 1500 candidates are displayed along with the culled Pareto front in dashed magenta. The accuracy x-axis plots  $R(\theta; \hat{P}_{\text{test}})$  and the fairness y-axis plots  $U(\theta; \hat{P}_{\text{test}})$

## 5.7 | Results

In summary, we have three datasets and apply Algorithm 1 with two possibilities of  $U_n$ . We also apply the adversarial technique to each of the three datasets for comparison.

We present the results for the UCI (gender) dataset in Figure 4. The results for UCI (race) and recidivism can be found in Figure B1 and Figure B2 in Appendix A, respectively. Each three-by-one block in Figure 4 corresponds to a different fairness measure  $U(\theta; \hat{P}_{\text{test}})$  in Table 3. Within each block, the rows correspond to a different approach to finding the Pareto candidates.

Since 15 values of  $\lambda$  and 100 training–testing splits are considered, the output of Algorithm 1, CANDIDATES, consists in total 1500 learned network parameters. The 1500 experiments were run in parallel using high performance computing resources. It should be noted that the individual nodes used are less powerful than many personal workstations, and even so, a single training run takes only about 30 min.

The results of applying Algorithm 1 with  $U_n(\theta) = |\hat{\tau}_{\text{ATO}}(\mathbf{h}^{(L-1)})|$  and  $U_n(\theta) = \sum_{l=1}^{L-1} |\hat{\tau}_{\text{ATO}}(\mathbf{h}^{(l)})|$  are shown in the first and second rows of each three-by-one block

**TABLE 3** In visualizing the estimated Pareto front, we plot for each candidate  $\theta$  its value in the fairness-accuracy space  $R(\theta; \hat{P}_{\text{test}})$  versus  $U(\theta; \hat{P}_{\text{test}})$ , as given below for four different choices of  $U(\theta; \hat{P}_{\text{test}})$

ATO	$ \tau_{\text{ATO}}(f_{\theta}(\mathbf{x}); \hat{P}_{\text{test}}) $
Equal odds (EO)	$MV(f_{\theta}(\mathbf{x}), \mathbf{a}   \mathbf{y}; \hat{P}_{\text{test}})$
Equal opportunity (EOpp)	$MV(f_{\theta}(\mathbf{x}), \mathbf{a}   \mathbf{y} = 1; \hat{P}_{\text{test}})$
Demographic parity (DP)	$MV(f_{\theta}(\mathbf{x}), \mathbf{a}; \hat{P}_{\text{test}})$

Note: The development of the MV estimators are introduced in the Appendix B.

in Figure 4, respectively. In each sub-figure, we plot for each  $\theta \in \text{CANDIDATES}$ , its value in the fairness-accuracy space,  $R(\theta; \hat{P}_{\text{test}})$  versus  $U(\theta; \hat{P}_{\text{test}})$ . We further display the Pareto front culled from these 1500 Pareto candidates where the culling simply discards the dominated points.

The result of the adversarial approach is displayed in the third row of each three-by-one block of Figure 4. Fifteen values of  $\lambda$  in (9) along with 100 training–testing splits are considered, producing in total 1500  $\theta_{\text{clf}}$ , each of which

**TABLE 4** Each row corresponds to a different approach of finding the fairness-accuracy Pareto front

	UCI (gender)	UCI (race)	Recidivism
ATO one layer	44	33	89
ATO all layers	27	27	53
Adversarial	13	9	27

Note: Each column corresponds to a different dataset. The individual cells report the number of *nondominated* points (higher is better) calculated in the space  $(R(\theta, \hat{P}_{\text{test}}), U(\theta, \hat{P}_{\text{test}}))$  where

$$U(\theta, \hat{P}_{\text{test}}) = \left| \tau_{\text{ATO}}(f_{\theta}(\mathbf{x}); \hat{P}_{\text{test}}) \right|. \text{ ATO one layer refers to Algorithm 1}$$

where  $U_n$  is the ATO calculated on the penultimate layer, while ATO all layers refers to the ATO calculated over all intermediate layers.

is plotted in the fairness-accuracy space,  $R(\theta, \hat{P}_{\text{test}})$  versus  $U(\theta, \hat{P}_{\text{test}})$ .

We can immediately see from Figure 4 that compared to Algorithm 1, the adversarial approach is less capable of finding a Pareto front that spans the fairness-accuracy space. Indeed, Table 4 shows that the set of nondominated points found by the adversarial approach is much smaller relative to Algorithm 1.

We also observe better Pareto front estimation when Algorithm 1 is applied with  $U_n$  as the ATO causal measure calculated in the penultimate layer, compared to the ATO calculated over all layers. This seems to be true of the two other data settings as well, as can be seen in Appendix A. This suggests using (8) for  $U_n$  makes training the neural network more difficult. It may be worthwhile to explore training the network one layer at a time in future work.

Finally, from Figure 4, we can see that the demographic parity, equal odds, and equal opportunity fairness measures tend to form distinct clumps in the fairness-accuracy space, relative to the ATO causal measure. Results for UCI (race) and recidivism contained in Appendix A indicate much of the same conclusions. Appendix A also contains further visualization on the distributions of the prediction probabilities as  $\lambda$  in Algorithm 1 is dialed between 0 and 1.

## 6 | FUTURE WORK

We have found standardization of the objectives to be crucial as a preprocessing step. Recent works such as Grad-Norm [45] may offer a way to simultaneously standardize as we train rather than as a preprocessing step. We also discovered that the Chebyshev scalarization scheme, though superior to the linear scalarization scheme, is still seen to often produce dominated points. This may not be a particular fault of the Chebyshev scalarization scheme,

rather the phenomenon is likely due to the fact that we can never perform optimization perfectly, especially for deep neural networks using stochastic optimization algorithms. Barring this practical difficulty, as Chebyshev scalarization only leads to weakly Pareto optimal solutions, there still may be theoretical grounds for improvement. We may do better by using for example, the hypervolume indicator [46] to find points that are Pareto optimal. This has its own challenges, as the hypervolume indicator is expensive to calculate in higher dimensions.

Our experiments also reveal the difficulty of selecting weights in the Chebyshev scalarization scheme in such a way that would lead to a wide diversity of Pareto optimal points. In fact, many different weights often produced nearly identical solutions. Appropriately specifying the weights in a scalarization scheme without a priori information on the shape of the Pareto front itself is an open problem in multiobjective optimization. We had stated earlier that adaptive weights would be an improvement to the current methodology. This certainly merits further study, but the following challenges can be anticipated. Adapting the weights while training the neural network fundamentally changes the convergence behavior of the optimization algorithm since the optimization problem itself is gradually changing. It is already a nontrivial affair to train a deep neural network, and so we expect the addition of weight adaptation will require careful tinkering to the training process. Furthermore, many adaptive weight schemes seem to depend on the Pareto front shape itself. Without a better understanding of the Pareto front, we might do more harm by using an unsuitable weight adaptation technique.

## ACKNOWLEDGMENTS

Dr Susan Wei is the recipient of an Australian Research Council Discovery Early Career Award (project number DE200101253) funded by the Australian Government.

## CONFLICT OF INTEREST

The authors declare no potential conflict of interests.

## DATA AVAILABILITY STATEMENT

The data that support the findings of this study are openly available.

## ORCID

Susan Wei  <https://orcid.org/0000-0002-6842-2352>

## REFERENCES

1. M. T. Emmerich and A. H. Deutz, *A tutorial on multiobjective optimization: Fundamentals and evolutionary methods*, Nat. Comput. 17 (2018), no. 3, 585–609.

2. I. Das and J. E. Dennis, *A closer look at drawbacks of minimizing weighted sums of objectives for pareto set generation in multicriteria optimization problems*, *Struct. Optim.* 14 (1997), no. 1, 63–69. <https://doi.org/10.1007/BF01197559>
3. M. Hardt, E. Price, and N. Srebro, 2016: Equality of opportunity in supervised learning. *Proceedings of the 30th International Conference on Neural Information Processing Systems*, Curran Associates Inc., USA, NIPS'16, 3323–3331. <http://dl.acm.org/citation.cfm?id=3157382.3157469>
4. M. Joseph, M. Kearns, J. H. Morgenstern, and A. Roth, “Fairness in learning: Classic and contextual bandits,” *Advances in neural information processing systems* 29, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (eds.), Curran Associates, Inc., Barcelona, Spain, 2016, pp. 325–333. <http://papers.nips.cc/paper/6355-fairness-in-learning-classic-and-contextual-bandits.pdf>
5. M. B. Zafar, I. Valera, M. Gomez-Rodriguez, and K. P. Gummadi, 2017: Fairness constraints: Mechanisms for fair classification. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, AISTATS 2017, 20–22 April 2017, Fort Lauderdale, FL, USA, 962–970. <http://proceedings.mlr.press/v54/zafar17a.html>
6. M. B. Zafar, I. Valera, M. Gomez Rodriguez, and K. P. Gummadi, 2017: Fairness beyond disparate treatment &#38; disparate impact: Learning classification without disparate mistreatment. *Proceedings of the 26th International Conference on World Wide Web*, International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, WWW'17, 1171–1180. <https://doi.org/10.1145/3038912.3052660>
7. Y. Jin and B. Sendhoff, *Pareto-based multiobjective machine learning: an overview and case studies*, *IEEE Trans. Syst. Man Cyber. Part C* 38 (2008), no. 3, 397–415. <https://doi.org/10.1109/TSMCC.2008.919172>
8. R. Zhang and D. Golovin, 2020: Random hypervolume scalarizations for provable multi-objective black box optimization. *International Conference on Machine Learning*, PMLR, 11096–11105.
9. K. V. Moffaert and A. Nowe, *Multi-objective reinforcement learning using sets of Pareto dominating policies*, *J. Mach. Learn. Res.* 15 (2014), 3663–3692.
10. C. Li, M. Georgiopoulos, and G. C. Anagnostopoulos, *Pareto-path multitask multiple kernel learning*, *IEEE Trans. Neu. Netw. Learn. Syst.* 26 (2014), no. 1, 51–61.
11. A. Shah and Z. Ghahramani, “Pareto frontier learning with expensive correlated objectives,” *International conference on machine learning*, PMLR, New York, NY, 2016, pp. 1919–1927.
12. O. Sener and V. Koltun, *Multi-task learning as multi-objective optimization*, *Adv. Neural Inf. Proces. Syst.* (2018), 527–538.
13. X. Lin, H.-L. Zhen, Z. Li, Q.-F. Zhang and S. Kwong, *Pareto multi-task learning*, *Adv. Neural Inf. Process Syst.* Vancouver, CN, 2019, pp. 12060–12070.
14. F. Kamiran and T. Calders, *Data preprocessing techniques for classification without discrimination*, *Knowl. Inf. Syst.* 33 (2012), no. 1, 1–33. <https://doi.org/10.1007/s10115-011-0463-8>
15. M. Feldman, S. A. Friedler, J. Moeller, C. Scheidegger, and S. Venkatasubraan, 2015: Certifying and removing disparate impact. *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ACM, New York, NY, USA, KDD'15, 259–268. <https://doi.org/10.1145/2783258.2783311>
16. F. Calmon, D. Wei, B. Vinzamuri, K. Natesan Ramamurthy, and K. R. Varshney, 2017: Optimized pre-processing for discrimination prevention. *Advances in Neural Information Processing Systems* 30, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., Curran Associates, Inc., 3992–4001. <http://papers.nips.cc/paper/6988-optimized-pre-processing-for-discrimination-prevention.pdf>
17. J. E. Johndrow and K. Lum, *An algorithm for removing sensitive information: Application to race-independent recidivism prediction*, *Ann. Appl. Stat.* 13 (2019), no. 1, 189–220. <https://doi.org/10.1214/18-AOAS1201>
18. T. Calders and S. Verwer, *Three naive bayes approaches for discrimination-free classification*, *Data min. Knowl. Disc.* 21 (2010), no. 2, 277–292. <https://doi.org/10.1007/s10618-010-0190-x>
19. T. Kamishima, S. Akaho, and J. Sakuma, 2011: Fairness-aware learning through regularization approach. *Proceedings of the 2011 IEEE 11th International Conference on Data Mining Workshops*, IEEE Computer Society, Washington, DC, USA, ICDMW '11, 643–650. <https://doi.org/10.1109/ICDMW.2011.83>
20. M. B. Zafar, I. Valera, M. G. Rodriguez, K. P. Gummadi, and A. Weller, 2017: From parity to preference-based notions of fairness in classification. *Proceedings of the 31st International Conference on Neural Information Processing Systems*, Curran Associates Inc., Red Hook, NY, USA, NIPS'17, 228–238.
21. Y. Bechavod and K. Ligett, *Learning fair classifiers: A regularization-inspired approach*, *CoRR* (2017) <http://arxiv.org/abs/1707.00044>
22. A. Agarwal, A. Beygelzimer, M. Dudík, J. Langford, and H. M. Wallach, 2018: A reductions approach to fair classification. *Proceedings of the 35th International Conference on Machine Learning*, ICML 2018, Stockholm, Sweden, July 10–15, 2018, 60–69. <http://proceedings.mlr.press/v80/agarwal18a.html>
23. H. Narasimhan, 2018: Learning with complex loss functions and constraints. *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, A. Storkey and F. Perez-Cruz, Eds., PMLR, Playa Blanca, Lanzarote, Canary Islands, volume 84 of *Proceedings of Machine Learning Research*, 1646–1654. <http://proceedings.mlr.press/v84/narasimhan18a.html>
24. R. Berk, H. Heidari, S. Jabbari, M. Joseph, M. J. Kearns, J. Morgenstern, S. Neel, and A. Roth, *A convex framework for fair regression*, *CoRR* (2017) <http://arxiv.org/abs/1706.02409>
25. A. Beutel, J. Chen, Z. Zhao, and E. H. H. Chi, *Data decisions and theoretical implications when adversarially learning fair representations*, *CoRR* (2017) <http://arxiv.org/abs/1707.00075>
26. C. Wadsworth, F. Vera, and C. Piech, *Achieving fairness through adversarial learning: an application to recidivism prediction*, *CoRR* (2018) <http://arxiv.org/abs/1807.00199>
27. D. Madras, E. Creager, T. Pitassi, and R. Zemel, 2018: Learning adversarially fair and transferable representations. *Proceedings of the 35th International Conference on Machine Learning*, J. Dy and A. Krause, Eds., PMLR, Stockholm, Sweden, volume 80 of *Proceedings of Machine Learning Research*, 3384–3393. <http://proceedings.mlr.press/v80/madras18a.html>

28. P. Manisha and S. Gujar, *A neural network framework for fair classifier*, CoRR (2018) <http://arxiv.org/abs/1811.00247>
29. M. Ehrgott, *Multicriteria optimization*, Springer-Verlag, Lecture Notes in Economics and Mathematical Systems, 2000.
30. I. Giagkiozis and P. Fleming, *Methods for multi-objective optimization: an analysis*, Inform. Sci. 293 (2015), 338–350. <https://doi.org/10.1016/j.ins.2014.08.071> <http://www.sciencedirect.com/science/article/pii/S0020025514009074>
31. I. Das and J. Dennis, *Normal-boundary intersection: A new method for generating the pareto surface in nonlinear multicriteria optimization problems*, SIAM J. Optim. 8 (2000), 631–657. <https://doi.org/10.1137/S1052623496307510>
32. J. Kleinberg, 2018: Inherent trade-offs in algorithmic fairness. Abstracts of the 2018 ACM International Conference on Measurement and Modeling of Computer Systems, ACM, New York, NY, USA, SIGMETRICS'18, 40–40. <https://doi.org/10.1145/3219617.3219634>
33. Y. Ritov, Y. Sun, and R. Zhao, 2017: On conditional parity as a notion of non-discrimination in machine learning. To appear in Statistical Science.
34. P. J. Bickel, E. A. Hammel, and J. W. O'Connell, *Sex bias in graduate admissions: Data from Berkeley*, Science 187 (1975), no. 4175, 398–404.
35. H. Cui, R. Li, and W. Zhong, *Model-free feature screening for ultrahigh dimensional discriminant analysis*, J. Am. Stat. Assoc. 110–510 (2015), 630–641.
36. M. J. Kusner, J. Loftus, C. Russell, and R. Silva, “Counterfactual fairness,” *Advances in neural information processing systems 30*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), Curran Associates, Inc., Long Beach, CA, 2017, pp. 4066–4076. <http://papers.nips.cc/paper/6995-counterfactual-fairness.pdf>
37. N. Kilbertus, M. Rojas-Carulla, G. Parascandolo, M. Hardt, D. Janzing, and B. Schölkopf, 2017: Avoiding discrimination through causal reasoning. Proceedings of the 31st International Conference on Neural Information Processing Systems, Curran Associates Inc., USA, NIPS'17, 656–666. <http://dl.acm.org/citation.cfm?id=3294771.3294834>
38. J. R. Loftus, C. Russell, M. J. Kusner, and R. Silva, *Causal reasoning for algorithmic fairness*, CoRR (2018) <http://arxiv.org/abs/1805.05859>
39. A. Khademi, S. Lee, D. Foley, and V. Honavar, 2019: Fairness in algorithmic decision making: An excursion through the lens of causality. The World Wide Web Conference, ACM, New York, NY, USA, WWW '19, 2907–2914. <https://doi.org/10.1145/3308558.3313559>
40. F. Li, K. L. Morgan, and A. M. Zaslavsky, *Balancing covariates via propensity score weighting*, J. Am. Stat. Assoc. 113 (2018), no. 521, 390–400. <https://doi.org/10.1080/01621459.2016.1260466>
41. G. W. Imbens and D. B. Rubin, *Causal inference for statistics, social, and biomedical sciences: An Introduction*, Cambridge University Press, New York, NY, 2015.
42. C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, 2017: On calibration of modern neural networks. Proceedings of the 34th International Conference on Machine Learning - Volume 70, JMLR.org, ICML'17, 1321–1330. URL <http://dl.acm.org/citation.cfm?id=3305381.3305518>
43. G. Louppe, M. Kagan, and K. Cranmer, “Learning to pivot with adversarial networks,” *Advances in neural information processing systems 30*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), Curran Associates, Inc., Long Beach, CA, 2017, pp. 981–990 <http://papers.nips.cc/paper/6699-learning-to-pivot-with-adversarial-networks.pdf>
44. D. Kingma and J. Ba, 2014: Adam: A method for stochastic optimization. International Conference on Learning Representations.
45. Z. Chen, V. Badrinarayanan, C.-Y. Lee, and A. Rabinovich, 2018: Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. International Conference on Machine Learning, PMLR, 794–803.
46. A. Auger, J. Bader, D. Brockhoff, and E. Zitzler, *Hypervolume-based multiobjective optimization: Theoretical foundations and practical implications*, Theor. Comput. Sci. 425 (2012), 75–103. <https://doi.org/10.1016/j.tcs.2011.03.012> <http://www.sciencedirect.com/science/article/pii/S0304397511002428>

**How to cite this article:** S. Wei, and M. Niethammer, *The fairness-accuracy Pareto front*, Stat. Anal. Data Min.: ASA Data Sci. J. **15** (2022), 287–302. <https://doi.org/10.1002/sam.11560>

## APPENDIX A. ADDITIONAL FIGURES

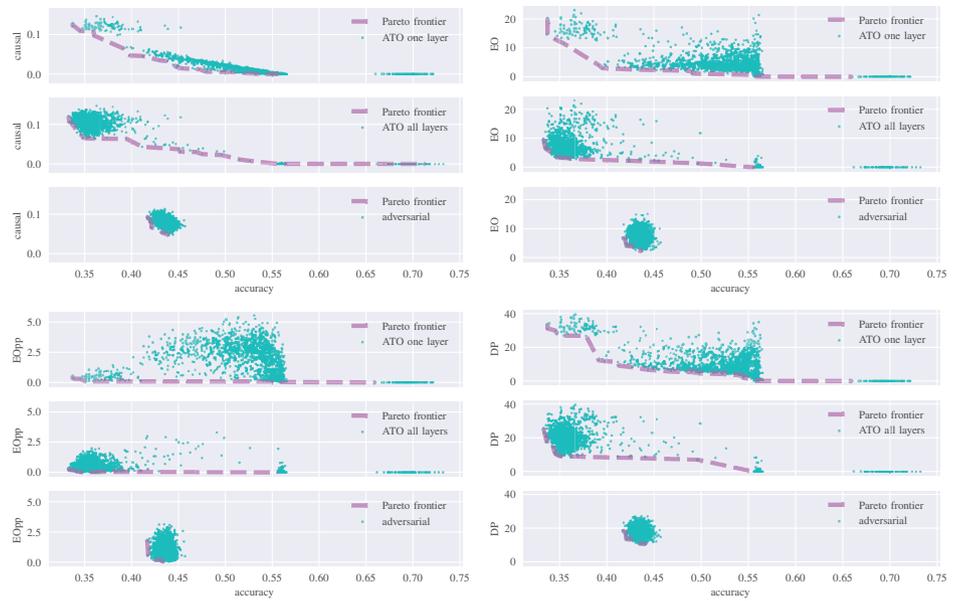
In this section, we provide additional figures for the experiments conducted in “The Fairness-Accuracy Pareto Front.” Figures B1 and B2 here are analogous to Figure 4 in the main text. The figures report, respectively, the results for the UCI (race) and recidivism dataset. Within each three-by-one block, the results of various Pareto estimation procedures are reported. Each block of subfigures corresponds to a different  $U(\theta; \hat{P}_{\text{test}})$ .

We also provide figures that help visualize the effect of dialing  $\lambda$  in Algorithm 1 of the main text from 0 to 1. In Figure B3, we display the distribution of the classifier’s prediction in the UCI (gender) dataset broken down by class label and sensitive attribute. Each panel of Figure B3 is a different  $\lambda$  value. In addition to reporting the ATO measure of fairness, we also indicate other noncausal fairness metrics including Equalized Odds, Equal Opportunity, and Demographic Parity. Similar visualization for the UCI (race) and recidivism dataset can be found in Figures B4 and B5, respectively.

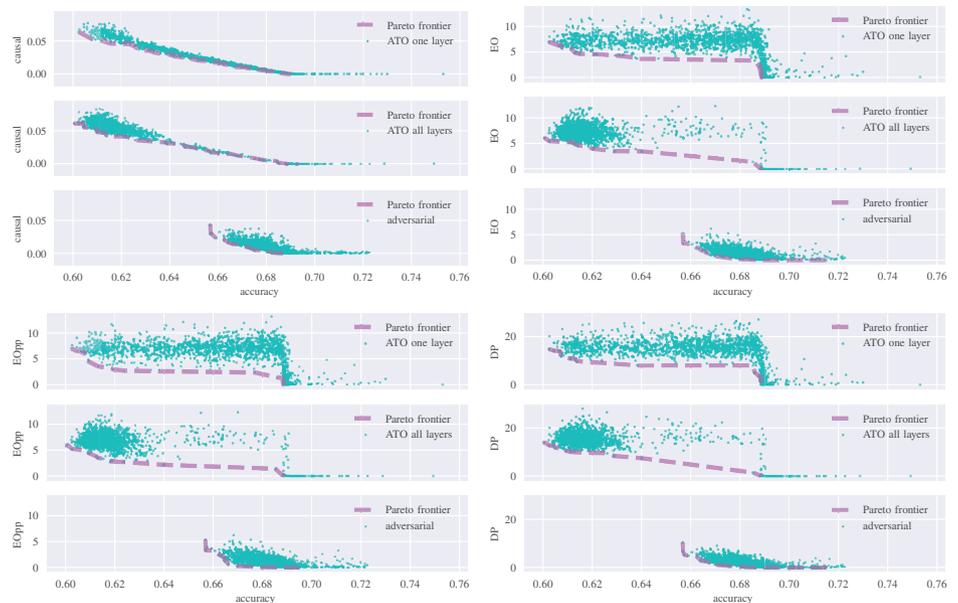
## APPENDIX B. ESTIMATION OF DEMOGRAPHIC AND CONDITIONAL PARITY

We describe a quantitative index that measures the degree to which conditional parity holds. First, recall we say a prediction score  $\hat{S}$  exhibits conditional parity with respect

**FIGURE B1** Pareto front estimation in the UCI (race) data set. Each three by one block corresponds to a fairness measure in Table 3 of the main document. Within each block, the row corresponds to different Pareto front estimation procedure—“ATO one layer” refers to Algorithm 1 in the main document employing the ATO causal estimand in the penultimate layer as  $U_n$ , “ATO all layers” refers to Algorithm 1 with the ATO causal estimand calculated over all layers and lastly the adversarial approach. In each subplot, 1500 candidates are displayed along with the culled Pareto front displayed in dashed magenta



**FIGURE B2** Pareto front estimation in the recidivism data set. Each three by one block corresponds to a fairness measure in Table 3 of the main document. Within each block, the row corresponds to different Pareto front estimation procedure—“ATO one layer” refers to Algorithm 1 in the main document employing the ATO causal estimand in the penultimate layer as  $U_n$ , “ATO all layers” refers to Algorithm 1 with the ATO causal estimand calculated over all layers and lastly the adversarial approach. In each subplot, 1500 candidates are displayed along with the culled Pareto front displayed in dashed magenta



to sensitive attribute  $Z$  conditioned on  $U$  if  $Z$  and  $U$  are independent conditional on  $U$ . We will limit the scope to categorical  $U$  and categorical  $Z$ . The conditional independence statement  $\hat{S} \perp\!\!\!\perp Z \mid U$  is typically assessed using the Cochran-Mantel-Haenszel test when  $\hat{S}$  and  $Z$  are binary variables and  $U$  is another, let us say  $k$ -level, categorical variable. Conditional independence in this case simply reduces to equality of odds ratios in each of the  $k$  two-by-two contingency tables.

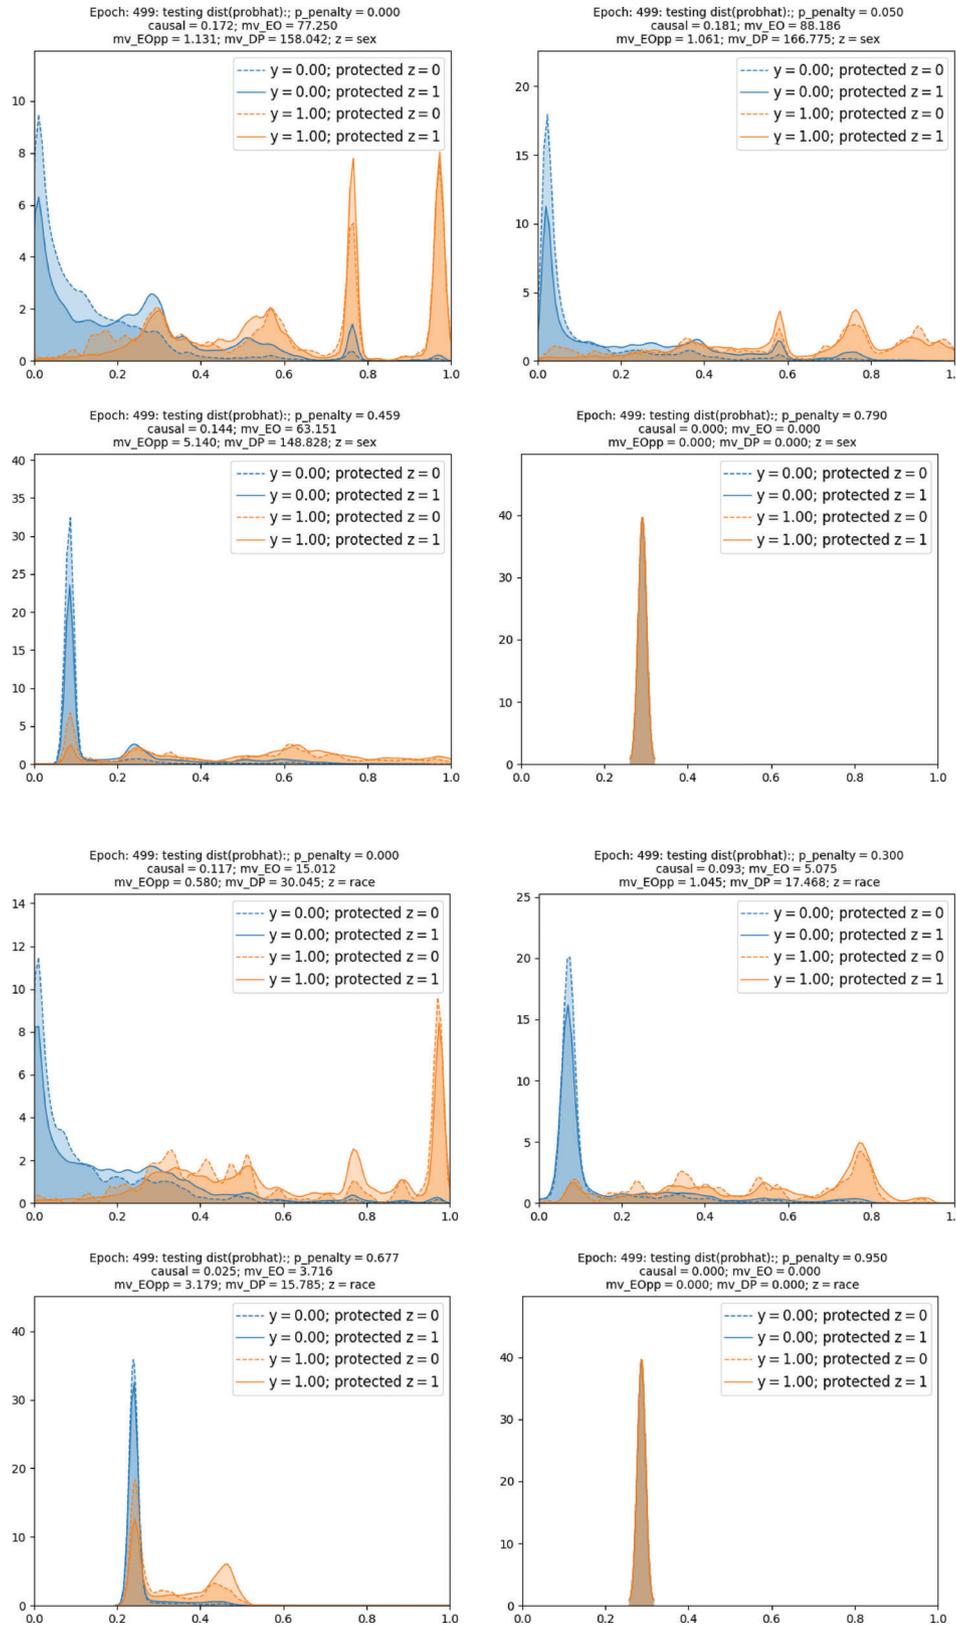
A more sophisticated method is required to assess conditional parity when  $\hat{S}$  is continuous. In the following, we introduce an index for conditional parity by adapting the mean–variance statistic of [35], which is not directly applicable as it is designed for *unconditional*

independence assessment between a continuous variable and a categorical variable.

### B.1. The mean–variance index

Without loss of generality, suppose the categorical random vector  $Z$  takes value in the set  $\{z_1, \dots, z_R\}$ . Let  $F_r(s) = P(\hat{S} \leq s \mid Z = z_r)$  be the conditional distribution of  $\hat{S}$  given  $Z = z_r$ . Let  $F(s) = P(\hat{S} \leq s)$  be the unconditional distribution of  $\hat{S}$ . The mean–variance index Cui et al. [35] is given by

$$MV = \sum_{r=1}^R P(Z = z_r) \int [F_r(s) - F(s)]^2 dF(s), \quad (\text{B1})$$



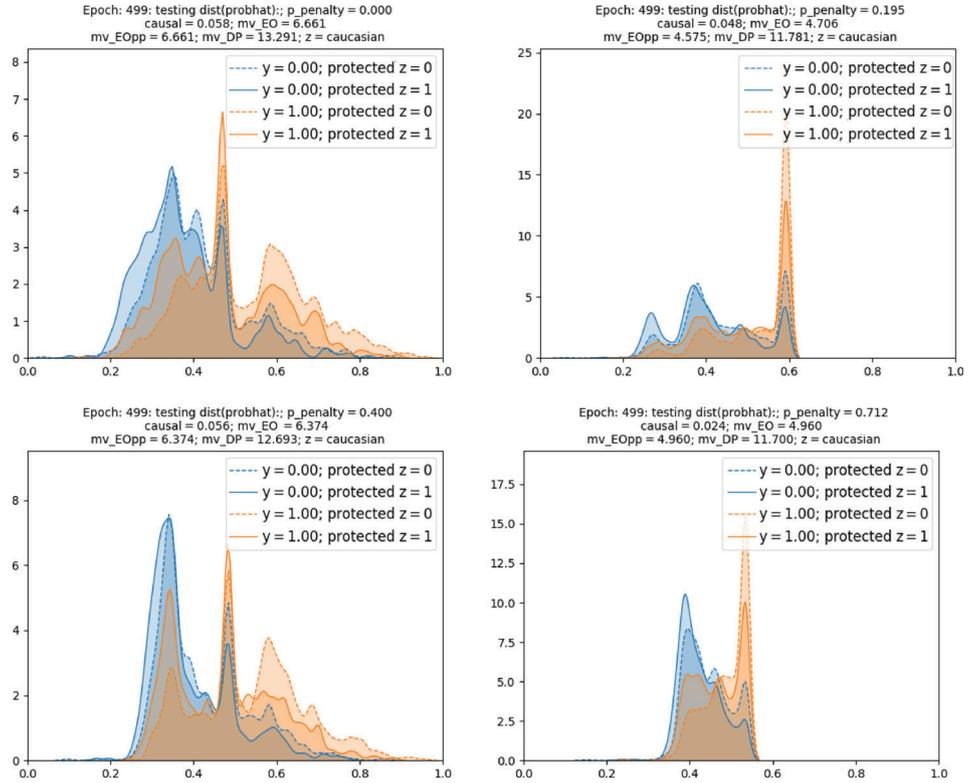
**FIGURE B3** For a particular training–testing split of the UCI (gender) dataset, we display the distributions of the predicted probabilities in the test set for four different values of  $\lambda$ , where  $\lambda$  is indicated by  $P\_PENALTY$  in the heading of each plot. The distributions are broken down by different values of the true target label  $y$  and the sensitive attribute  $z$ . Besides the ATO measure (CAUSAL), we also indicate equalized odds (MV\_EO), equality of opportunity (MV\_EOPP), and demographic parity (MV\_DP) in the headings of the subplots

**FIGURE B4** For a particular training–testing split of the UCI (race) dataset, we display the distributions of the predicted probabilities in the test set for four different values of  $\lambda$ , where  $\lambda$  is indicated by  $P\_PENALTY$  in the heading of each plot. The distributions are broken down by different values of the true target label  $y$  and the sensitive attribute  $z$ . Besides the ATO measure (CAUSAL), we also indicate equalized odds (MV\_EO), equality of opportunity (MV\_EOPP), and demographic parity (MV\_DP) in the headings of the subplots

where  $P(Z = z_r) > 0$  for all  $j = 1, \dots, r$ . Note that the integral in (B1) is simply the Cramér-von Mises distances between  $F_r$  and  $F$ , and so the mean–variance statistic is the weighted average of these distances, weighted by how

likely a particular value of  $Z$  is. If  $\hat{S}$  and  $Z$  are independent, then  $F_r(s) = F(s)$  for all  $r = 1, \dots, R$ . Thus the mean–variance index in (B1) has the salient property that it is zero if and only if  $\hat{S} \perp\!\!\!\perp Z$ .

**FIGURE B5** For a particular training–testing split of the recidivism dataset, we display the distributions of the predicted probabilities in the test set for four different values of  $\lambda$ , where  $\lambda$  is indicated by P\_PENALTY in the heading of each plot. The distributions are broken down by different values of the true target label  $y$  and the sensitive attribute  $z$ . Besides the ATO measure (CAUSAL), we also indicate equalized odds (MV\_EO), equality of opportunity (MV\_EOPP), and demographic parity (MV\_DP) in the headings of the subplots



Let the notation  $1\{\cdot\}$  denote indicator function of an event, that is, it is 1 if the event happens and 0 otherwise. The plug-in estimator for (B1) based on a sample  $\{(\hat{S}_i, Z_i), i = 1, \dots, n\}$  is

$$\widehat{MV} = \sum_{r=1}^R \hat{P}(Z = z_r) \left[ \frac{1}{n} \sum_{i=1}^n [\hat{F}_r(\hat{S}_i) - \hat{F}(\hat{S}_i)]^2 \right] \quad (\text{B2})$$

where

$$\begin{aligned} \hat{P}(Z = z_r) &= n^{-1} \sum_{i=1}^n 1\{Z_i = z_r\}, \\ \hat{F}_r(s) &= \frac{\sum_{i=1}^n 1\{\hat{S}_i \leq s, Z_i = z_r\}}{\sum_{i=1}^n 1\{Z_i = z_r\}}, \\ \hat{F}(s) &= \frac{1}{n} \sum_{i=1}^n 1\{\hat{S}_i \leq s\}. \end{aligned}$$

That the estimator in (B2) is consistent for its theoretical counterpart in (B1), that is,  $\widehat{MV}$  converges in probability to  $MV$  as  $n \rightarrow \infty$ , is established in Cui et al. [35].

### B.2.A new index to measure conditional parity

Now let us return to the assessment of  $\hat{S} \perp\!\!\!\perp Z \mid U$ . Suppose the categorical random vector  $U$  takes value in the set  $\{u_1, \dots, u_K\}$ . Then, we propose the following as a direct extension of the mean–variance statistic of Cui et al. [35]

to the conditional case:

$$MV^k = \sum_{r=1}^R P_k(Z = z_r) \int [F_{r,k}(s) - F_k(s)]^2 dF_k(s) \quad (\text{B3})$$

where

$$\begin{aligned} P_k(Z = z_r) &= P(Z = z_r \mid U = u_k) \\ F_{r,k}(s) &= P(\hat{S} \leq s \mid Z = z_r, U = u_k) \\ F_k(s) &= P(\hat{S} \leq s \mid U = u_k). \end{aligned}$$

We employ the following statistic to assess the degree to which  $\hat{S} \perp\!\!\!\perp Z \mid U$  holds:

$$\max_{k=1, \dots, K} MV^k(\hat{S}, Z). \quad (\text{B4})$$

Since  $MV^k$  is zero if and only if  $\hat{S}$  is independent of  $Z$  conditioned on  $U$  and the same is true for the maximum. In other words, (B4) is zero if and only if  $\hat{S} \perp\!\!\!\perp Z \mid U$ .

Since (B3) and (B4) are both unknown population quantities, they require estimation. Let  $\widehat{MV}_k$  be the plug-in estimator for (B3) based on the empirical measure. Namely, let

$$\widehat{MV}^k(\hat{S}, Z) = \sum_{r=1}^R \hat{P}_k(Z = z_r) \int [\hat{F}_{r,k}(s) - \hat{F}_k(s)]^2 d\hat{F}_k(s)$$

where

$$\begin{aligned}\hat{P}_k(Z = z_r) &= \frac{\sum_{i=1}^n 1 \{Z_i = z_r, U_i = u_k\}}{\sum_{i=1}^n 1 \{U_i = u_k\}} \\ \hat{F}_{r,k}(s) &= \frac{\sum_{i=1}^n 1 \{\hat{S}_i \leq s, Z_i = z_r, U_i = u_k\}}{\sum_{i=1}^n 1 \{Z_i = z_r, U_i = u_k\}} \\ \hat{F}_k(s) &= \frac{\sum_{i=1}^n 1 \{\hat{S}_i \leq s, U_i = u_k\}}{\sum_{i=1}^n 1 \{U_i = u_k\}}.\end{aligned}$$

To estimate (B4), we simply take the maximum over  $k$ :

$$\max_{k=1, \dots, K} \widehat{MV}^k. \quad (\text{B5})$$

We will call the quantity in (B5) the **conditional mean-variance index**. Note that this measure will be zero if conditional independence holds and will increase with increasing dependence. Hence, it is a suitable measure to assess conditional parity.