
Fair Bayesian Optimization

Valerio Perrone, Michele Donini, Krishnaram Kenthapadi, Cédric Archambeau

Amazon

{vperrone, donini, kenthk, cedrica}@amazon.com

Abstract

Given the increasing importance of machine learning (ML) in our lives, algorithmic fairness techniques have been proposed to mitigate biases that can be amplified by ML. Commonly, these specialized techniques apply to a single family of ML models and a specific definition of fairness, limiting their effectiveness in practice. We introduce a general constrained Bayesian optimization (BO) framework to optimize the performance of any ML model while enforcing one or multiple fairness constraints. BO is a global optimization method that has been successfully applied to automatically tune the hyperparameters of ML models. We apply BO with fairness constraints to a range of popular models, including random forests, gradient boosting, and neural networks, showing that we can obtain accurate and fair solutions by acting solely on the hyperparameters. We also show empirically that our approach is competitive with specialized techniques that explicitly enforce fairness constraints during training, and outperforms preprocessing methods that learn unbiased representations of the input data. Moreover, our method can be used in synergy with such specialized fairness techniques to tune their hyperparameters. Finally, we study the relationship between hyperparameters and fairness of the generated model. We observe a correlation between regularization and unbiased models, explaining why acting on the hyperparameters leads to ML models that generalize well and are fair.

1 Introduction

With the increasing use of machine learning (ML) in domains such as financial lending, hiring, criminal justice, and college admissions, there has been a major concern for the potential for ML to unintentionally encode societal biases and result in systematic discrimination [4, 5, 6, 7, 9]. For example, a classifier that is only tuned to maximize performance can unfairly predict a high credit risk for some subgroups of the population applying for a loan. Extensive work has been done to measure and mitigate biases during different stages of the ML life-cycle [5].

In many practical ML settings, one needs to optimize the performance of ML models in a black-box manner while enforcing fairness constraints. For example, several cloud platforms allow users to bring their own proprietary model training code and data, perform model training, and tune the model hyperparameters treating them as a black-box [22].¹ However, being tailored to specific models and fairness definitions, most existing fairness techniques are inapplicable to these settings. Motivated by this, we present a general constrained Bayesian Optimization (BO) framework to tune the performance of ML models while satisfying fairness constraints. We demonstrate its effectiveness on several classes of ML models, including random forests, gradient boosting, and neural networks, showing that we can obtain accurate and fair models simply by acting on their hyperparameters. Figure 1 illustrates

¹SigOpt (<https://sigopt.com/>); Cloud AutoML (<https://www.blog.google/products/google-cloud/cloud-automl-making-ai-accessible-every-business/>); Optuna (<https://optuna.org/>); Amazon AMT (<https://docs.aws.amazon.com/sagemaker/latest/dg/automatic-model-tuning.html>).

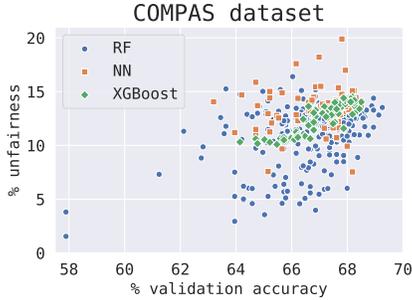


Figure 1: Unfairness-accuracy trade-off by varying the hyperparameters of XGBoost, RF, and NN on a recidivism prediction task. Each dot corresponds to a different hyperparameter configuration. For a given level of accuracy, models with very different levels of unfairness can be generated simply by changing the model hyperparameters.

this idea by plotting the accuracy and unfairness levels achieved by trained gradient boosted tree ensembles (XGBoost [13]), random forests (RF) and fully-connected feed-forward neural networks (NN), with each dot corresponding to a random hyperparameter configuration. The key observation is that, given a level of accuracy, one can reduce unfairness just by tuning the model hyperparameters. As an example, paying 0.02 points of accuracy (-2.5%, from 0.69 to 0.67) can give a RF model with 0.08 fewer unfairness points (-70%, from 0.13 to 0.04).

Our methodology supports arbitrary fairness definitions, allows for multiple constraints to be enforced simultaneously, and is complementary to existing bias mitigation techniques. Experiments show that our black-box approach is more effective than preprocessing techniques that learn unbiased representations of the input data, and is competitive with methods that have access to the model internals and incorporate fairness constraints as part of the objective during model training.

The paper is organized as follows. Section 2 presents the main ideas of algorithmic fairness, such as state-of-the-art methods and common definitions of fairness. Section 3 introduces our model-agnostic methodology to optimize model hyperparameters while satisfying fairness constraints. The experimental results in Section 4 show that our method compares favourably with state-of-the-art techniques. We also analyze the importance of hyperparameters controlling regularization in the search for fair and accurate models. Finally, Section 5 presents conclusions and future directions.

2 Algorithmic Fairness

The goal of algorithmic fairness is to develop machine learning methods that are accurate and fair. There is extensive work on identifying and measuring the extent of discrimination (e.g., [4, 9]), and on mitigation approaches in the form of fairness-aware algorithms (e.g., [8, 11, 14, 16, 18, 19, 23, 27, 30, 41, 43, 45, 46]). We can divide these methods into three main families. Methods in the first family modify a pre-trained model to make it less biased while trying to keep performance high (i.e., post-processing the model) [17, 23, 36]. The second family consists of methods that enforce fairness constraints during training (e.g., [2, 14, 43, 44]). The third family of methods achieves fairness by modifying the data representation (i.e., pre-processing the data), and then applying standard machine learning methods [10, 45]. All algorithmic fairness methods require a measure of fairness to be defined *a priori*, and most of them only work with one fairness measure at a time.

2.1 Fairness Definitions

Today, there is no consensus on a unique definition of fairness, and some of the most common definitions are conflicting [40]. Our goal is not to introduce yet another fairness definition, but to propose a flexible methodology that is able to output fair models regardless of the selected criterion we want to enforce. As we will show, in our black-box framework we can seamlessly incorporate different definitions either independently or simultaneously.

Let Y be the true label (binary), S the protected (or sensitive) attribute (binary), and \hat{Y} the predicted label. The most common definitions can be grouped into three categories: (i) considering the predicted outcome given the true label; (ii) considering the true label given predicted outcome; (iii) considering the predicted outcome only. The following are examples of the most used definitions.

Equal Opportunity (EO): requires equal True Positive Rates (TPR) across subgroups, i.e., $P(\hat{Y} = 1|Y = 1, S = 0) = P(\hat{Y} = 1|Y = 1, S = 1)$;

Equalized Odds (EOdd): requires equal False Positive Rates (FPR), in addition to EO;

Statistical Parity (SP): requires positive predictions to be unaffected by the value of the protected attribute, regardless of the actual true label, i.e., $P(\hat{Y} = 1|S = 0) = P(\hat{Y} = 1|S = 1)$.

Our goal is to find accurate models with a controlled (small) violation of the fairness constraint. Hence, following [14], we consider the family of the ϵ -fair models. A model \hat{Y} is ϵ -fair if it violates the fairness definition by at most $\epsilon \geq 0$. In the case of EO, a model \hat{Y} is ϵ -fair if the difference in EO (DEO) is at most ϵ :

$$|P(\hat{Y} = 1|Y = 1, S = 0) - P(\hat{Y} = 1|Y = 1, S = 1)| \leq \epsilon. \quad (1)$$

For EOdd, we have two different constraints simultaneously. The first one is equivalent to DEO, and the second one is the difference of FPR (DFP). Finally, we can similarly define the difference in SP (DSP):

$$|P(\hat{Y} = 1|S = 0) - P(\hat{Y} = 1|S = 1)| \leq \epsilon. \quad (2)$$

The inherent trade-offs underlying different notions of fairness has been studied extensively [16, 18, 31]. Carefully picking the correct definition of fairness for the problem at hand has critical importance, and cannot be delegated to an automatic agent. Instead, a human decision is required (e.g., with human-in-the-loop approaches [42]).

2.2 Black-box Algorithmic Fairness

Black-box approaches to enforce fairness have been proposed in the literature, usually consisting of data pre-processing or model post-processing techniques. Pre-processing methods aim to change the data representation to make it less biased with respect to one or more sensitive attributes. For example, [45] learns a fair representation of the data on top of any training procedure. This is achieved by solving an optimization problem with a two-fold goal: encode the data by preserving as much information as possible and obfuscate the membership to specific subgroups. Common practice is also to apply the following two steps [29]: (i) remove the sensitive attribute from the feature set, and (ii) rebalance the dataset, i.e., increase the number of observations using synthetic oversampling via SMOTE [12]. Pre-processing methods are black-box, but their hyperparameters, as well as the ones of the underlying base methods, still need to be tuned for performance. As we will show shortly, this can still impact the fairness and accuracy of the returned solution.

Other approaches are described in [2, 3], where a fair classification task is reduced to a sequence of cost-sensitive classification problems. The solutions to these problems yield a randomized classifier with low empirical error subject to fairness constraints. The connection between randomized classifiers and fairness (and consequently the connection to differential privacy) has also been studied [33]. In the context of empirical risk minimization algorithms, these methods are black-box ones with respect to the base model, but still need specific implementations based on the fairness definition at hand and output an ensemble of models.² In contrast, we show that our constrained BO approach is agnostic, both to the underlying base methods and to the selected fairness constraint. Moreover, our proposal can be used in synergy with these methods to tune their hyperparameters.

The second group of black-box methods are post-processing techniques. These include adjusting the threshold of the learned classifier to make the model more fair with respect to a given fairness definition while retaining high accuracy. One of the most important methods in this family is [23], which also introduces the concept of EO. This method is black-box with respect to the base model but only works with specific statistical definitions of fairness. The idea is to optimally adjust any learned model to mitigate discrimination by adding a flipped predicted label with a certain probability, which is adjusted to minimize unfairness. This can be viewed as tossing a biased coin to enforce a certain amount of positive discrimination (also known as affirmative action in public policy) to mitigate the negative bias in the original model. The main drawback is that post-processing techniques are inherently sub-optimal. They are allowed to act only on the previously learned information, without generally collecting new data or re-using the original data. The assumption is that the model is already trained and the original data eventually discarded, which is not the setting we focus on in this work.

²Fairlearn code at: <https://github.com/fairlearn/fairlearn>.

3 Fair Bayesian Optimization

Bayesian optimization (BO) is a well-established methodology to optimize expensive black-box functions (see [38] for an overview). It relies on a probabilistic model of the unknown target $f(\mathbf{x})$ one wishes to optimize. The black-box $f(\mathbf{x})$ is repeatedly queried until one runs out of budget (e.g., time). Queries consist of evaluations of f at hyperparameter configurations $\mathbf{x}^1, \dots, \mathbf{x}^n$ selected according to an explore-exploit trade-off criterion or *acquisition function* [28]. The hyperparameter configuration corresponding to the best query is then returned. A popular approach is to impose a Gaussian process (GP) prior over f and then compute the posterior GP based on the observed queries $f(\mathbf{x}^1), \dots, f(\mathbf{x}^n)$ [37]. The posterior GP is characterized by a posterior mean and a posterior variance function that are required when evaluating the acquisition function for each new query of f .

A widely used acquisition function is the Expected Improvement (EI) [32]. This is defined as the expected amount of improvement of an evaluation with respect to the current minimum $f(\mathbf{x}_{min})$. For a Gaussian predictive distribution, EI is defined in closed-form as

$$EI(\mathbf{x}) = \mathbf{E}[\max(0, f(\mathbf{x}_{min}) - f(\mathbf{x}))] = \sigma^2(\mathbf{x})(z(\mathbf{x})\Phi_n(z(\mathbf{x})) + \phi_n(z(\mathbf{x}))),$$

where $z(\mathbf{x}) := \frac{\mu(\mathbf{x}) - f(\mathbf{x}_{min})}{\sigma^2(\mathbf{x})}$. Here, μ and σ^2 are respectively the posterior GP mean and variance, and Φ_n and ϕ_n respectively the CDF and PDF of the standard normal. Alternative acquisition functions based on information gain criteria have also been developed [24]. Standard acquisitions focus only on the objective $f(\mathbf{x})$ and do not account for additional constraints. In this work, we aim to optimize a black-box function $f(\mathbf{x})$ subject to fairness constraints $c(\mathbf{x}) \leq \epsilon$, $\epsilon \in \mathbb{R}$, with ϵ determining how strictly the corresponding fairness definition should be enforced.

We next describe Fair Bayesian Optimization (FairBO), our approach to optimize the hyperparameters of a black-box function while satisfying arbitrary fairness constraints (Algorithm 1). We leverage the constrained EI (cEI), an established acquisition function to extend BO to the constrained case [20, 21, 39]. We place an additional GP on the fairness constraint and weight the EI with the posterior probability of the constraint being satisfied, giving $cEI(\mathbf{x}) = EI(\mathbf{x})P(c(\mathbf{x}) \leq \epsilon)$. In our setting, feasible hyperparameter configurations are those satisfying the desired fairness constraint (e.g., the DSP across subgroups should be lower than 0.1). We define $cEI(\mathbf{x})$ with respect to the current *fair* best, which may not be available in the initial iterations. Therefore, we start by greedily optimizing $P(c(\mathbf{x}) \leq \epsilon)$ and then switch to $cEI(\mathbf{x})$ (lines 4-10) when the first fair hyperparameter configuration is found.

FairBO is straightforward to extend to handle D fairness constraints simultaneously, each with its own upper bound $\{\epsilon_d\}_{d=1}^D$. One option is to merge the fairness constraints into a single binary feedback encoding whether all constraints are satisfied. Assuming independence, one can alternatively place a fairness model $\{c_d(\mathbf{x})\}_{d=1}^D$ on every fairness constraint and let $P(\mathbf{x}) = \prod_{d=1}^D P(c_d(\mathbf{x}) \leq \epsilon_d)$, each term being the probability of satisfying a fairness constraint. FairBO can also be easily implemented through alternative, entropy-based acquisition functions [25, 35]. We leave an empirical comparison with this variant for future work.

Algorithm 1 FairBO

Input: Initial and total budgets T_0, T ; unfairness bound ϵ ; GP prior on objective $f(\mathbf{x})$ and fairness $c(\mathbf{x})$ model.

- 1: Evaluate $f(\mathbf{x}_i)$ and $c(\mathbf{x}_i)$ for $i = 1, \dots, T_0$ hyperparameters \mathbf{x}_i from the search space (e.g., drawn uniformly at random or from a fixed initial design) and set the used budget $t = T_0$.
 - 2: Define the set of evaluated hyperparameters $\mathcal{C} = \{(\mathbf{x}_i, f(\mathbf{x}_i), c(\mathbf{x}_i))\}_{i=1}^{T_0}$
 - 3: Compute the posterior GP for the objective and the fairness models based on \mathcal{C} .
 - 4: **while** $t < T$ **do**
 - 5: $\mathbf{x}_{new} = \arg \max_{\mathbf{x}} EI(\mathbf{x})P(c(\mathbf{x}) \leq \epsilon)$.
 - 6: Evaluate $f(\mathbf{x}_{new})$ and $c(\mathbf{x}_{new})$.
 - 7: Update $\mathcal{C} = \mathcal{C} \cup \{(\mathbf{x}_{new}, f(\mathbf{x}_{new}), c(\mathbf{x}_{new}))\}$
 - 8: Compute the posterior GP for the objective and the fairness models based on \mathcal{C} .
 - 9: $t = t + 1$
 - 10: **end while**
 - 11: **return** Best fair hyperparameter configuration in \mathcal{C} .
-

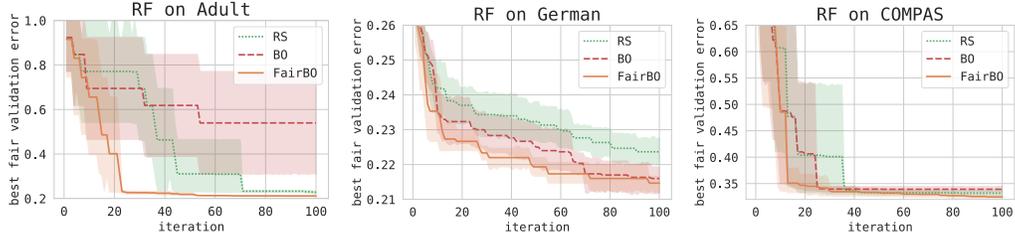


Figure 2: Comparison of RS, BO, and FairBO over the validation error (vertical axis) of the best feasible solution on Adult, German, and COMPAS, tuning RF as base model. The fairness constraint is $DSP \leq 0.05$. FairBO finds a more accurate fair model in fewer iterations (horizontal axis).

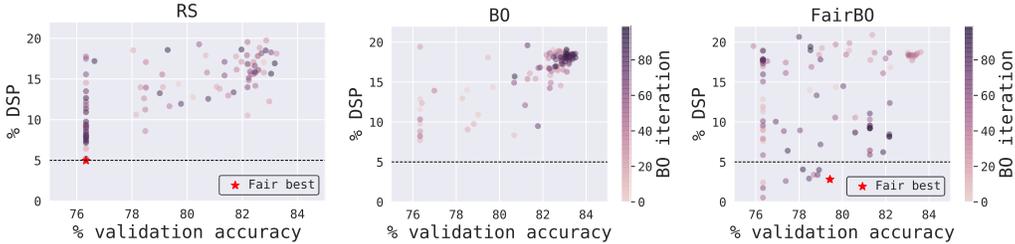


Figure 3: Comparison of RS, BO, and FairBO on the task of tuning RF on Adult. The horizontal line is the fairness constraint, set to $DSP \leq 0.05$, and darker dots correspond to later BO iterations. Standard BO can get stuck in high-performing yet unfair regions, failing to return a feasible solution. RS is more robust than BO, but only finds a fair model with low validation accuracy.

4 Experiments

We consider three datasets widely used in the context of fairness: (i) Adult – Census Income [15], a binary classification task with binary gender as sensitive attribute, where the task is to predict if income exceeds \$50K/yr based on census data; (ii) German Credit Data [15], a binary classification problem with binary gender as sensitive attribute, where the goal is to classify people described by a set of attributes as good or bad credit risks; (iii) COMPAS, a binary classification problem concerning recidivism risk, with binarized ethnic group as sensitive attribute (one group for “white” and one for all other ethnic groups).³ We tune four popular ML algorithms implemented in `scikit-learn` [34]: XGBoost, Random Forest (RF), a fully-connected neural network (NN), and Linear Learner (LL), optimizing the hyperparameters in Appendix A. We optimize for validation accuracy, with a random 70%/30% split into train/validation, and place an upper bound on unfairness (e.g., defined via DSP as per inequality (2)). All hyperparameter optimization methods are initialized with 5 random hyperparameter configurations. BO and FairBO are implemented in GPyOpt [1], with the GP using a Matérn-5/2 covariance kernel with automatic relevance determination hyperparameters, optimized by type-II maximum likelihood [37]. Results are averaged across 10 repetitions, with 95% confidence intervals obtained via bootstrapping. All experiments are run on AWS with `m4.xlarge` machines.

4.1 FairBO performance

We first compare FairBO to Random Search (RS) and standard BO (based on the EI acquisition). Figure 2 compares the validation error of the fair solution found on Adult, German, and COMPAS with a fairness constraint of $DSP \leq 0.05$. As expected, FairBO finds an accurate and fair model more quickly than RS and BO. On Adult, FairBO reaches the fair (local) optimum five times faster than RS. Figure 3 shows an example run with all tried hyperparameter configurations. Standard BO can get stuck in high-performing yet unfair regions, failing to return a feasible solution. While RS is more robust, it only finds a fair solution with the same accuracy of the trivial model always predicting the majority class (i.e., the set of points with accuracy ≈ 0.763). Analogous results are shown in

³COMPAS link: <https://github.com/propublica/compas-analysis>.

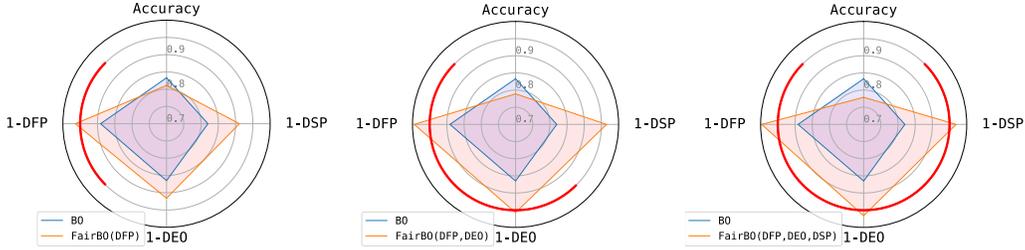


Figure 4: Best RF hyperparameter configuration found on Adult by BO and FairBO with progressively more fairness constraints, represented by the red arches (left: $\text{DFP} \leq 0.05$; center: $\{\text{DFP} \leq 0.05, \text{DEO} \leq 0.05\}$; right: $\{\text{DFP} \leq 0.05, \text{DEO} \leq 0.05, \text{DSP} \leq 0.05\}$). Compared to BO, FairBO can trade off accuracy for a fairer solution, and can do so with respect to all definitions simultaneously.

Appendix B for XGBoost and NN, as well as with a looser fairness constraint $\text{DSP} \leq 0.15$ and a fairness constraint on DEO (inequality (1)), noting that other fairness definitions can be plugged in.

4.2 Multiple fairness constraints

In contrast to most algorithmic fairness techniques, FairBO can seamlessly handle multiple fairness definitions simultaneously. We consider 100 iterations of standard BO and FairBO on the problem of tuning RF on Adult, progressively adding more fairness constraints. Specifically, we first impose a constraint on DFP, then on both DFP and DEO, and finally on DFP, DEO, and DSP together. All constraint thresholds are set to 0.05 and results are averaged over 10 independent repetitions. Figure 4 shows accuracy and three fairness metrics (i.e., one minus unfairness, namely, $(1 - \text{DFP})$, $(1 - \text{DEO})$, $(1 - \text{DSP})$ respectively) of the returned fair solution for RF. Analogous results for XGBoost, NN, and LL are given in Appendix B. Interestingly, FairBO allows us to trade off relatively little accuracy for a more fair solution, which gets progressively more fair as we add more constraints.

4.3 Hyperparameters and fairness

We showed that hyperparameter tuning can mitigate unfairness effectively. We now investigate more closely the role of each hyperparameter on the unfairness of the resulting model. For each algorithm, we apply fANOVA [26] to study hyperparameter importance on fairness, defined as DSP (analogous results with DEO are given in Appendix B). Hyperparameter configurations and unfairness metrics are collected from 100 iterations of random search and 10 random seeds, for a total of 1000 data points per algorithm-dataset pair. Figure 5 indicates that the hyperparameters controlling the regularization level tend to have the largest impact on fairness. In the case of RF, the most important hyperparameter is the maximum tree depth; for XGBoost, this is either the L1 weight regularizer α or the number of boosting rounds; for NN, Adam’s initial learning rate eps plays the biggest role (as we keep the number of epochs fixed). Finally, for LL the most relevant hyperparameter influencing fairness is precisely the regularization factor α . Figure 6 shows the DSP and accuracy for 100 random hyperparameter configurations for each algorithm, before and after fixing the most relevant hyperparameter detected by fANOVA. As expected, fixing these hyperparameters limits the ability of FairBO to provide fair and accurate solutions, leading to fewer fair solutions.

We conjecture that, by preventing overfitting, the hyperparameters controlling regularization generate models with a lower ability to discriminate among the different values of the sensitive attribute. For example, consider the simple case in which the sensitive feature is uncorrelated with the other features. Assuming we have a linear model w , where the entry $w_s \in \mathbb{R}$ is the weight assigned to the sensitive feature, we can bound the DSP as follows:

$$\text{DSP}(w) = |P_x(w \cdot x \geq 0 | S = 0) - P_x(w \cdot x \geq 0 | S = 1)| \leq P_x(-w_s \leq w \cdot x < 0 | S = 0).$$

The idea is that, given x in subgroup $S = 0$ and being x' the same as x where we flipped the value of the sensitive feature to $S = 1$, we have $w \cdot x - w \cdot x' = -w_s$. Consequently, a smaller w_s helps obtain a less biased model. Indeed, unfairness is correlated with the weight assigned to the sensitive feature (or to the sum of the weights assigned to all the features correlated with the sensitive one), and regularization tends to alleviate this. Increasing the regularization in a cross-entropy loss will have a similar effect, generating models progressively less data dependent and steering DSP closer to zero.

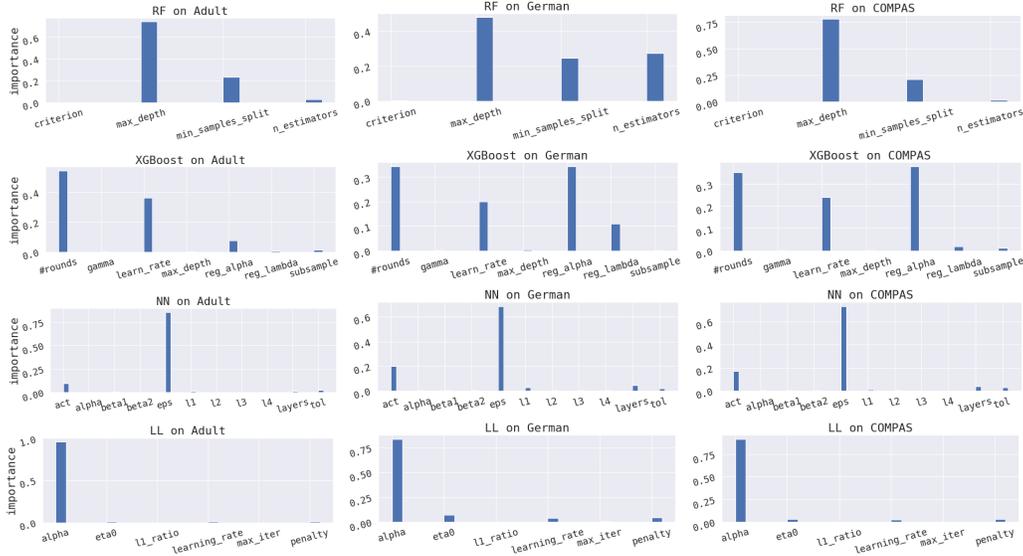


Figure 5: Hyperparameter importance on fairness: the role of each tuned hyperparameter on the unfairness (DSP) of the resulting model is evaluated with fANOVA [26]. Statistics are collected from 100 iterations of random search and 10 seeds. Regularization hyperparameters tend to impact fairness the most (e.g., the most relevant hyperparameter in LL is precisely the regularization factor α).

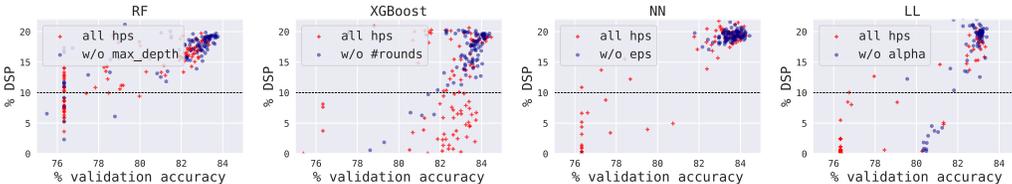


Figure 6: Unfairness vs Accuracy on Adult for 100 random hyperparameter configurations of RF, XGBoost, NN, and LL. For each algorithm, we either vary all hyperparameters or fix the one most correlated with fairness (detected by fANOVA [26]). As expected, fixing these hyperparameters limits the ability of FairBO to provide fair and accurate solutions, yielding fewer fair configurations.

4.4 Model-agnostic and model-specific techniques

In the context of algorithmic fairness several *ad-hoc* methods have been proposed. We compare to the method of Zafar [43], Adversarial debiasing [46], and Fair Empirical Risk Minimization (FERM) [14].⁴ These methods enforce fairness during training and optimize the parameters of a linear model to make it both accurate and fair with respect to a fixed fairness definition. These methods are not model-agnostic and only apply to linear models. As alternative black-box approaches we compare to SMOTE, which preprocesses the data by removing the sensitive feature and rebalancing the observations, as well as FERM preprocessing [14], which learns a fair representation of the data before fitting a linear model. We allocate 100 hyperparameter tuning iterations for all approaches.

Table 1 shows the best fair model found by FairBO on LL compared to the best fair model found by each baseline. As expected, FERM achieves higher accuracy, due to the constraint applied directly while training the parameters (as opposed to the hyperparameters) of the linear model. However, the gap in performance with FairBO is modest, and FairBO outperforms both Zafar and Adversarial Debiasing. While conceptually simple, FairBO emerges as a surprisingly competitive baseline that can outperform or compete against these highly specialized techniques. We note that all model-specific techniques tend to find solutions that are more fair than the required constraint. FairBO is

⁴Code for Zafar from <https://github.com/mbilalzafar/fair-classification>; code for Adversarial Debiasing from <https://github.com/IBM/AIF360>; code for FERM from https://github.com/jmikko/fair_ERM.

Table 1: Validation error of the best fair models for model-specific (first three rows) and model-agnostic fairness methods. We use the fairness constraint, $DSP \leq 0.1$.

| Method | Adult | German | COMPAS |
|-----------------|-------------------|-------------------|-------------------|
| FERM | 0.164 ± 0.010 | 0.185 ± 0.012 | 0.285 ± 0.009 |
| Zafar | 0.187 ± 0.001 | 0.272 ± 0.004 | 0.411 ± 0.063 |
| Adversarial | 0.237 ± 0.001 | 0.227 ± 0.008 | 0.327 ± 0.002 |
| FERM preprocess | 0.228 ± 0.013 | 0.231 ± 0.015 | 0.343 ± 0.002 |
| SMOTE | 0.178 ± 0.005 | 0.206 ± 0.004 | 0.321 ± 0.002 |
| FairBO (ours) | 0.175 ± 0.007 | 0.196 ± 0.005 | 0.307 ± 0.001 |

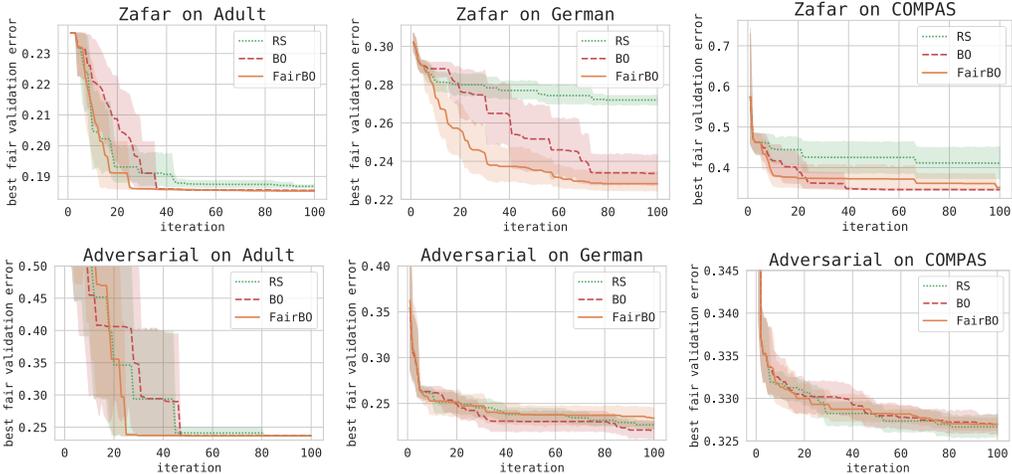


Figure 7: Comparison of RS, BO, and FairBO where Zafar and Adversarial Debiasing are used as base learners. We use the fairness constraint, $DSP \leq 0.1$. Hyperparameter tuning on top of model-specific techniques helps find better performing fair solutions, and the FairBO approach tends to do so in fewer iterations (horizontal axis) than RS and BO.

also the best model-agnostic method, outperforming both SMOTE and FERM preprocessing. This shows that we can remove bias with a smaller impact on accuracy.

As FairBO only acts on the hyperparameters, it can be used *on top* of model-specific techniques, which come with their own hyperparameters. Blindly tuning these hyperparameters can negatively impact the fairness of the resulting solution. We demonstrate this by combining FairBO with Zafar and Adversarial Debiasing, which we found to be sensitive to their hyperparameter settings (unlike FERM). Figure 7 shows that hyperparameter tuning on top of model-specific techniques yields better performing fair solutions, and FairBO tends to find them more quickly than random search and standard BO. In other words, FairBO is the method of choice when automating the tuning of alternative fairness techniques, finding superior fair solutions.

5 Conclusions

We showed that tuning model hyperparameters is surprisingly effective to mitigate unfairness in ML and proposed FairBO, a constrained Bayesian optimization framework to jointly tune ML models for accuracy and fairness. FairBO is model agnostic, can be used with arbitrary fairness definitions, and allows for multiple fairness definitions to be applied simultaneously. The proposed methodology empirically finds more accurate fair solutions than data-debiasing techniques, while being competitive with state-of-the-art algorithm-specific fairness techniques. We also showed that FairBO is preferable over standard BO when tuning the hyperparameters of specialized techniques. Finally, we demonstrated the importance of regularization hyperparameters in yielding fair and accurate models. Potential directions for future work include applying our framework to regression, image recognition, and natural language processing problems, and covering more complex fairness definitions, such as with continuous sensitive attributes.

Broader Impact

Algorithmic fairness has the potential for a profound impact on society. This paper aims to make it safer to use automatic agents that generate decisions affecting critical domains such as justice, financial landing, and hiring. We believe that simplifying the process of training accurate and fair models can help spread good practice in our field and foster the generation of unbiased models. Our work pursues exactly this goal.

More fair machine learning is needed in our society, especially in light of the many discoveries of negative bias in commonly used machine learning models. With less biased and more fair machine learning, we can improve the trust of our automatic agents and increase awareness on this topic among the colleagues in our community. But to maximize impact, fair machine learning also needs to be accessible to non-experts. Ultimately, we have the possibility to enhance the benefits that machine learning can bring to society without translating human biases to the learned models.

We are aware that statistical measures of fairness, such as statistical parity or equal opportunity, cannot be considered as the unique definitions. Indeed, any definition of fairness applied to the task at hand has to be carefully understood and chosen by a human, not by an automatic agent. It is well-known that some of the definitions are in contrast with each other so that, by enforcing one, we are simultaneously forcing other definitions to be violated. The choice of the right definition is fundamental but out of the scope of our proposal, and requires a human-in-the-loop approach.

References

- [1] GPyOpt: A Bayesian optimization framework in Python. <http://github.com/SheffieldML/GPyOpt>, 2016.
- [2] A. Agarwal, A. Beygelzimer, M. Dudik, J. Langford, and H. Wallach. A reductions approach to fair classification. *ICML*, 2018.
- [3] A. Agarwal, M. Dudik, and Z. S. Wu. Fair regression: Quantitative definitions and reduction-based algorithms. *ICML*, 2019.
- [4] J. Angwin, J. Larson, S. Mattu, and L. Kirchner. Machine bias. *ProPublica*, 2016.
- [5] S. Barocas, M. Hardt, and A. Narayanan. Fairness and machine learning. URL: www.fairmlbook.org, 2018.
- [6] T. Bolukbasi, K. Chang, J. Y. Zou, V. Saligrama, and A. T. Kalai. Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. *NeurIPS*, 2016.
- [7] J. Buolamwini and T. Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. *FAT**, 2018.
- [8] T. Calders, F. Kamiran, and M. Pechenizkiy. Building classifiers with independency constraints. *IEEE ICDM*, 2009.
- [9] A. Caliskan, J. J. Bryson, and A. Narayanan. Semantics derived automatically from language corpora contain human-like biases. *Science*, 2017.
- [10] F. Calmon, D. Wei, B. Vinzamuri, K. N. Ramamurthy, and K. R. Varshney. Optimized pre-processing for discrimination prevention. *NeurIPS*, 2017.
- [11] L. E. Celis, D. Straszak, and N. K. Vishnoi. Ranking with fairness constraints. *ICALP*, 2018.
- [12] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer. SMOTE: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 2002.
- [13] T. Chen and C. Guestrin. XGBoost: A scalable tree boosting system. *ACM SIGKDD*, 2016.
- [14] M. Donini, L. Oneto, S. Ben-David, J. S. Shawe-Taylor, and M. Pontil. Empirical risk minimization under fairness constraints. *NeurIPS*, 2018.
- [15] D. Dua and C. Graff. UCI machine learning repository, 2017.
- [16] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel. Fairness through awareness. *ITCS*, 2012.
- [17] M. Feldman, S. A. Friedler, J. Moeller, C. Scheidegger, and S. Venkatasubramanian. Certifying and removing disparate impact. *ACM SIGKDD*, 2015.

- [18] S. A. Friedler, C. Scheidegger, and S. Venkatasubramanian. On the (im)possibility of fairness. *arXiv preprint arXiv:1609.07236*, 2016.
- [19] S. A. Friedler, C. Scheidegger, S. Venkatasubramanian, S. Choudhary, E. P. Hamilton, and D. Roth. A comparative study of fairness-enhancing interventions in machine learning. *FAT**, 2019.
- [20] J. Gardner, M. Kusner, Z. Xu, K. Weinberger, and J. Cunningham. Bayesian optimization with inequality constraints. *ICML*, 2014.
- [21] M. A. Gelbart, J. Snoek, and R. P. Adams. Bayesian optimization with unknown constraints. *UAI*, 2014.
- [22] D. Golovin, B. Solnik, S. Moitra, G. Kochanski, J. E. Karro, and D. Sculley. Google Vizier: A service for black-box optimization. *KDD*, 2017.
- [23] M. Hardt, E. Price, and N. Srebro. Equality of opportunity in supervised learning. *NeurIPS*, 2016.
- [24] P. Hennig and C. J. Schuler. Entropy search for information-efficient global optimization. *JMLR*, 2012.
- [25] J. M. Hernández-Lobato, M. A. Gelbart, M. W. Hoffman, R. P. Adams, and Z. Ghahramani. Predictive entropy search for Bayesian optimization with unknown constraints. *ICML*, 2015.
- [26] F. Hutter, H. Hoos, and K. Leyton-Brown. An efficient approach for assessing hyperparameter importance. *ICML*, 2014.
- [27] S. Jabbari, M. Joseph, M. Kearns, J. Morgenstern, and A. Roth. Fairness in reinforcement learning. *ICML*, 2017.
- [28] D. R. Jones, M. Schonlau, and W. J. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 1998.
- [29] F. Kamiran and T. Calders. Data preprocessing techniques for classification without discrimination. *KAIS*, 2012.
- [30] T. Kamishima, S. Akaho, H. Asoh, and J. Sakuma. Fairness-aware classifier with prejudice remover regularizer. *ECML PKDD*, 2012.
- [31] J. Kleinberg. Inherent trade-offs in algorithmic fairness. *SIGMETRICS*, 2018.
- [32] J. Mockus, V. Tiesis, and A. Zilinskas. The application of Bayesian methods for seeking the extremum. *Towards Global Optimization*, 1978.
- [33] L. Oneto, M. Donini, M. Pontil, and J. Shawe-Taylor. Randomized learning and generalization of fair and private classifiers: From PAC-Bayes to stability and differential privacy. *Neurocomputing*, 2020.
- [34] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *JMLR*, 2011.
- [35] V. Perrone, I. Shcherbatyi, R. Jenatton, C. Archambeau, and M. Seeger. Constrained Bayesian optimization with max-value entropy search. *arXiv preprint arXiv:1910.07003*, 2019.
- [36] G. Pleiss, M. Raghavan, F. Wu, J. Kleinberg, and K. Q. Weinberger. On fairness and calibration. *NeurIPS*, 2017.
- [37] C. Rasmussen and C. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- [38] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas. Taking the human out of the loop: A review of Bayesian optimization. *IEEE*, 2016.
- [39] J. Snoek, O. Rippel, K. Swersky, R. Kiros, N. Satish, N. Sundaram, M. Patwary, M. Prabhat, and R. Adams. Scalable Bayesian optimization using deep neural networks. *ICML*, 2015.
- [40] S. Verma and J. Rubin. Fairness definitions explained. *FairWare*, 2018.
- [41] B. Woodworth, S. Gunasekar, M. I. Ohanessian, and N. Srebro. Learning non-discriminatory predictors. *COLT*, 2017.
- [42] M. Yaghini, H. Heidari, and A. Krause. A human-in-the-loop framework to construct context-dependent mathematical formulations of fairness. *arXiv preprint arXiv:1911.03020*, 2019.

- [43] M. B. Zafar, I. Valera, M. Gomez-Rodriguez, and K. P. Gummadi. Fairness constraints: Mechanisms for fair classification. *AISTATS*, 2017.
- [44] M. B. Zafar, I. Valera, M. Gomez-Rodriguez, and K. P. Gummadi. Fairness constraints: A flexible approach for fair classification. *JMLR*, 2019.
- [45] R. Zemel, Y. Wu, K. Swersky, T. Pitassi, and C. Dwork. Learning fair representations. *ICML*, 2013.
- [46] B. H. Zhang, B. Lemoine, and M. Mitchell. Mitigating unwanted biases with adversarial learning. *AIES*, 2018.

A Optimized hyperparameters

A.1 Algorithms

We considered the problem of tuning four popular ML algorithms, as implemented in `scikit-learn`: XGBoost (XGB), Random Forest (RF), Neural Network (NN), Linear Learner (LL). In this section, we give more details on the search space over which each hyperparameter was optimized.

XGBoost We consider a 7-dimensional search space: number of boosting rounds in $\{1, 2, \dots, 256\}$ (log scaled), learning rate in $[0.01, 1.0]$ (log scaled), minimum loss reduction to partition leaf node γ in $[0.0, 0.1]$, L1 weight regularization α in $[10^{-3}, 10^3]$ (log scaled), L2 weight regularization λ in $[10^{-3}, 10^3]$ (log scaled), subsampling rate in $[0.01, 1.0]$, maximum tree depth in $\{1, 2, \dots, 16\}$.

RF We consider a 4-dimensional search space: number of trees in $\{1, 2, \dots, 64\}$ (log scaled), tree split threshold in $[0.01, 0.5]$ (log scaled), tree maximum depth in $\{1, 2, 3, 4, 5\}$, criterion for quality of split in {Gini, Entropy}.

NN We consider an 11-dimensional search space: number of layers in $\{1, 2, 3, 4\}$, each layer size in $\{2, 3, \dots, 32\}$ (log scaled), activation in {Logistic, Tanh, ReLU}, tolerance in $[10^{-5}, 10^{-2}]$ (log scaled), L2 regularization in $[10^{-6}, 0.1]$ (log scaled), and Adam parameters: initial learning rate ϵ in $[10^{-6}, 10^{-2}]$ (log scaled), β_1 and β_2 in $[10^{-3}, 0.99]$ (log scaled).

LL We consider a 6-dimensional search space: iteration count in $\{1, 2, \dots, 128\}$, regularization type in {L1, L2, ElasticNet}, Elastic Net mixing parameter in $[0, 1]$, regularization factor α in $[10^{-3}, 10^3]$ (log scaled), initial learning rate η_0 in $[10^{-4}, 0.1]$ (log scaled), learning rate schedule in {Constant, Optimal, Invscaling, Adaptive}.

A.2 Baselines

In addition to the hyperparameters of the tuned algorithms, we allocated 100 iterations of random search to tune each baseline, namely FERM, Zafar, Adversarial, FERM preprocess and SMOTE.

FERM We considered 2 hyperparameters for FERM. The L2 regularization coefficient C in $[0.001, 10.0]$ (log scaled), and FERM’s epsilon-fairness threshold in $[0.0, 1.0]$.

Zafar We tuned 3 hyperparameters: L1 regularization coefficient in $[10^{-3}, 10.0]$ (log scaled), L2 regularization coefficient in $[0.001, 10.0]$ (log scaled), and Zafar’s epsilon-fairness threshold in $[0.0, 1.0]$.

Adversarial We tuned 4 different hyperparameters: the adversary loss weight in $[0.0, 1.0]$, the number of epochs for training in $\{1, 2, \dots, 500\}$ (log scaled), the batch size in $\{32, 33, \dots, 256\}$ (log scaled), and the number of hidden units of the network between in $\{20, 21, \dots, 5000\}$ (log scaled).

FERM preprocess FERM preprocessing learns a fair representation of the dataset, which is then fed to LL. Hence, we tuned the same 6 hyperparameters as per the original LL.

SMOTE In addition to the 6 hyperparameters of LL, we jointly tuned 2 hyperparameters controlling the degree of dataset rebalancing: oversampling rate of the less frequent class in $[0.3, 1.0]$ and number of neighbors to generate synthetic examples in $\{1, 2, \dots, 20\}$.

B Additional Experiments

Compared to Random Search (RS) and standard BO, FairBO explores the fair regions of the hyperparameter space more quickly and tends to return a more accurate, fair solution. This appendix presents additional experiments with different fairness thresholds, definitions and algorithms. We also study the impact of each hyperparameter on the deviation from equal opportunity (DEO) of the resulting model.

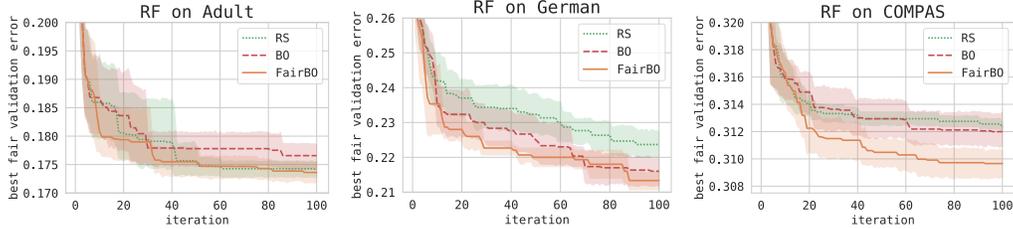


Figure 8: Comparison of RS, BO, and FairBO over the validation error (vertical axis) of the best feasible solution on Adult, German, and COMPAS data using RF as base model. The fairness constraint is set to a looser value of $DSP \leq 0.15$. FairBO tends to find a well-performing fair model in fewer iterations than RS and BO.

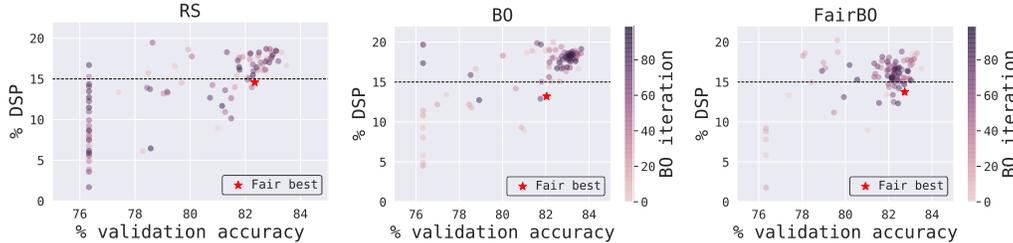


Figure 9: Comparison of FairBO, BO, and RS when tuning RF on Adult. The horizontal line is the fairness constraint, set to a looser value of $DSP \leq 0.15$. Darker dots correspond to later BO iterations. Both RS and BO find a less accurate fair solution after 100 iterations. Additionally, BO can get stuck in high-performing yet unfair regions.

Fairness thresholds We investigate the impact of varying the fairness threshold ϵ . We repeat the experiment of tuning a RF model on Adult, German and COMPAS, this time with a looser fairness constraint of $DSP \leq 0.15$. Figure 8 shows the validation error of the fair solution on the three datasets. As expected, the performance gap between FairBO and the baselines is still clear but overall less pronounced compared to the experiments with a stricter fairness constraint $DSP \leq 0.05$. Additionally, due to the looser constraint, the accuracy of the best fair solution is significantly higher in Adult and COMPAS. At the same time, FairBO still tends to find a well-performing fair model more quickly than RS and BO. Figure 9 illustrates the behavior on Adult, indicating that standard BO can still get stuck in high-performing but unfair regions. Although BO also finds a fair solution, this is less accurate than the one found by FairBO. RS also tends to require more resources to find an accurate and fair solution.

Fairness definitions FairBO can be applied directly to arbitrary fairness definitions. In this section, we repeat the RF experiments by replacing the initial, strict constraint on statistical parity ($DSP \leq 0.05$) with an analogous constraint on equal opportunity ($DEO \leq 0.05$). As previously, 100 iterations are allocated to RS, BO and FairBO. Figure 10 shows the best validation error of the fair solution found after each hyperparameter evaluation on the three datasets. Consistently with the DSP constraint, FairBO finds fair and accurate solutions more quickly than BO and RS. In addition, constraining DEO on COMPAS allows for fair solutions with higher accuracy compared to constraining DSP. This is not surprising as DEO is generally a better proxy of accuracy than DSP (a perfect classifier has DEO equal to zero, unlike DSP).

FairBO can also handle multiple fairness definitions simultaneously. As in the experiments with RF in the main paper, we consider 100 iterations of standard BO and FairBO on the problem of tuning NN, XGBoost and LL on Adult while satisfying the three fairness definitions. Specifically, we impose that DFP, DEO and DSP should all be less than 0.05. Results are averaged over 10 independent repetitions. Figure 11 shows accuracy and three fairness metrics of the returned fair solutions, with the red arches indicating the constraints. FairBO allows us to trade off a relatively small degree of accuracy to get a more fair solution. Interestingly, when FairBO is applied to XGBoost, the fair solution comes with the slightest accuracy loss.

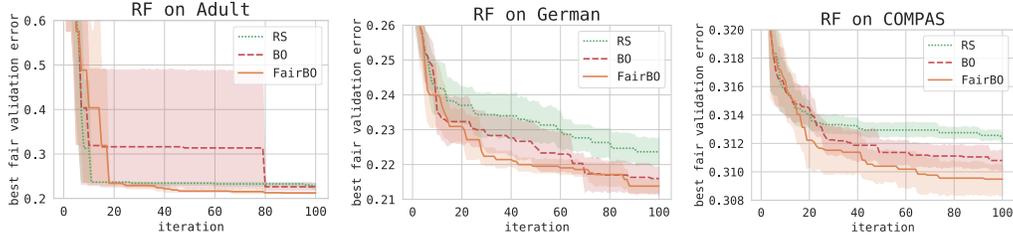


Figure 10: Comparison of RS, BO, and FairBO over the validation error (vertical axis) of the best feasible solution on Adult, German, and COMPAS data using RF as base model. The fairness constraint is set to $DEO \leq 0.05$. FairBO tends to find a well-performing fair model more with fewer iterations than RS and BO.

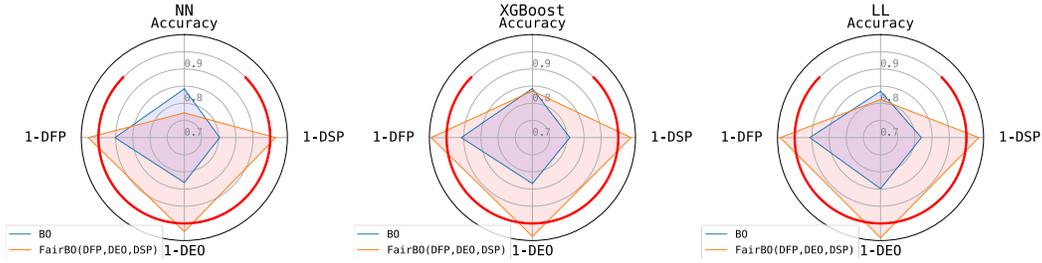


Figure 11: Best hyperparameter configuration found on Adult by BO and FairBO on NN, XGBoost, and LL with three simultaneous fairness constraints, represented by the red arches: $\{DFP < 0.05, DEO < 0.05, DSP < 0.05\}$. Compared to BO, FairBO can trade off accuracy for a fairer solution with respect to each definition. When applied to XGBoost, a slight accuracy degradation is enough to find solutions that are fair across all definitions.

Tuning XGboost and NN Figure 12 compares RS with BO and FairBO on the problem of tuning XGboost and a NN on Adult, COMPAS, and German. Solutions are considered fair if they satisfy the strict requirement of $DSP \leq 0.05$. As in the case of RF, FairBO tends to find an accurate and fair solution faster than baselines. Consistently with the multiple constraint experiment, tuning XGBoost allows for more accurate fair solutions than RF and NN.

Hyperparameters and equal opportunity Previous experiments shed light on the role of hyperparameter tuning on unfairness. We now study the contribution of each hyperparameter when unfairness is defined as difference in equal opportunity (DEO). As in the DSP experiments in the main paper, for each algorithm we study hyperparameter importance via fANOVA. Hyperparameter configurations and unfairness metrics are collected from 100 iterations of random search and 10 random seeds for each dataset. Figure 13 shows hyperparameter importance on DEO, confirming the results obtained with DSP and indicating that the hyperparameters controlling regularization tend to play the largest role. In the case of RF, the most important hyperparameter is the maximum tree depth on Adult and COMPAS, and number of trees on German; for XGBoost, the L1 weight regularizer α , number of boosting rounds or learning rate are the most important; for NN, the most impactful is Adam’s initial learning rate ϵ (as the number of epochs is kept fixed to the default value); finally, the most relevant hyperparameter in LL is, once again, exactly the regularization factor α .

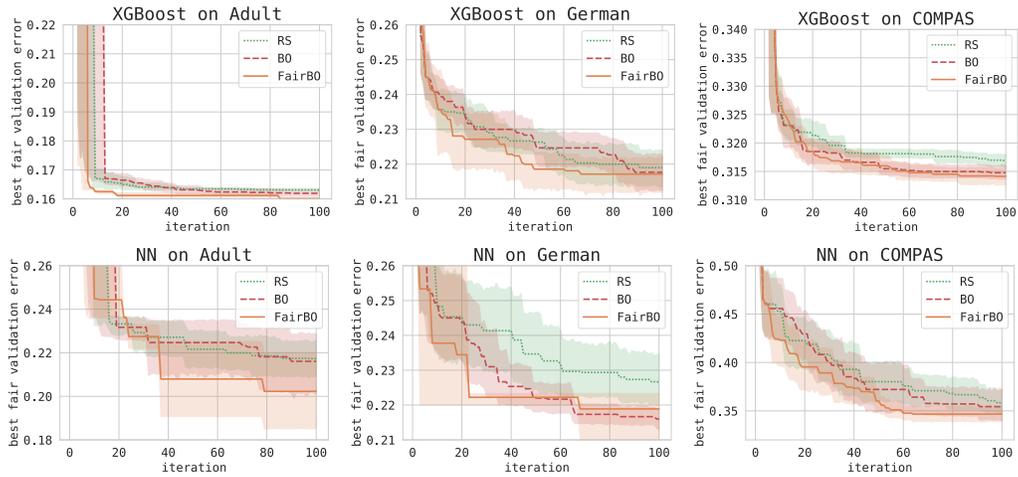


Figure 12: Comparison of RS, BO, and FairBO over the validation error (vertical axis) of the best feasible solution on Adult, German, and COMPAS data using XGBoost and NN as base models. The fairness constraint is set to $DSP \leq 0.05$. FairBO tends to find a well-performing fair model in fewer iterations than RS and BO.

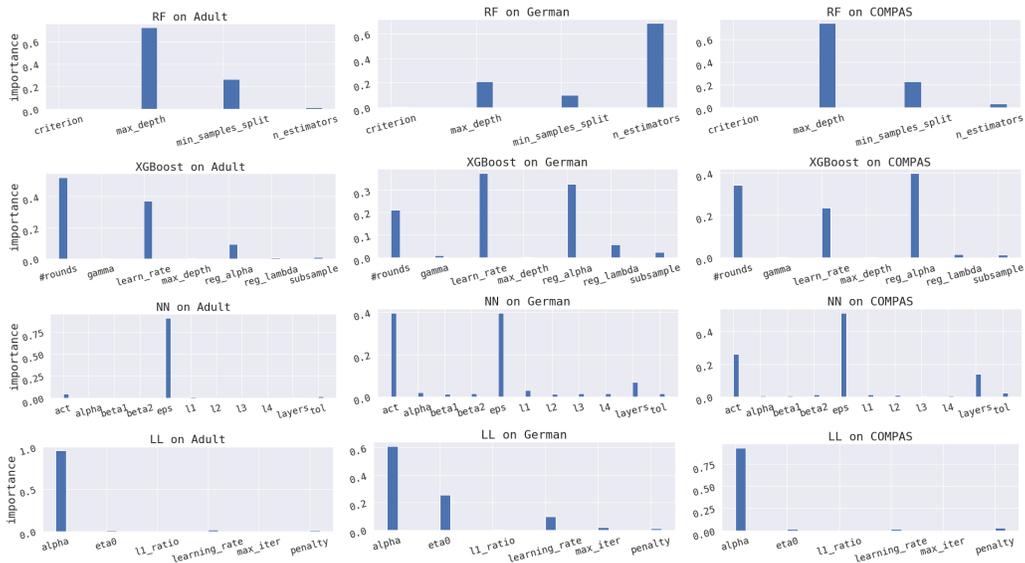


Figure 13: Hyperparameter importance on equal opportunity. The role of each hyperparameter on the unfairness (DEO) of the resulting model is evaluated via fANOVA. Statistics are collected from 100 iterations of random search and 10 seeds. Regularization hyperparameters tend to impact fairness the most (e.g., LL's most important hyperparameter is precisely the regularization factor alpha).