A Bandit-Based Algorithm for Fairness-Aware Hyperparameter Optimization

André F. Cruz
Feedzai
andre.cruz@feedzai.com

Pedro Saleiro
Feedzai
pedro.saleiro@feedzai.com

Catarina Belém
Feedzai
catarina.belem@feedzai.com

Carlos Soares
Universidade do Porto
csoares@fe.up.pt

Pedro Bizarro
Feedzai
pedro.bizarro@feedzai.com

Abstract

Considerable research effort has been guided towards algorithmic fairness but there is still no major breakthrough. In practice, an exhaustive search over all possible techniques and hyperparameters is needed to find optimal fairness-accuracy trade-offs. Hence, coupled with the lack of tools for ML practitioners, real-world adoption of bias reduction methods is still scarce. To address this, we present Fairband, a bandit-based fairness-aware hyperparameter optimization (HO) algorithm. Fairband is conceptually simple, resource-efficient, easy to implement, and agnostic to both the objective metrics, model types and the hyperparameter space being explored. Moreover, by introducing fairness notions into HO, we enable seamless and efficient integration of fairness objectives into real-world ML pipelines. We compare Fairband with popular HO methods on four real-world decision-making datasets. We show that Fairband can efficiently navigate the fairness-accuracy trade-off through hyperparameter optimization. Furthermore, without extra training cost, it consistently finds configurations attaining substantially improved fairness at a comparatively small decrease in predictive accuracy.

1 Introduction

Artificial Intelligence (AI) has increasingly been used to aid decision-making in sensitive domains, including healthcare (Rajkomar et al., 2019), criminal justice (Berk et al., 2018), and financial services (Board, 2017). These algorithmic decision-making systems are accumulating societal responsibilities, often without human oversight. At the same time, Machine Learning (ML) models are usually optimized for a single global metric of predictive accuracy (*e.g.*, binary cross-entropy loss on the training set), without taking into account possible side-effects and their real-world implications.

One potential side-effect is algorithmic bias, *i.e.*, disparate predictive and error rates across subgroups of the population, hurting individuals based on ethnicity, age, gender, or any other sensitive attribute (Angwin et al., 2016; Bartlett et al., 2019; Buolamwini and Gebru, 2018). This has several causes, from historical biases encoded in the data, to misrepresented populations in data samples, noisy labels, development decisions (*e.g.*, missing values imputation), or simply the nature of learning under severe class-imbalance (Suresh and Guttag, 2019).

Algorithmic fairness (Kleinberg et al., 2018) is an emerging sub-field in AI that aims at reducing bias in decision-making systems. Although a focus of extensive research in recent years, there are still no major breakthroughs in automatic bias reduction techniques (Friedler et al., 2019). Additionally, existing bias reduction techniques only target specific stages of the ML pipeline (*e.g.*, data sampling,
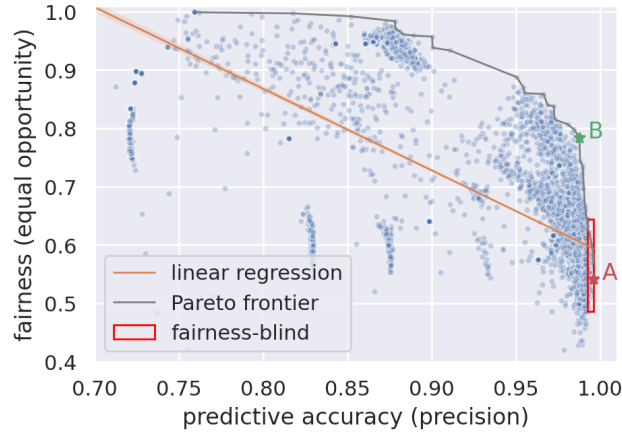
Preprint. Under review.

Figure 1: Fairness-accuracy trade-off of thousands of models on the Adult dataset. In orange, the linear regression relationship between accuracy and fairness; in the red rectangle, the top 10% models with highest accuracy (the target region for a fairness-blind process); in light gray, the fairness-accuracy Pareto frontier; marked with an A, the model with highest accuracy; marked with a B, a model with 0.8% lower accuracy and 44.8% higher fairness than A, arguably a better trade-off, and one that would not be found by traditional fairness-blind techniques.

model training), and often only apply to a single fairness definition or family of ML models, limiting their adoption in practice.

There is a lack of practical methodologies and tools to seamlessly integrate fairness objectives and bias reduction techniques in existing real-world ML pipelines. As a consequence, treating fairness as a primary objective when developing AI systems is not standard practice yet (Saleiro et al., 2020). Moreover, the absence of major breakthroughs in algorithmic fairness suggests that an exhaustive search over all possible ML models and bias reduction techniques may be necessary in order to find optimal fairness-accuracy trade-offs, hence discouraging AI practitioners. However, in algorithmic fairness, model selection becomes a multi-criteria problem. We must provide fairness-aware AI development efficiently, with limited computational resources, minimizing predictive accuracy impact, and without the explicit need for expert knowledge. Therefore, this work contributes to the state-of-the-art in algorithmic fairness by bridging the gap between research and practice from an unexplored dimension: efficient fairness-aware hyperparameter optimization.

Figure 1 illustrates the fairness-accuracy trade-off over thousands of models trained on the Adult dataset (Kohavi, 1996), each with a different hyperparameter configuration. Currently employed model selection processes are fairness-blind, solely optimizing for predictive accuracy. By doing so, these methods unknowingly target models with low fairness (region marked with a red rectangle). However, as shown by the plotted Pareto frontier (Pareto, 1906), we can achieve significant fairness improvements at small accuracy costs. For instance, model B achieves 44.8% higher fairness than model A (the model with highest predictive accuracy), at a cost of 0.8% decrease in predictive accuracy, arguably a better trade-off. While current fairness-blind techniques target model A, we target the region of optimal fairness-accuracy trade-offs to which model B belongs. Indeed, we observe a large spread over the fairness metric at any level of predictive accuracy, even within this fairness-blind region. Thus, it is absolutely possible to select fairer hyperparameter configurations without significant decrease in predictive accuracy. With this in mind, we present Fairband, a bandit-based fairness-aware hyperparameter optimization algorithm.

Fairband is a resource-aware HO algorithm that targets optimal fairness-accuracy trade-offs, without increasing training budget or sacrificing parallelization. By making the hyperparameter search fairness-aware while maintaining resource-efficiency, we are enabling AI practitioners to adapt pre-existing business operations to accommodate fairness with controllable extra cost, and without significant friction.

The summary of our contributions are as follows:

- Fairband, a flexible and efficient fairness-aware HO method for multi-objective optimization of the fairness-accuracy trade-off that is agnostic to both the explored hyperparameter space and the objective metrics (described in Section 3).

- A dynamic method to automatically search for good fairness-accuracy trade-offs without requiring manual weight parameterization (described in Section 3.1).

- A competitive baseline for fairness-aware HO: random search with balanced fairness-accuracy ranking of hyperparameter configurations.

- Strong empirical evidence that hyperparameter optimization is an effective way to navigate the fairness-accuracy trade-off.

- Competitive results on 4 real world datasets: Fairband achieves significantly improved fairness at a small predictive accuracy cost, and no extra budget when compared to literature HO baselines.

2   Related Work

Algorithmic fairness research work can be broadly divided into three families: pre-processing, in-processing, and post-processing.

Pre-processing methods aim to improve fairness before the model is trained, by modifying the input data such that it no longer exhibits biases. The objective is often formulated as learning new representations that are invariant to changes in specified factors (*e.g.*, membership in a protected group) (Calmon et al., 2017; Creager et al., 2019; Edwards and Storkey, 2016; Zemel et al., 2013). However, by acting on the data itself, and in the beginning of the ML pipeline, fairness may not be guaranteed on the end model that will be used in the real-world.

In-processing methods alter the model's learning process in order to penalize unfair decision-making. The objective is often formulated as optimizing predictive accuracy under fairness constraints (or optimizing fairness under predictive accuracy constraints) (Cotter et al., 2018). Another approach optimizes for complex predictive accuracy metrics which include some fairness notion (Zafar et al., 2017a), akin to regularization. However, these approaches are highly model-dependent and metric-dependent, and even non-existent for numerous state-of-the-art ML algorithms.

Post-processing methods aim to adjust an already trained classifier such that fairness constraints are fulfilled. This is usually done by calibrating the decision threshold (Fish et al., 2016; Hardt et al., 2016). As such, these approaches are flexible and applicable to any score-based classifier. However, one may argue that by acting on the model *after* it was learned this process is inherently sub-optimal (Woodworth et al., 2017). It is akin to knowingly learning a biased model and then correcting these biases, instead of learning an unbiased model from the start.

Although a largely unexplored direction, algorithmic fairness can also be tackled from a hyperparameter optimization perspective. One of the simplest and most flexible HO methods is random search (RS). This method iteratively selects combinations of random hyperparameter values and trains them on the full training set until the allocated budget is exhausted. Although simple in nature, RS has several advantages that keep it relevant nowadays, including having no assumptions on the hyperparameter space, on the objective function, or even the allocated budget (*e.g.*, it may run indefinitely). Additionally, RS is known to generally perform better than grid search (Bergstra and Bengio, 2012), and to converge to the optimum as budget increases.

Bayesian Optimization (BO) is a state-of-the-art HO method that consists in placing a prior (usually a Gaussian process) over the objective function to capture beliefs about its behavior (Shahriari et al., 2016). It iteratively updates the *prior* distribution, $p(\lambda)$, using the evidence *trial*, X, and forms a *posterior* distribution of its behavior, $p(\lambda|X)$. Afterwards, an acquisition function a is constructed to determine the next query point $\lambda^{(i+1)} = \arg\max_\lambda a(\lambda)$. This process is repeated in a sequential manner, continuously improving the approximation of the underlying objective function (Hutter et al., 2011).

Previous work has extended BO to constrained optimization (Gardner et al., 2014; Gelbart et al., 2014), in which the goal is to optimize a given metric subject to any number of data-dependent constraints. Recently, Perrone et al. (2020) applied this approach to the fairness setting by weighing the acquisition function by the likelihood of fulfilling the fairness constraints. However, constrained

i	$s = 4$		$s = 3$		$s = 2$		$s = 1$		$s = 0$	
	n_i	r_i	n_i	r_i	n_i	r_i	n_i	r_i	n_i	r_i
0	81	1.2	34	3.7	15	11	8	33	5	100
1	27	3.7	11	11	5	33	2	100		
2	9	11	3	33	1	100				
3	3	33	1	100						
4	1	100								

Table 1: Number of sampled configurations, n_i, and budget per configuration, r_i, for each Hyperband bracket (when the ratio of budget increase $\eta = 3$, and the maximum budget per configuration $R = 100$). One budget unit equals 1% of the train dataset.

optimization approaches inherently target a single fairness-accuracy trade-off (which itself may not be feasible), leaving practitioners unaware of the possible fairness choices and their accuracy costs. Additionally, BO has been shown to scale poorly (cubically in the number of data points) and to lack behind bandit-based methods when under budget constraints (Hutter et al., 2011; Li et al., 2016; Falkner et al., 2018).

Successive Halving (SH) (Karnin et al., 2013; Jamieson and Talwalkar, 2016) casts the task of hyperparameter optimization as identifying the best arm in a multi-armed bandit setting. Given a budget for each iteration, B_i, SH (1) uniformly allocates it to a set of arms (hyperparameter configurations), (2) evaluates their performance, (3) discards the worst half, and repeats from step 1 until a single arm remains. Thereby, the budget for each surviving configuration is effectively doubled at each iteration. SH's key insight stems from extrapolating the rank of configurations' performances from their rankings on diminished budgets (low-fidelity approximations). However, SH itself carries two parameters for which there is no clear choice of values: the total budget B and the number of sampled configurations n. We must consider the trade-off between evaluating a higher number of configurations (higher n) on an averaged lower budget per configuration (B/n), or evaluating a lower number of configurations (lower n) on an averaged higher budget. The higher the average budget, the more accurate the extrapolated rankings will be, but a lower number of configurations will be explored (and vice-versa). SH's performance has been shown to compare favorably to several competing bandit strategies from the literature (Audibert and Bubeck, 2010; Even-Bar et al., 2006; Jamieson et al., 2014; Kalyanakrishnan et al., 2012).

Hyperband (HB) (Li et al., 2016) addresses this "n *versus* B/n" trade-off by dividing the total budget into different instances of the trade-off, and then calling SH as a subroutine for each one. This is essentially a grid search over feasible values of n. HB takes two parameters R, the maximum amount of resources allocated to any single configuration; and η, the ratio of budget increase in each SH round ($\eta = 2$ for the original SH). Each SH run, dubbed a bracket, is parameterized by the number of sampled configurations n, and the minimum resource units allocated to any configuration r. The algorithm features an outer loop that iterates over possible combinations of (n,r), and an inner loop that executes SH with the aforementioned parameters fixed. The outer loop is executed $s_{max} + 1$ times, $s_{max} = \lfloor log_\eta(R) \rfloor$, and the inner loop (SH) takes approximately B resources. Thus, the execution of Hyperband takes a budget of $(s_{max} + 1)B$. Table 1 displays the number of configurations and budget per configuration within each bracket when considering $\eta = 3$ and $R = 100$.

Current bias reduction methods either (1) act on the input data and cannot guarantee fairness on the end model, (2) act on the model's training phase and can only be applied to specific model types and fairness metrics, or (3) act on a learned model's predictions and thus are limited to act on a sub-optimal space. On the other hand, hyperparameter optimization is simultaneously model independent, metric independent, and already a compoennt of existing real-world ML pipelines. By introducing fairness objectives on the HO phase in a effient way, we aim to help real-world practitioners to seamlessly finding optimal fairness-accuracy tradeoffs regardless of the underlying model type or bias reduction method.

3 Fairband

Our stated goal is to enable flexible and efficient fairness-aware hyperparameter optimization. To this end, we present Fairband (FB), a novel bandit-based algorithm for multi-criteria hyperparameter optimization.

By acting on the algorithms' hyperparameter space, we benefit from the advantages of in-processing bias reduction methods, while avoiding its shortcomings (*e.g.*, dependency on the metrics and model). That is, we aim to improve model fairness by acting on the models' training phase, instead of simply correcting the data (pre-processing), or correcting the predictions (post-processing). Nonetheless, as any multi-objective optimization problem with competing metrics, we can only aim to identify good fairness-accuracy trade-offs (Zafar et al., 2017b). The decision on which trade-off to employ should be left to the model's stakeholders.

Aiming to benefit from the efficiency of state-of-the-art resource-aware HO methods, we build our method on top of Successive Halving (SH) (Karnin et al., 2013) and Hyperband (HB) (Li et al., 2016). Thus, Fairband benefits from these methods' advantages: being both model- and metric-agnostic, having efficient resource usage, and trivial parallelization. Furthermore, these methods are highly exploratory and therefore prone to inspect broader regions of the hyperparameter space. For instance, in our experiments, Hyperband evaluates approximately six times more configurations than Random Search with the same budget[1]. Most importantly, these algorithms are easily extendable: by changing the sampling strategy, by changing how we evaluate a given hyperparameter configuration, or by changing how we select the top configurations to be kept between iterations. As an example, BOHB (Falkner et al., 2018) is an extension of Hyperband that introduces Bayesian optimization into the sampling strategy.

In order to introduce fairness objectives into the hyperparameter optimization process, we assume our goal is the maximization of a predictive accuracy metric a and a fairness metric f, for $a, f \in [0, 1]$ (or, equivalently, the minimization of $1 - a$ and $1 - f$).

Our method weighs the optimization metric by both predictive accuracy and fairness, parameterized by the relative importance of predictive accuracy $\alpha \in [0, 1]$ (see Equation 1). This is a popular method for multi-objective optimization known as *weighted-sum scalarization* (Deb, 2014). By employing this technique in a bandit-based setting, we rely on the hypothesis that if model m_a represents a better fairness-accuracy trade-off than model m_b with a short training budget, then this distinction is likely to be maintained with a higher training budget. Thus, by selecting models based on both fairness and predictive accuracy, we are guiding the search towards fairer and better performing models. These low-fidelity estimates of future metrics on lower budget sizes is what drives Hyperband and SH's efficiency in hyperparameter search. Our proposed optimization metric, o, is given by the following equation:

$$o = \alpha \cdot a + (1 - \alpha) \cdot f \tag{1}$$

Accordingly, all models are evaluated in both fairness and predictive accuracy metrics on a holdout validation set. Computing fairness does not imply significant extra computational cost, as it is based on the same predictions used to estimate predictive accuracy. Additionally, fairness assessment libraries are readily available (Saleiro et al., 2018).

3.1 Dynamic α

As a multi-objective optimization problem, we aim to identify configurations that represent a balanced trade-off between the two target metrics, $a, f \in [0, 1]$. Simply employing $\alpha = 0.5$ would assume one unit of improvement on one metric is effectively equal to one unit of improvement on the other, and would guide the search towards a single region of the fairness-accuracy trade-off.

Thus, aiming for a complete out-of-the-box experience without the need for specific domain knowledge, we propose a heuristic for automatically setting α values targeting a broader exploration of the Pareto frontier (Pareto, 1906) and a balance between searching for fairer or more accurate configurations. We dub this variant *FB-auto*. Assuming that α values can indeed guide the search

---

[1]With the parameters used on the Hyperband seminal paper (Li et al., 2016), it evaluates 128 configurations compared with 21 configurations evaluated by Random Search on an equal budget.

towards different regions of the fairness-accuracy trade-off (which we will empirically see to be true), our aim is to efficiently explore the Pareto frontier in order to find a comprehensive selection of balanced trade-offs. As such, if our currently explored trade-offs correspond to high accuracy but low fairness, we want to guide the search towards higher fairness (by choosing a lower α). Conversely, if our currently explored trade-offs correspond to high fairness but low accuracy, we want to guide the search towards higher accuracy (by choosing a higher α).

To achieve the aforementioned balance we need a proxy-metric of our target direction of change. This direction is given by the difference, δ, between the average model fairness, $\overline{f}$, and average predictive accuracy, $\overline{a}$, as shown in Equation 2:

$$\delta = \overline{f} - \overline{a}, \ \delta \in [-1, 1] \tag{2}$$

Hence, when this difference is negative, $\overline{f} < \overline{a}$, the models we sampled thus far tend towards better-performing but unfairer regions of the hyperparameter space. Consequently, we want to decrease α to direct our search towards fairer configurations. Conversely, when this difference is positive, $\overline{f} > \overline{a}$, we want to direct our search towards better-performing configurations, increasing α. We want this change in α to be proportional to δ by some constant $k > 0$, such that

$$\frac{d\alpha}{d\delta} = k, \ k \in \mathbb{R}^{+} \tag{3}$$

and by integrating this equation we get

$$\alpha = k \cdot \delta + c, \ c \in \mathbb{R} \tag{4}$$

with c being the constant of integration. Given that $\delta \in [-1, 1]$, and together with the constraint that $\alpha \in [0, 1]$, the only feasible values for k and c are $k = 0.5$ and $c = 0.5$. Hence, the computation of dynamic-α is given as follows by Equation 5:

$$\alpha = 0.5 \cdot (\overline{f} - \overline{a}) + 0.5 \tag{5}$$

Moreover, earlier iterations are expected to have lower predictive accuracy (as these are trained on a lower budget), while later iterations are expected to have higher predictive accuracy. By computing new values of α at each Fairband iteration, we promote a dynamic balance between these metrics as the search progresses, predictably giving more importance to accuracy on earlier iterations but continuously moving importance to fairness as accuracy increases (a natural side-effect of increasing training budget). Thus, we enable parameter-free fairness-aware optimization via *dynamic α*.

3.2 Algorithm

Firstly, we consider a broad hyperparameter space as a requirement for the effective execution of Fairband. We consider as hyperparameters any decision in the ML pipeline, as bias can be introduced at any stage of this pipeline (Barocas and Selbst, 2016). Thus, an effective search space includes which model type to use, the model hyperparameters which dictate how it is trained, and the sampling hyperparameters which dictate the distribution and prevalence rates of training data.

Fairband is detailed in Algorithm 1. It has three parameters: R, the maximum amount of resources allocated to any single configuration; η, which dictates both the budget increase and the proportion of configurations discarded in each SH round; and α, the relative importance of predictive accuracy versus fairness. The first two parameters are "inherited" from Hyperband. On the other hand, α may be omitted, relying on our proposed dynamic computation of α.

Our method incorporates Hyperband's exploration of SH's brackets, iterating through different parameters of SH corresponding to different instances of the "n versus B/n" trade-off (whether to evaluate more configurations on a lower budget, or less configurations on a higher budget). For each SH run (or bracket), our method (1) randomly samples n hyperparameter configurations, (2) trains each sampled configuration on the allocated budget r_i, (3) evaluates their predictive accuracy and fairness, and (4) selects the top k configurations on the objective metric o. Optionally, between steps 3 and 4, a dynamic value of α will be computed as described in Section 3.1. Note that when using a

Algorithm 1 Fairband

Input: maximum budget per configuration R,
        η (default $\eta = 3$),
        α (default $\alpha = auto$)

1: $s_{max} \leftarrow \lfloor \log_\eta (R) \rfloor$ ▷ define number of brackets
2: $B \leftarrow (s_{max} + 1) \cdot R$ ▷ compute budget per bracket
3: **for** $s \in \{s_{max}, s_{max} - 1, ..., 0\}$ **do**   ▷ iterate through SH brackets, as per Li et al. (2016)
4:     $n \leftarrow \lceil \frac{B}{R} \cdot \frac{\eta^s}{s+1} \rceil, r \leftarrow R \cdot \eta^{-s}$ ▷ choice of n *versus* B/n trade-off
5:     $T \leftarrow get\_hyperparameter\_configurations(n)$
6:     **for** $i \in \{0, ..., s\}$ **do**   ▷ run Successive Halving
7:         $n_i \leftarrow \lfloor n \cdot \eta^{-i} \rfloor$ ▷ train n_i configurations
8:         $r_i \leftarrow r \cdot \eta^i$ ▷ r_i training budget per config.
9:         $M \leftarrow \{train\_with\_budget(\boldsymbol{\lambda}, r_i) : \boldsymbol{\lambda} \in T\}$
10:        $A \leftarrow \{evaluate\_accuracy(m_{\boldsymbol{\lambda}}) : m_{\boldsymbol{\lambda}} \in M\}$
11:        $F \leftarrow \{evaluate\_fairness(m_{\boldsymbol{\lambda}}) : m_{\boldsymbol{\lambda}} \in M\}$
12:        **if** $\alpha = auto$ **then**   ▷ compute dynamic α if applicable
13:           $\overline{f} \leftarrow sum(F)/|F|$ ▷ average fairness
14:           $\overline{a} \leftarrow sum(A)/|A|$ ▷ average accuracy
15:           $\alpha \leftarrow 0.5 \cdot (\overline{f} - \overline{a}) + 0.5$
16:        $O \leftarrow \{\alpha \cdot A[m_{\boldsymbol{\lambda}}] + (1 - \alpha) \cdot F[m_{\boldsymbol{\lambda}}] : m_{\boldsymbol{\lambda}} \in M\}$ ▷ compute objective metric, o
17:        $I \leftarrow argsort(O)$ ▷ sorted in descending order
18:        $k \leftarrow \lfloor n_i / \eta \rfloor$ ▷ number of configurations to keep
19:        $T \leftarrow T[I[0 : k]]$ ▷ select top k configurations
20: **return** $\boldsymbol{\lambda}^*$, configuration with maximal intermediate goal seen so far

---

static value of $\alpha = 1$ our method is functionally equivalent to the traditional Hyperband algorithm (only taking predictive accuracy into account).

The result of our method's execution is a collection of hyperparameter configurations that effectively represent the fairness-accuracy trade-off. One could plot all available choices on the fairness-accuracy space and manually pick a trade-off, according to whichever business constraints or legislation are in place (see examples of Figure 3). For Fairband with static α, a target trade-off has already been chosen for the method's search phase, and we once again employ this trade-off for model selection (selection-α). For the FB-auto variant of Fairband, aiming for an automated balance between both metrics, we employ the same strategy for setting α as that used during search. By doing so, the weight of each metric is pondered by an approximation of their true range instead of blindly applying a pre-determined weight. For instance, if the distribution of fairness is in range $f \in [0, 0.9]$ but that of accuracy is in range $a \in [0, 0.3]$, then a balance would arguably be achieved by weighing accuracy higher, as each unit increase in accuracy represents a more significant relative change (this mechanism is achieved by Equation 5). However, at this stage we can use information from all brackets, as we no longer want to promote exploration of the search space but instead aim for a consistent and stable model selection. Thus, for FB-auto, the selection-α is chosen from the average fairness and predictive accuracy of all sampled configurations.

4 Experimental Setup

In order to validate our proposal, we evaluate Fairband on a search space spanning multiple ML algorithms, model hyperparameters, and sampling choices on four different datasets.

We compare our method to several baselines in the hyperparameter optimization community, including Random Search (RS) and Hyperband (HB). We study two versions of Fairband: FB-auto (employing the dynamic α strategy) and FB-bal (employing $\alpha = 0.5$). In addition, we consider an RS variant in which we introduce fairness-awareness into the final model selection criteria, giving equal importance to both metrics, *i.e.*, $\alpha = 0.5$ (RS-bal).

Dataset	Setting	Acc. Metric	Fairness Metric	Target Threshold	Sensitive Attribute
Donors Choose	assistive	precision	equal opportunity	1000 PP	poverty level
Adult	assistive	precision	equal opportunity	50% TPR	gender
COMPAS	punitive	precision	predictive equality	2% FPR	race
AOF	punitive	recall	predictive equality	5% FPR	age

Table 2: Details on used datasets and their metrics.

4.1 Datasets

We validate our methodology on three datasets from the fairness literature, and one large-scale case study on online bank account opening fraud. Both accuracy and fairness metrics are highly dependent on the task for which a given model is trained, and the real-world setting in which it will be deployed (Rodolfa et al., 2020). Thus, we detail a task for each of the datasets we employ, subsequently deriving the metrics we use for each.

The **Donors Choose** dataset (Wijesinghe et al., 2014) consists in data pertaining to thousands of projects proposed for/by K-12 schools. The objective is to identify projects at risk of getting underfunded to provide tailored interventions. As an assistive funding setting, we set a limit of 1000 positive predictions (PP), and select balanced true positive rates across schools from different poverty levels as the fairness metric, also known as equal opportunity (Hardt et al., 2016). We use precision as a metric of predictive accuracy.

The **Adult** dataset (Kohavi, 1996) consists on data from the 1994 US census, including age, gender, race, occupation, and income, among others. In order to properly employ this dataset, we devise a scenario of a social security program targeting low-income individuals. In this setting, a positive prediction indicates an income of less than $50K per year, thus making that person eligible for the assistive program. As an assistive setting, we select balanced true positive rates across genders as the fairness metric (equal opportunity). We target a global true positive rate (recall) of 50%, catching half of the total universe of individuals at need of assistance. Additionally, aiming to help only those that need it, we use precision as a metric of predictive accuracy.

The **COMPAS** dataset (Angwin et al., 2016) is a criminal justice dataset whose objective is to predict whether someone will re-offend based on the person's criminal history, demographics, and jail time. As a punitive setting, we select balanced false positive rates for individuals of different races as the fairness metric, also known as predictive equality (Corbett-Davies et al., 2017). At the same time, we target a global false positive rate of 2%, in order to maintain a very low number of unjustly jailed individuals. Regarding predictive accuracy, we use the precision metric.

The **AOF** dataset is a large-scale (500K instances) real-world dataset on online bank account opening fraud. The objective is to predict whether an individual's request for opening a bank account is fraudulent. As a punitive setting, we select balanced false positive rates across age groups as the fairness metric (predictive equality). In this setting, a false positive is a genuine customer that sees her/his request unjustly denied, costing the company potential earnings, and disturbing the customer's life. To maintain a low costumer attrition, we target a global false positive rate of 5%. In addition, we use recall (true positive rate) as a metric for predictive accuracy, as we want to promote models that correctly catch a high volume of fraud. Table 2 summarizes the task details for all datasets.

4.2 Search Space

First and foremost, we define a hyperparameter as any parameter used to tune an algorithm's learning process (Hutter et al., 2019). Both accuracy and fairness metrics are seen as (possibly noisy) black-box functions of these hyperparameters. This broad definition of hyperparameters includes the model type (which can be abstracted as a categorical hyperparameter), data sampling criteria, as well as regular hyperparameters pertaining only to the model.

In order to validate our methods, we define a comprehensive hyperparameter search space, allowing us to effectively navigate the fairness-accuracy trade-off. We select five ML model types: Random Forest (RF) (Breiman, 2001), Decision Tree (DT) (Breiman et al., 1984), Logistic Regression (LR) (Walker and Duncan, 1967), LightGBM (LGBM) (Ke et al., 2017), and feed-forward neural networks (NN).

Each model type has the same likelihood of being selected when randomly sampling configurations to test. When sampling a hyperparameter configuration, after a model type is randomly picked, we take a random sample of that model's hyperparameters. Besides the model type and model hyperparameters, we also experiment with three different undersampling strategies: targeting 20%, 10%, and 5% positive samples. This type of hyperparameter is only used in the AOF dataset, as its high class imbalance poses a challenge (it features 99 negatively-labeled samples per each positively-labeled sample). Thus, we define a *model* as the result of training a given *hyperparameter configuration* on a given train dataset.

4.3 Hyperparameter Parameters

We configure Fairband to execute SH brackets as in the original Hyperband algorithm. Following the authors (Li et al., 2016), we set $\eta = 3$. We define 1 budget unit as 1% of the train dataset, and we set the maximum budget allocated to any configuration as 100 budget units, $R = 100$ (100% of the training dataset). These settings result in:

$$s_{max} = \lfloor \log_\eta (R) \rfloor = 4, \quad B = (s_{max} + 1) \cdot R = 500 \tag{6}$$

The outer loop will run $s_{max} + 1 = 5$ times, for $s \in \{s_{max}, ..., 0\}$. Each run (or bracket) will consume at most $B = 500$ budget units. Accordingly, each bracket will use at most $s_{max} + 1 = 5$ training slices of increasing size, corresponding to the following dataset percentages: 1.23%, 3.70%, 11.1%, 33.3%, 100%. These training slices are sampled such that smaller slices are contained in larger slices, and such that the class-ratio is maintained (by stratified sampling). The number of configurations and budget per configuration within each bracket are displayed in Table 1. Throughout a complete Hyperband run, 143 unique hyperparameter configurations will be randomly sampled (total n_i for $i = 0$), and 206 models will be trained and evaluated (total n_i for $i \in \{0..., s_{max}\}$).

5 Results & Discussion

In this section, we present and analyze the results from our fairness-aware HO experiments. To validate the methodology, we guide hyperparameter search by evaluating all sampled configurations on the same validation dataset, while evaluating the best-performing configuration (according to the objective function o) on a held-out test dataset in the end. Likewise, the model thresholds are set on the validation dataset, and then used on both the validation and test datasets. All studied HO methods are given the same training budget: 2400 budget units.

Table 3 shows the validation and test results of running different HO methods on four chosen datasets[2]. Results are averaged over 15 runs, and statistical significant differences against the baseline methods are shown with $\triangle$ (HB) and $\diamond$ (RS). This table comprises the training and evaluation of 40K unique models, one of the largest studies of the fairness-accuracy trade-off to date. Our method (FB-auto) consistently achieves higher fairness than the baselines on all datasets, at a small cost in predictive accuracy (statistically significant on all datasets). The same trend is observed with the remaining fairness-aware methods (FB-bal and RS-bal), when compared with the fairness-blind methods (RS and HB).

The differences in predictive accuracy and fairness between the proposed fairness-aware methods and the fairness-blind baselines have strong statistical significance on the Donors Choose, Adult, and AOF datasets. The same trend is visible on the COMPAS dataset, although on this dataset fairness-blind RS does not achieve better performance than FB, while achieving substantially lower fairness. We note that the COMPAS dataset is the smallest by a significant margin (approximately one order of magnitude smaller than Adult and Donors Choose, and two orders smaller than AOF). We find the large variability in fairness between validation and test results on COMPAS to be best explained by its size, together with the strict target threshold of 2% FPR (which is set on the validation data, and used on both validation and test data).

Overall, FB-auto arguably achieves the best fairness-accuracy trade-off on three (Donors Choose, Adult, AOF) out of the four datasets. On the COMPAS dataset, FB-bal dominates the remaining

---

[2]Data, plots, and ML artifacts available at `https://github.com/feedzai/fair-automl`

Algo.	Validation		Test	
	Predictive Acc.	Fairness	Predictive Acc.	Fairness
Donors Choose				
FB-auto	53.8▲♦	97.9▲♦	50.0▲♦	**87.3▲♦**
FB-bal	52.8▲♦	**98.6▲♦**	49.9▲♦	86.1▲♦
RS-bal	52.3▲♦	95.2▲♦	50.4▲♦	82.2▲♦
RS	59.9	26.8	53.1	32.8
HB	**60.6**	28.8	**53.4**	34.7
Adult				
FB-auto	90.6▲♦	**95.6▲♦**	90.1▲♦	**93.9▲♦**
FB-bal	91.9▲♦	94.2▲♦	79.7▲♦	79.0▲♦
RS-bal	89.8▲♦	83.7▲♦	90.5▲♦	83.4▲♦
RS	99.3	55.1	**99.4**	55.1
HB	**99.4**	53.2	**99.4**	53.3
COMPAS				
FB-auto	83.9▲◊	**97.0▲♦**	79.2▲	41.6▲
FB-bal	84.2▲◊	96.0▲♦	79.4▲	**42.7▲**
RS-bal	80.8▲♦	77.2▲♦	74.8▲◊	40.7▲♦
RS	86.8	29.0	79.7	25.8
HB	**88.6**	17.7	**82.6**	24.1
AOF				
FB-auto	61.0▲♦	**69.7▲♦**	62.6▲♦	**74.9▲♦**
FB-bal	61.0▲♦	68.8▲♦	62.1▲♦	72.9▲♦
RS-bal	62.3▲♦	61.5▲♦	63.8▲	66.6▲♦
RS	67.3	37.8	67.4	40.4
HB	**68.7**	40.7	**69.0**	43.9

Table 3: Validation and test results for all algorithms on all datasets. Statistical significance is tested against the baseline models (RS and HB) with a Kolmogorov-Smirnov test (Lilliefors, 1967). Comparison with RS is indicated by ◊ when $p < 0.05$ (♦ when $p < 0.01$), and comparison with HB is indicated by △ when $p < 0.05$ (▲ when $p < 0.01$).

Dataset	Abs. Difference (pp)		Rel. Difference (%)	
	Predictive Acc.	Fairness	Predictive Acc.	Fairness
Donors Choose	-3.4	+52.6	-6.4	+152
Adult	-9.3	+40.6	-9.4	+76.2
COMPAS	-3.4	+17.5	-4.1	+72.6
AOF	-6.4	+31.0	-9.3	+70.6
Average	-5.6	+35.4	-7.3	+92.9

Table 4: Comparison of using Fairband (FB-auto) versus Hyperband (HB), on test results.

fairness-aware methods (although differences to FB-auto are not statistically significant). Unsurprisingly, HB achieves the highest predictive accuracy on all datasets. However, by using FB-auto we achieve 92.9% improvement in fairness at a cost of only 7.3% drop in predictive accuracy, averaged over all datasets. Table 4 summarizes the comparison between HB and FB-auto on all datasets.
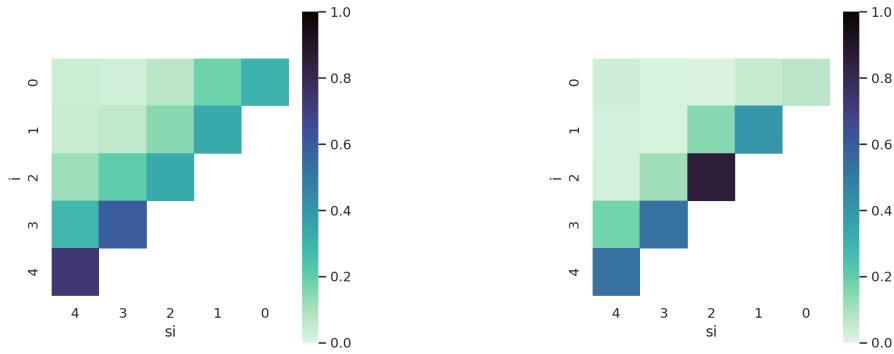
Figure 2: Average density of Pareto optimal models per FB-auto iteration (Adult dataset on left plot, AOF on right plot). Refer to Table 1 for information on the configurations at each iteration.
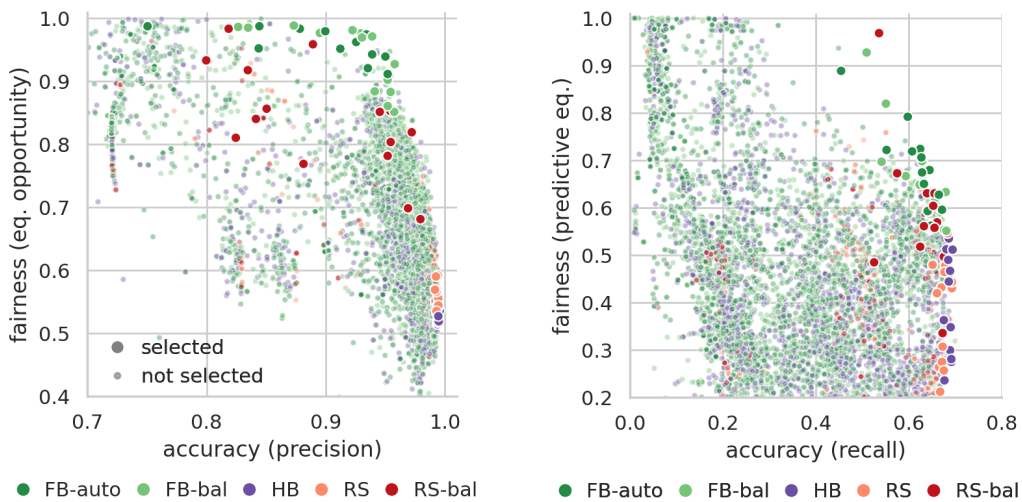

Figure 3: Fairness and predictive accuracy of selected models by hyperparameter optimization algorithm (Adult dataset on left plot, AOF on right plot).

5.1 Search Strategy

We evaluate the search strategy by analyzing the evolution of fairness and performance simultaneously, and whether we can effectively extend the practical Pareto Frontier as the search progresses. That is, whether optimal trade-offs are more likely to be found as we discard the worst performing models and further increase the allocated budget for top-performing models within each bracket (see Table 1). Figure 2 shows a heat map of the average density of Pareto optimal models[3] in each FB iteration, for 15 runs of the FB-auto algorithm, on the Adult and AOF datasets. As the iterations progress, under-performing configurations are pruned, and the density of Pareto optimal models steadily increases, confirming the effectiveness of the search strategy.

5.2 Model Selection

Figure 3 shows the final selected models for 15 runs of each method, respectively for the Adult (left) the AOF (right) datasets. The remaining models considered during the search are also shown, with

---

[3]Fraction of Pareto optimal models within each bracket, with optimality assessed within each run.

lower opacity and smaller size. As can be seen by the plots, Fairband consistently identifies good fairness-accuracy trade-offs from the universe of available configurations. Moreover, Fairband often achieves higher fairness than RS-bal for the same predictive accuracy. Indeed, the models selected by Fairband are consistently close to or form the Pareto frontier. Most importantly, as evident by the spread of selected models, we can successfully navigate the fairness-accuracy Pareto frontier solely by means of HO.

5.3 Efficiency over Budget

A different perspective of the HO methods' execution is their progression as the budget increases. We expect that increasingly better trade-offs are found as the budget increases, guided by our method's α. Figures 4 and 5 show this progression for the Adult and AOF datasets, respectively (error bands show 95% confidence intervals)[4].

Across all datasets, Fairband (both FB-auto and FB-bal) is able to provide strong anytime fairness, quickly converging to fairer regions of hyperparameter space, while RS-bal finds fairer configurations only at later stages. Indeed, FB-auto achieves better fairness (difference is statistically significant) at no cost to predictive accuracy (difference is not statistically significant) when compared to RS-bal.

On the AOF dataset, both versions of Fairband show an acute drop in fairness accompanied by symmetric increase in predictive accuracy by the 500 budget mark. This initial budget allocation corresponds to the left-most bracket ($s = 4$ in Table 1), which is highly exploratory (samples 81 different hyperparameter configurations) and thus trains each configuration on a small initial budget (1.2% of the training dataset). As configurations are pruned and the budget per configuration increases, the discovered trade-offs are progressively more accurate but less fair. This steep increase in the training budget per configuration on the first 500 budget units leads to the visibly high variability in objective metrics.

Regarding fairness-blind methods, RS and HB show similar behavior, with predictive accuracy increasing asymptotically, typically at the cost of fairness. However, HB consistently achieves higher predictive accuracy than RS at all stages.

When compared to literature HO baselines, both versions of Fairband achieve significantly improved fairness at a comparatively small cost in predictive accuracy. This trend is repeated by RS-bal, although achieving lower fairness than Fairband. It is important to consider that these fairness-blind baseline methods are the current standard in HO. By unfolding the fairness dimension, we show that strong predictive accuracy carries an equally strong real-world cost in unfairness, and this is hidden by traditional HO methods.
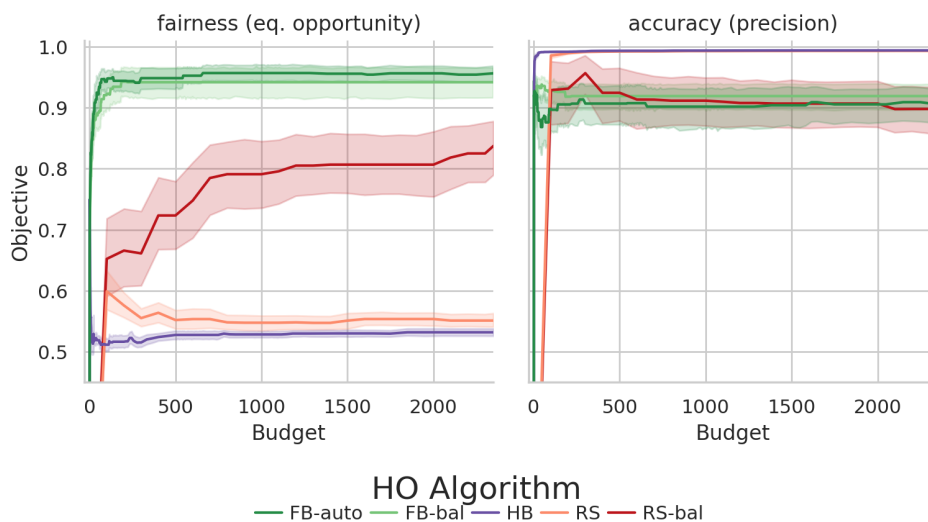

Figure 4: Fairness and predictive accuracy as a function of budget in the Adult validation set.

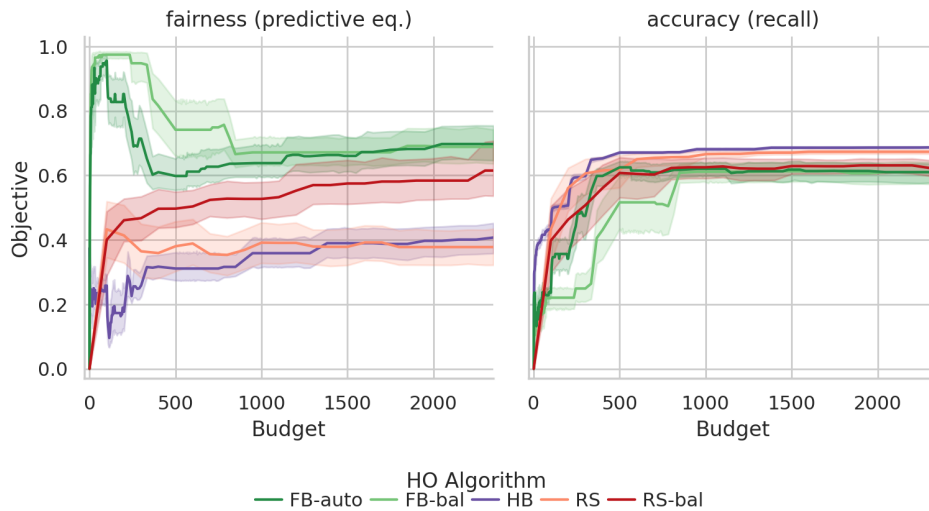[4]Plots for all datasets are available at `https://github.com/feedzai/fair-automl`

Figure 5: Fairness and predictive accuracy as a function of budget in the AOF validation set.

5.4 Optimizing Bias Reduction Hyperparameters

As a hyperparameter optimization method, one fruitful approach to bias-mitigation is adding bias reduction methods to Fairband's search space. As such, we introduce the Exponentiated Gradient (EG) reduction for fair classification algorithm[5] (Agarwal et al., 2018) into our search space on the Adult dataset. EG is a state-of-the-art bias reduction algorithm that optimizes predictive accuracy subject to fairness constraints, and is compatible with any cost-sensitive binary classifier. In our setting, we target equal opportunity, and apply EG over a Decision Tree classifier.

Figure 6 shows a plot of the models selected by FB-auto over 15 runs on the Adult dataset. The introduction of EG creates a new cluster of models in our search space (shown in orange), consisting of possible fairness-accuracy trade-offs in a previously unoccupied region (compare with Figure 3, left plot). However, even though these models were trained specifically targeting our fairness metric (equality of opportunity) while the remaining models were trained in a fairness-blind manner, Fairband chooses other model types more often than not. Indeed, the selected NNs and DTs arguably represent the best fairness-accuracy trade-offs.
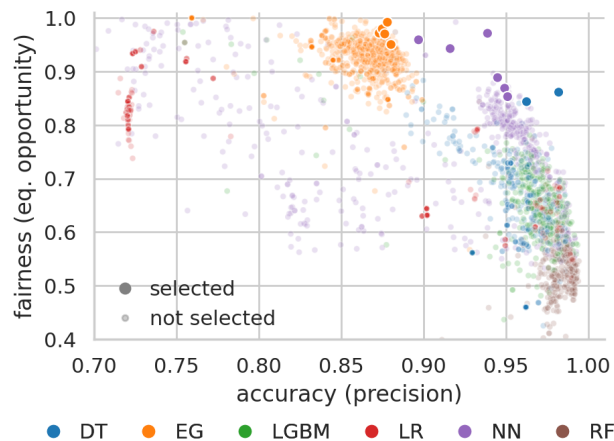

Figure 6: Fairness and predictive accuracy of selected models for FB-auto, discriminated by model type, on the Adult dataset.

---

[5]Implemented on the open-source *Fairlearn* package.

Essentially, these results show that blindly applying bias reduction techniques may lead to sub-optimal fairness-accuracy trade-offs. Overall, these results support the fact that Fairband should be employed in all ML pipelines that aim for fair decision-making, together with bias-blind methods and bias reduction methods alike, to properly explore the fairness-accuracy search space.

6 Conclusion

There have been widespread reports of real-world AI systems shown to be biased, causing serious disparate impact across different sub-groups, unfairly affecting people based on race, gender or age. The AI research community has embraced this issue and has been doing extensive research work. However, the current landscape of algorithmic fairness lacks (1) practical methodologies and (2) tools for real-world practitioners.

This work aims to bridge that gap by providing a simple and flexible hyperparameter optimization technique to foster the incorporation of fairness objectives in real-world ML pipelines. Fairband is a bandit-based fairness-aware hyperparameter optimization method that extends the Successive Halving and Hyperband algorithms by guiding the hyperparameter search towards fairer configurations[6].

Fairband enables targeting a specific fairness-accuracy trade-off (by means of an α parameter), which is often dictated by business restrictions or regulatory law. Aiming for a complete out-of-the-box experience, we alternatively propose an algorithm for setting α automatically, eliminating the need to tune this parameter.

By introducing fairness notions into hyperparameter optimization, our method can be seamlessly integrated into real-world ML pipelines, at no extra training cost. Moreover, our method is easy to implement, resource-efficient, and both model- and metric-agnostic, providing no obstacles to its widespread adoption.

We evaluate our method on four real-world decision-making datasets, and show that it is able to provide significant fairness improvements at a small cost in predictive accuracy, when compared to traditional HO techniques. We show that it is both possible and effective to navigate the fairness-accuracy trade-off through hyperparameter optimization.

Crucially, we observe that there is a wide spread of attainable fairness values at any level of predictive accuracy. At the same time, we once again document the known inverse relation between fairness and predictive accuracy. Hence, by blindly optimizing a single predictive accuracy metric (as is standard practise in real-world ML systems) we are inherently targeting unfairer regions of the hyperparameter space.

Acknowledgements

References

Agarwal, A., Beygelzimer, A., Dudik, M., Langford, J., and Wallach, H. (2018). A Reductions Approach to Fair Classification. In Dy, J. and Krause, A., editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 60–69, Stockholmsmässan, Stockholm Sweden. PMLR.

Angwin, J., Larson, J., Kirchner, L., and Mattu, S. (2016). Machine bias: There's software used across the country to predict future criminals. and it's biased against blacks. `https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing`. Accessed: 2020-01-09.

---

[6]Fairband is completely agnostic to the hyperparameter search space and therefore it is subject to the fairness-accuracy trade-offs the input hyperparameter configurations are able to attain.

Audibert, J.-Y. and Bubeck, S. (2010). Best Arm Identification in Multi-Armed Bandits. In *COLT - 23th Conference on Learning Theory - 2010*, page 13 p., Haifa, Israel.

Barocas, S. and Selbst, A. D. (2016). Big data's disparate impact. *Calif. L. Rev.*, 104:671.

Bartlett, R., Morse, A., Stanton, R., and Wallace, N. (2019). Consumer-Lending Discrimination in the FinTech Era. Technical report, National Bureau of Economic Research.

Bergstra, J. and Bengio, Y. (2012). Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, 13:281–305.

Berk, R., Heidari, H., Jabbari, S., Kearns, M., and Roth, A. (2018). Fairness in criminal justice risk assessments: The state of the art. *Sociological Methods & Research*, page 0049124118782533.

Board, F. S. (2017). Artificial intelligence and machine learning in financial services. *Market developments and financial stability implications*, 1.

Breiman, L. (2001). Random Forests. *Machine Learning*, 45(1):1–33.

Breiman, L., Friedman, J., Stone, C. J., and Olshen, R. A. (1984). *Classification and regression trees*. CRC press.

Buolamwini, J. and Gebru, T. (2018). Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. In Friedler, S. A. and Wilson, C., editors, *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, volume 81 of *Proceedings of Machine Learning Research*, pages 77–91, New York, NY, USA. PMLR.

Calmon, F., Wei, D., Vinzamuri, B., Ramamurthy, K. N., and Varshney, K. R. (2017). Optimized pre-processing for discrimination prevention. In *Advances in Neural Information Processing Systems*, pages 3992–4001.

Corbett-Davies, S., Pierson, E., Feller, A., Goel, S., and Huq, A. (2017). Algorithmic Decision Making and the Cost of Fairness. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '17*, pages 797–806, New York, New York, USA. ACM Press.

Cotter, A., Jiang, H., and Sridharan, K. (2018). Two-Player Games for Efficient Non-Convex Constrained Optimization. *Proceedings of the 30th International Conference on Algorithmic Learning Theory*, 98:300–332.

Creager, E., Madras, D., Jacobsen, J.-H. H., Weis, M. A., Swersky, K., Pitassi, T., and Zemel, R. (2019). Flexibly Fair Representation Learning by Disentanglement. In Chaudhuri, K. and Salakhutdinov, R., editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 2019-June, pages 2600–2612. PMLR.

Deb, K. (2014). *Multi-objective Optimization*, pages 403–449. Springer US, Boston, MA.

Edwards, H. and Storkey, A. (2016). Censoring Representations with an Adversary. In *4th International Conference on Learning Representations, ICLR 2016 - Conference Track Proceedings*, pages 1–14.

Even-Bar, E., Mannor, S., and Mansour, Y. (2006). Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *Journal of Machine Learning Research*, 7:1079–1105.

Falkner, S., Klein, A., and Hutter, F. (2018). BOHB: Robust and Efficient Hyperparameter Optimization at Scale. *35th International Conference on Machine Learning, ICML 2018*, 4:2323–2341.

Fish, B., Kun, J., and Lelkes, Á. D. (2016). A confidence-based approach for balancing fairness and accuracy. *Proceedings of the 2016 SIAM International Conference on Data Mining*, pages 144–152.

Friedler, S. A., Scheidegger, C., Venkatasubramanian, S., Choudhary, S., Hamilton, E. P., and Roth, D. (2019). A comparative study of fairness-enhancing interventions in machine learning. In *Proceedings of the conference on fairness, accountability, and transparency*, pages 329–338.

15

Gardner, J., Kusner, M., Zhixiang, Weinberger, K., and Cunningham, J. (2014). Bayesian optimization with inequality constraints. volume 32 of *Proceedings of Machine Learning Research*, pages 937–945, Bejing, China. PMLR.

Gelbart, M. A., Snoek, J., and Adams, R. P. (2014). Bayesian optimization with unknown constraints. In *Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence*, UAI'14, page 250–259, Arlington, Virginia, USA. AUAI Press.

Hardt, M., Price, E., Price, E., and Srebro, N. (2016). Equality of opportunity in supervised learning. In Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I., and Garnett, R., editors, *Advances in Neural Information Processing Systems 29*, pages 3315–3323. Curran Associates, Inc.

Hutter, F., Hoos, H. H., and Leyton-Brown, K. (2011). Sequential model-based optimization for general algorithm configuration. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 6683 LNCS:507–523.

Hutter, F., Kotthoff, L., and Vanschoren, J. (2019). *Automated Machine Learning: Methods, Systems, Challenges*. Springer Publishing Company, Incorporated, 1st edition.

Jamieson, K., Malloy, M., Nowak, R., and Bubeck, S. (2014). lil' UCB : An Optimal Exploration Algorithm for Multi-Armed Bandits. In Balcan, M. F., Feldman, V., and Szepesvári, C., editors, *Proceedings of The 27th Conference on Learning Theory*, volume 35, pages 423–439, Barcelona, Spain. PMLR.

Jamieson, K. and Talwalkar, A. (2016). Non-stochastic best arm identification and hyperparameter optimization. *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics, AISTATS 2016*, pages 240–248.

Kalyanakrishnan, S., Tewari, A., Auer, P., and Stone, P. (2012). PAC subset selection in stochastic multi-armed bandits. *Proceedings of the 29th International Conference on Machine Learning, ICML 2012*, 1:655–662.

Karnin, Z., Koren, T., and Somekh, O. (2013). Almost optimal exploration in multi-armed bandits. *30th International Conference on Machine Learning, ICML 2013*, 28(PART 3):2275–2283.

Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., and Liu, T.-Y. (2017). Lightgbm: A highly efficient gradient boosting decision tree. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R., editors, *Advances in Neural Information Processing Systems 30*, pages 3146–3154. Curran Associates, Inc.

Kleinberg, J., Ludwig, J., Mullainathan, S., and Rambachan, A. (2018). Algorithmic fairness. In *Aea papers and proceedings*, volume 108, pages 22–27.

Kohavi, R. (1996). Scaling up the accuracy of naive-bayes classifiers: A decision-tree hybrid. In *Kdd*, volume 96, pages 202–207.

Li, L., Jamieson, K., DeSalvo, G., Rostamizadeh, A., and Talwalkar, A. (2016). Hyperband: A Novel Bandit-Based Approach to Hyperparameter Optimization. *Journal of Machine Learning Research*, 18:1–52.

Lilliefors, H. W. (1967). On the Kolmogorov-Smirnov Test for Normality with Mean and Variance Unknown. *Journal of the American Statistical Association*, 62(318):399–402.

Pareto, V. (1906). Manuale di economica politica, societa editrice libraria. *Manual of political economy*, 1971.

Perrone, V., Donini, M., Kenthapadi, K., and Archambeau, C. (2020). Fair bayesian optimization.

Rajkomar, A., Dean, J., and Kohane, I. (2019). Machine Learning in Medicine. *New England Journal of Medicine*, 380(14):1347–1358.

Rodolfa, K. T., Saleiro, P., and Ghani, R. (2020). Chapter 11: Bias and Fairness. In Foster, I., Ghani, R., Jarmin, R. S., Kreuter, F., and Lane, J., editors, *Big data and social science: A practical guide to methods and tools*. crc Press.

Saleiro, P., Kuester, B., Stevens, A., Anisfeld, A., Hinkson, L., London, J., and Ghani, R. (2018). Aequitas: A bias and fairness audit toolkit. *arXiv preprint arXiv:1811.05577*.

Saleiro, P., Rodolfa, K. T., and Ghani, R. (2020). Dealing with bias and fairness in data science systems: A practical hands-on tutorial. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 3513–3514.

Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., and De Freitas, N. (2016). Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175.

Suresh, H. and Guttag, J. V. (2019). A framework for understanding unintended consequences of machine learning.

Walker, S. H. and Duncan, D. B. (1967). Estimation of the Probability of an Event as a Function of Several Independent Variables. *Biometrika*, 54(1/2):167.

Wijesinghe, D. B., Paraneetharan, S., Gogulan, T., and Ahiladas, B. (2014). Predicting excitement at donorschoose.org - kdd cup 2014.

Woodworth, B., Gunasekar, S., Ohannessian, M. I., and Srebro, N. (2017). Learning Non-Discriminatory Predictors. *Proceedings of the 2017 Conference on Learning Theory*, 65(1):1920–1953.

Zafar, M. B., Valera, I., Gomez Rodriguez, M., and Gummadi, K. P. (2017a). Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In *Proceedings of the 26th international conference on world wide web*, pages 1171–1180.

Zafar, M. B., Valera, I., Rogriguez, M. G., and Gummadi, K. P. (2017b). Fairness constraints: Mechanisms for fair classification. In *Artificial Intelligence and Statistics*, pages 962–970. PMLR.

Zemel, R. S., Wu, Y., Swersky, K., Pitassi, T., and Dwork, C. (2013). Learning fair representations. *30th International Conference on Machine Learning, ICML 2013*, 28(PART 2):325–333.