

A Human-in-the-Loop Fairness-Aware Model Selection Framework for Complex Fairness Objective Landscapes

Jake Robertson^{1, 2}, Thorsten Schmidt¹, Frank Hutter^{1, 3}, Noor Awad¹

¹University of Freiburg,

²Zuse School ELIZA,

³ELLIS Institute Tübingen

robertsj@cs.uni-freiburg.de, thorsten.schmidt@stochastik.uni-freiburg.de,

fh@cs.uni-freiburg.de, awad@cs.uni-freiburg.de

Abstract

Fairness-aware Machine Learning (FairML) applications are often characterized by complex social objectives and legal requirements, frequently involving multiple, potentially conflicting notions of fairness. Despite the well-known Impossibility Theorem of Fairness and extensive theoretical research on the statistical and socio-technical trade-offs between fairness metrics, many FairML tools still optimize or constrain for a single fairness objective. However, this one-sided optimization can inadvertently lead to violations of other relevant notions of fairness. In this socio-technical and empirical study, we frame fairness as a many-objective (MaO) problem by treating fairness metrics as *conflicting objectives*. We introduce *ManyFairHPO*, a human-in-the-loop, fairness-aware model selection framework that enables practitioners to effectively navigate complex and nuanced fairness objective landscapes. *ManyFairHPO* aids in the identification, evaluation, and balancing of fairness metric conflicts and their related social consequences, leading to more informed and socially responsible model-selection decisions. Through a comprehensive empirical evaluation and a case study on the Law School Admissions problem, we demonstrate the effectiveness of *ManyFairHPO* in balancing multiple fairness objectives, mitigating risks such as self-fulfilling prophecies, and providing interpretable insights to guide stakeholders in making fairness-aware modeling decisions.

1 Introduction

Instances of algorithmic discrimination are a growing concern in both the machine learning (ML) literature and, more recently, in the media and greater society. This is a consequence of the increasing prevalence of ML applications where individuals are disparately impacted by algorithmic decisions (Angwin et al. 2016; de Zwart 2022). Mirroring the complex and socio-technical nature of the machine bias problem, the field of Fairness-aware Machine Learning (FairML) has emerged, providing a collaborative space for political philosophers, social scientists, legislators, statisticians, and ML researchers. FairML has the overarching goals of defining, studying, detecting, and mitigating algorithmic bias.

Despite significant advancements, the FairML community has received widespread criticism from the Social Sciences,

Humanities, and Law for attempting to solve the complex, nuanced, and socio-technical problem of machine bias *algorithmically* (Hoffmann 2019; Selbst et al. 2019). Several arguments cite that real-world applications of FairML are often characterized by a complex set of social objectives and legal requirements (Ruf and Detyniecki 2021). Due to their complexity, these criteria are unlikely to be captured by single, coarsely-grained statistical measures of fairness. In such cases, FairML methods that only incorporate a single notion of fairness risk *fair-washing*, or proposing a so-called *fair* model that satisfies one notion of fairness while violates another potentially relevant one, potentially resulting in negative social consequences. Other criticisms of FairML cite the black-box nature of bias-mitigation techniques as a key concern (Robertson, Stinson, and Hu 2022), suggesting a crucial interplay between fairness, transparency, and interpretability (Barocas, Hardt, and Narayanan 2023; Schöffer 2023). Broadly effective FairML methods should not only cope with a diverse set of objectives and requirements but ideally offer interpretable insights.

Rather than resisting these criticisms, we embrace them, taking the perspective that FairML approaches that oversimplify the complex and socio-technical nature of FairML problems risk doing more harm than good. Instead, FairML approaches must adapt to the context in which they exist. In recent years, the topic of fairness has gained popularity in the Automated Machine Learning (AutoML) literature, which typically formulates fairness as a bi-objective (BiO) or constrained hyperparameter optimization (HPO) problem. Fairness-aware HPO varies common ML design decisions (tree depth, neural architecture, neural network width, etc.) to explore the Pareto Front of fair and accurate models (Wu and Wang 2021; Perrone et al. 2021; Schmucker et al. 2020; Dooley* et al. 2023). According to Weerts et al. (2023), fairness-aware AutoML holds key advantages in human-centricity and transparency, enabling practitioners to explore multiple fairness-accuracy trade-offs and gain interpretable insights into the fairness-objective landscape of the problem at hand. In addition, the BiO problem formulation extends naturally to the MaO case, encompassing three or more fairness and accuracy constraints or objectives, enabling practitioners to explore not only fairness-accuracy trade-offs but trade-offs between fairness metrics themselves. Although the prospect of constraining or op-

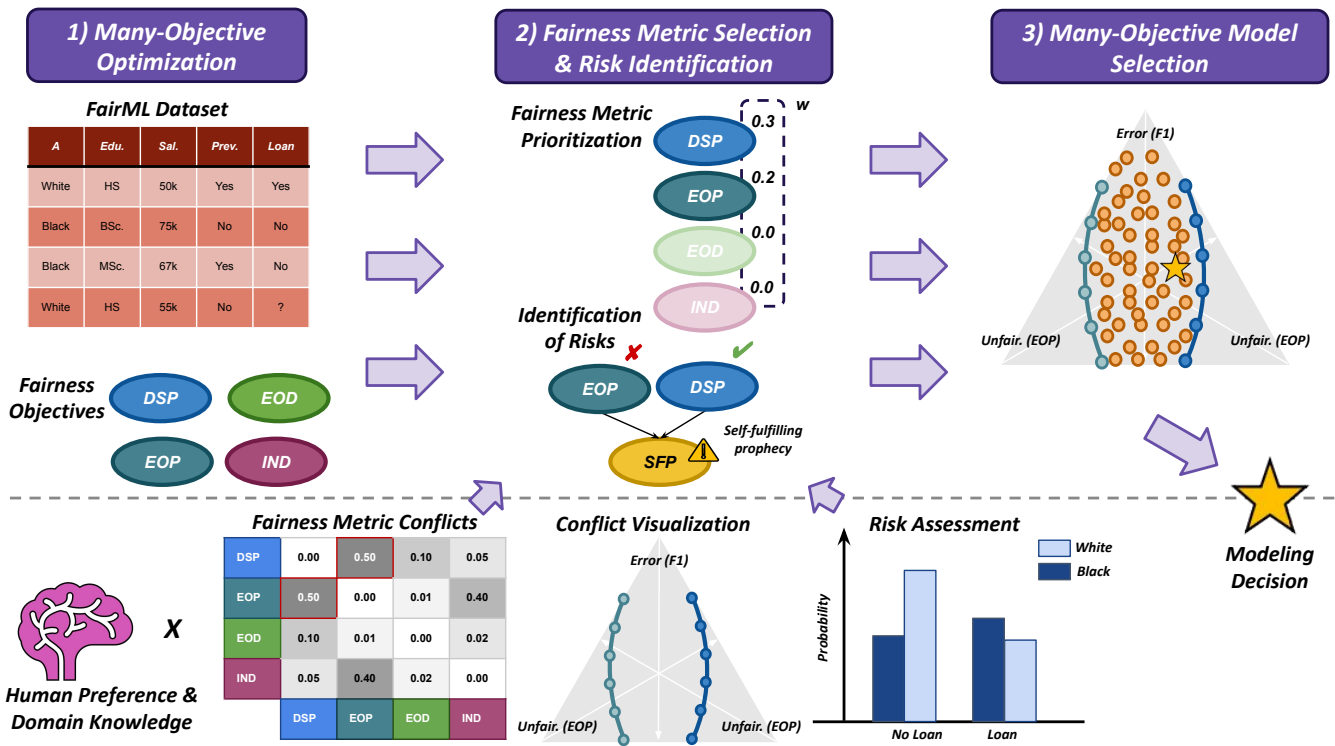


Figure 1: *ManyFairHPO*: We introduce *ManyFairHPO*, a human-in-the-loop, many-objective optimization framework for navigating complex fairness and performance trade-offs in machine learning. The approach consists of three main stages: 1) Many-Objective Optimization to explore trade-offs between a candidate set of fairness and performance metrics, 2) Fairness Metric Selection & Risk Identification to incorporate domain knowledge, human preferences, and insights from Pareto front analysis to prioritize metrics and assess fairness conflict risks, and 3) Many-Objective Model Selection to choose a final model that balances multiple objectives based on assigned weights.

optimizing for multiple notions of fairness has been mentioned several times in the literature (Perrone et al. 2021; Schmucker et al. 2020), previous studies lack a convincing socio-technical and empirical basis to answer the crucial question: “*how and why should we put this into practice?*” To bridge these gaps in the literature, our study motivates the MaO problem formulation for fairness, making the following contributions:

1. We propose a human-in-the-loop optimization framework for fairness, *ManyFairHPO*, that aids practitioners in the identification, evaluation, and balancing of fairness metric conflicts and their related social consequences, enabling fairness-aware model-selection decisions that encapsulate multiple conflicting fairness objectives.
2. In an empirical and socio-technical study, we evaluate and motivate the MaO HPO problem formulation for fairness on a diverse set of fairness benchmarks. The key message of our evaluation is *ManyFairHPO* performs competitively to the typical BiO problem formulation and state-of-the-art bias-mitigation techniques. Additionally, we find that *ManyFairHPO* thoroughly explores the trade-offs between fairness metrics themselves.
3. Exemplifying the efficacy of *ManyFairHPO* in fairness problems with a complex set of fairness objectives, we

perform a case study on the topical Law School Admissions problem. We provide a simulation with multiple stakeholders, demonstrating how our interpretable metrics and visualizations can help them understand the potential social consequences of fairness metric conflicts, such as the risk of self-fulfilling prophecy. We show how the interpretability of *ManyFairHPO* can lead to a thorough, stakeholder-informed model selection decision that mitigates this risk.

In the remainder of the paper, after summarizing related works that incorporate multiple fairness objectives and constraints (Section 2), we introduce the related terminology in algorithmic fairness as well as multi-objective optimization (Section 3). We then introduce our approach, *ManyFairHPO*, and our experimental setup (Section 4). Our results (Section 5) show that *ManyFairHPO* is both suitable as a bias-mitigation technique and enables practitioners to navigate and explore complex fairness objective landscapes. Finally, we conclude with possible future directions of research and a call to action for the fairness community to consider fairness metric conflicts and the MaO problem formulation as a crucial next step in fairness research.

2 Related Work

The characterization of fairness as an inherently BiO problem has received significant attention in the literature (Martinez, Bertran, and Sapiro 2020; Islam, Pan, and Foulds 2021), with several works citing the importance of considering fairness-accuracy trade-offs in the FairML problem landscape. In this section, we discuss the critical issue of optimization problem formulation for fairness, comparing the typical BiO framing to its counterpart, constrained optimization (CO).

We begin by discussing several works that formulate the FairML task as a CO problem and discuss its pros and cons compared to the BiO problem formulation. In Fair Bayesian Optimization, Perrone et al. (2021) expand upon common bias-mitigation techniques to integrate constraints across multiple fairness metrics. Relatedly, Hsu et al. (2022) employ Mixed Integer Linear Programming to enforce constraints across various fairness criteria, to explore and push the boundaries. Finally, Exponentiated Gradient Reduction (EGR) claims to “yield a randomized classifier with the lowest (empirical) error subject to the desired constraints” (Agarwal et al. 2018). Due to its strong reputation as a state-of-the-art constrained approach, we select EGR as a baseline in our empirical analysis (Section 5).

The first key difference between the BiO and constrained problem formulation is that the latter assumes the knowledge of appropriate fairness metrics and achievable constraints. Although previous socio-technical work (Ruf and Detyniecki 2021) provides useful guidelines for fairness metric selection, the range of achievable fairness metric values (especially in the presence of conflicts) is highly data and model-dependent. While we align with the argument of Perrone et al. (2021) that COtimization is more computationally efficient in cases where fairness objectives and reasonable constraints are clearly defined, we maintain that this is an unrealistic scenario, especially when constraining across multiple, potentially conflicting notions of fairness. We also highlight that preferences towards certain fairness metrics or the enforcement of fairness constraints can still be performed after BiO optimization, making it a more flexible method when computational resources are not extremely limited and prior knowledge of the fairness objective landscape is scarce.

Secondly, because constrained approaches focus the search on the constrained region of interest, they typically do not explore the entire Pareto Front of all objectives, often leaving quality (in this case either in terms of performance or fairness) on the table (Weerts et al. 2023). Additionally, by exploring the entire Pareto Front(s), approaches provide interpretable insights into the overall objective landscape, a key aspect in building trust (Schöffner 2023) and aiding practitioners in the iterative FairML design cycle (Weerts et al. 2023). Such insights include but are not limited to the degree of trade-offs between fairness metrics, the nature and shape of Pareto Fronts (e.g. knee points) (Dhiman and Kumar 2019), and the relative difficulty of different fairness objectives (Jordan, Cohen, and Thomas 2018).

A final key feature of the BiO problem formulation is its natural extension to the MaO setting, which enables practi-

tioners to incorporate multiple user-defined fairness metrics as additional objectives. Although Schmucker et al. (2020) argue that their methodology is “extensible” to the MaO setting, they do not verify this hypothesis or explore its benefits. In this study, we fill this gap in the FairML literature, providing a socio-technical and empirical basis for the MaO problem formulation.

3 Background

3.1 Fairness-aware Machine Learning

In this section, we provide an introduction to fairness and its many definitions, bias mitigation strategies, as well as the theoretical trade-offs between fairness and performance objectives. For a more comprehensive review of the fairness literature, we refer to (Pessach and Shmueli 2022).

FairML approaches seek to detect and mitigate machine bias at various stages of the ML pipeline and can be effectively divided into pre-processing, in-processing, and post-processing techniques, which detect and mitigate bias in either 1) the input data, 2) the training algorithm or 3) the model’s predictions (Pessach and Shmueli 2022). In recent years, the objective of fairness has gained increased popularity in the AutoML community, resulting in several studies that achieved competitive performance to specialized bias-mitigation techniques by simply varying the hyperparameters or neural architecture of ML models (Perrone et al. 2021; Schmucker et al. 2020; ?). For an in-depth discussion on fairness-aware AutoML, we refer to Weerts et al. (2023).

Despite the strong performance of bias-mitigation techniques and fairness-aware AutoML, a lingering question in the fairness literature remains an introspective one: how do we define fairness itself? Grappling with a philosophical question that has been debated for millennia (Binns 2018), as well as a growing set of social objectives and legal requirements, the FairML community has proposed over 20 different fairness metrics (Barocas, Hardt, and Narayanan 2023), which seek to quantify the fairness (or unfairness) of a model. The FairML problem is formally defined below.

Definition 3.1 (FairML Problem) *Given a data set $\mathcal{D} = (X, Y, A)$ of n features and m samples $X \in \mathbb{R}^{m \times n}$, a binary target $Y \in \{0, 1\}^m$, and a binary protected attribute $A \in \{0, 1\}^m$, find a model $\mathcal{M} : X \rightarrow \hat{Y}$ that satisfies or optimizes a set of fairness and performance objectives or constraints $\{f_0, f_1, f_2, \dots\}$*

The three most common metrics for group fairness are Demographic Statistical Parity (DDSP), Equalized Odds (DEOD), and Equal Opportunity (DEOP). Group fairness operates upon the *egalitarian* principle that valuable resources should be distributed equally across salient demographic groups, while individual fairness is based on the notion of *just deserts* (Binns 2018), requiring that individuals receive their deserved outcome. Associated individual fairness metrics are similarity-based and require that similar individuals (based on a set of *legitimate factors*) receive similar outcomes (Barocas, Hardt, and Narayanan 2023). Due to the challenge in defining legitimate factors, we opt for a simplified definition of Inverse Distance (INVD) as proposed by Berk et al. (2017), which treats the ground truth

as a legitimate factor (despite the potential for direct bias). We provide a mathematical definition of the above fairness metrics in (Table 1).

A foundational concept, especially in the fairness-aware AutoML literature is the *fairness-accuracy trade-off*, which states that improvements in fairness usually come at the cost of reduced predictive accuracy (Pessach and Shmueli 2022). This is because unfairness is caused primarily by *indirect bias*, which occurs due to an indirect influence of protected attributes on the target variable via a set of *indirect encodings* (Kamishima et al. 2012) which may also influence the outcome of interest. For example, standardized test scores are a partially useful predictor of student success. However, these scores can also be indirectly influenced by race via latent factors such as access to test-preparation resources. For a more mathematical discussion of the fairness-accuracy trade-off, we refer to (Cooper and Abrams 2021).

3.2 Fairness Metric Impossibility Theorem

A central concept to this study in particular is the *Impossibility Theorem of Fairness*, which states that the three main notions of group fairness (Sufficiency, Separation, and Independence) cannot all be satisfied at once (Miconi 2017), (Kleinberg, Mullainathan, and Raghavan 2016) (Chouldechova 2016). The Impossibility Theorem applies when there is a difference in base rates, and is thus widely applicable in real-world FairML problems. According to Barocas, Hardt, and Narayanan (2023):

“What this shows is that we cannot impose multiple criteria as hard constraints. This leaves open the possibility that meaningful trade-offs between these different criteria exist.”

3.3 Fairness Metric Conflict-Related Risks

Given the potential conflicts between fairness metrics themselves, we formalize the notion of conflict-related risks. In Figure 2 (left) we first outline our perspective on fairness metrics as a mapping to specific social objectives such as equity, equality, and individual justice. When one metric is satisfied by violating another (right), a *first-order* risk occurs, namely the reduction in the social benefits of the violated metric (e.g. reduced individual justice). However, more consequential are the *second-order* risks of fairness metric violations, which occur rather as a downstream effect. In their critical analysis of DDSP, (Dwork et al. 2012) describe three conflict-related risks, namely 1) reduced utility 2) self-fulfilling prophecy (SFP) and 3) subset targeting. In the context of this study, we focus on the risk of SFP due to its clear connection to a conflict between group fairness metrics DDSP and DEOP. We explain SFP with an example.

Imagine a university admissions algorithm that obtains equal acceptance rates across demographic groups by elevating the chances of underprivileged students getting accepted. This approach leads to the inclusion of several underqualified students from this group. Although this algorithm achieves low DDSP unfairness by achieving equal acceptance rates, it violates the unfairness metric DEOP, which

Fairness Metrics		Social Objectives		Violated Metric			
				DSP	EOP	EOD	IND
Statistical Parity (DSP)	Equality						
Equal Opportunity (EOP)	Equity						
Equalized Odds (EOD)	Equity						
Inverse Distance (IND)	Individual Justice						

Satisfied Metric	DSP	EOP	EOD	IND	
	DSP		SFP	SFP	SFP
	EOP	?		?	?
	EOD	?	?		?
IND	?	?	?		

Figure 2: *Conflict-Related Risks*: Fairness metrics serve as a statistical means to specify and optimize various social objectives, such as equality, equity, and individual justice. However, certain conflicts hold potential downstream consequences, such as the risk of Self-Fulfilling Prophecy (SFP) when DDSP is satisfied by violating other metrics.

measures the between-group acceptance rate of qualified applicants. In addition to the reduction in equality of opportunity, the second-order risks are more dangerous. If the underprivileged students who are accepted do not succeed at the university, such an admissions strategy could create positive feedback loops for future discrimination, ultimately reinforcing the self-fulfilling prophecy that the underprivileged group is less qualified (Dwork et al. 2012).

We emphasize, through the various question marks in Figure 2, that fairness metric conflict-related risks is a largely understudied topic. Additionally, we note that these conflicts (and their associated risks) are not symmetric, a phenomenon that we describe in Appendix Section A.2.

3.4 Multi-Objective Optimization

Multi-objective optimization is an increasingly relevant topic in ML research, as ML applications are increasingly scrutinized with respect to environmental, logistical, and social requirements (Karl et al. 2022). Multi-objective optimization seeks to provide solutions to an optimization problem under a set of two or more *conflicting* objectives (Sharma and Kumar 2022). Objectives are called conflicting when gains in one objective trade-off against losses in the other.

Let Λ be a design space representing all possible design decisions. The goal of multi-objective optimization is to find a set of design solutions $\lambda \in \Lambda$ that minimize a multi-criteria objective function $f : \Lambda \rightarrow R^d$, which returns a vector of costs $\langle f_1(\lambda), f_2(\lambda), \dots, f_d(\lambda) \rangle$ with respect to each objective. We consider the strict partial Pareto order \prec_{Pareto} where $x \prec_{\text{Pareto}} y$ iff $x_i \leq y_i$ for all $1 \leq i \leq d$ and $x_j < y_j$ for at least one j . If we consider all possible values $\mathcal{Y} = f(\Lambda)$ we obtain a set of minimal values called the *Pareto Front*.

$$\mathcal{P}(\mathcal{Y}) := \{y \in \mathcal{Y} : \{y' \in \mathcal{Y}, y' \prec_{\text{Pareto}} y\} = \emptyset\} \quad (1)$$

Any value in $\mathcal{P}(\mathcal{Y})$ is called *non-dominated*, i.e., there is no other solution that is as least as good as this value for all objectives but strictly better in at least one of the objectives. The *hypervolume* is a quality indicator and represents the volume of the dominated region of the objective space with respect to a reference point r .

$$\mathcal{H}(\mathcal{P}) := \text{Vol}(\mathcal{P}(\mathcal{Y})) \quad (2)$$

Name	Definition
Demographic Statistical Parity (DDSP)	$ P(\hat{Y} = 1 A = 0) - P(\hat{Y} = 1 A = 1) $
Equalized Opportunity (DEOP)	$ P(\hat{Y} = 1 A = 0, Y = 1) - P(\hat{Y} = 1 A = 1, Y = 1) $
Equalized (Average) Odds (DEOD)	$\frac{1}{2} \sum_{y \in \{0,1\}} \cdot P(\hat{Y} = 1 A = 0, Y = y) - P(\hat{Y} = 1 A = 1, Y = y) $
Inverse Distance (INVD)	$\frac{1}{m^2} \cdot \sum_{i,j=1}^m y_i - y_j \cdot \hat{y}_i - \hat{y}_j $

Table 1: *Fairness metrics*: Given protected attributes A , targets Y , and predictions \hat{Y} , we summarize measures of unfairness drawn from the `ai360` library. INVD is a simplified version of similarity-based individual fairness which treats ground truth labels as legitimate factors for the inverse distance function ($|y_i - y_j|$).

In this study, we focus on *normalized hypervolume*, an extension of the concept of hypervolume which scales the empirical value to a (0, 1) range, making it comparable across different optimization tasks. For a further discussion on normalized hypervolume and other multi-objective quality indicators, we refer to (Hansen et al. 2022).

MaO optimization is a subcategory of multi-objective optimization and is applicable to high-dimensional problems with three or more objectives. MaO algorithms are conceptually similar to low-dimensional multi-objective algorithms but often implement additional quality-diversity measures to ensure sufficient exploration of all objectives (Deb and Jain 2014). For simplicity and to avoid confusion, we divide multi-objective optimization into two problem settings: BiO (optimizing for two objectives) and MaO (optimizing for more than two objectives).

3.5 Hyperparameter Optimization

Hyperparameter optimization (HPO) seeks to automate the trial-and-error process of designing and deploying ML models (Bischl et al. 2023). HPO has been shown to consistently improve performance across a wide variety of ML tasks, with high-profile success stories (Chen et al. 2018) and has since become a crucial aspect of the ML design cycle. Another advantage of HPO is its extensibility to the BiO case, where varying ML hyperparameters can have an impact on not only performance, but other ML objectives, such as interpretability, energy efficiency, and namely, fairness (Karl et al. 2022). In the context of fairness-aware AutoML, recent studies have shown a significant impact of regularizing hyperparameters on commonly used fairness metrics. Fairness-aware HPO provides a convenient framework for the FairML problem and is also extensible to include preprocessing, in-processing, and post-processing techniques (as well as their own hyperparameters) as part of the search (Wu and Wang 2021).

3.6 Evolutionary Algorithms

Many design spaces are large and non-differentiable, rendering an exhaustive search or gradient-based methods computationally infeasible. A popular approach to multi-objective problems is the Evolutionary Algorithm (EA), a population-based optimization technique that draws inspiration from the process of biological evolution to solve black-box, non-differentiable optimization problems (Eiben and Smith 2015). By implementing bio-evolutionary concepts such as

selection, mutation, and crossover, EAs effectively balance exploration and exploitation, generating state-of-the-art results in a variety of domains.

Nondominated-Sorting Genetic Algorithm (NSGA-II) is a state-of-the-art MO-EA (Deb et al. 2002), which applies the notion of dominance in order to recursively divide the population into ranked fronts. Parent and survivor selection is performed using the heuristics of front rank and crowding distance, which measures the distance of solutions to their nearest neighbor on their respective front. Entire fronts are greedily selected until including the next front exceeds a pre-defined quota. At this stage, individuals are selected from the final front based on crowding distance, which serves as a heuristic for uniqueness, and encourages exploration through population diversity. NSGA-III (Deb and Jain 2014) incorporates the concept of reference directions in order to select individuals for crossover and survival from equally spaced regions of the many-dimensional objective space.

4 Methodology

In this section, we introduce ManyFairHPO, our MaO optimization approach for fairness that provides an interpretable, human-in-the-loop framework to help practitioners navigate the complex landscape of fairness metric conflicts and make informed and socially responsible model selection decisions. ManyFairHPO consists of three main stages: 1) MaO optimization for a candidate set of fairness metrics, 2) fairness metric selection and risk identification, and 3) fairness-aware model selection. The description below follows Figure 1, which provides a visual overview of the ManyFairHPO framework.

4.1 Many-Objective Optimization

ManyFairHPO takes a FairML dataset and a non-weighted candidate set of performance and fairness metrics f_0, f_1, f_2, \dots as input. We then employ the popular MaO optimization algorithm NSGA-III (Deb and Jain 2014) to efficiently explore the objective space of fairness-accuracy trade-offs and fairness metric conflicts, yielding a MaO Pareto Front $P(\mathcal{Y}_{MULTI})$. Starting with MaO optimization, ManyFairHPO differs from other approaches by delaying the selection and prioritization of performance and fairness objectives until the trade-offs between these metrics are known.

4.2 Fairness Metric Selection & Risk Identification

The second stage of ManyFairHPO focuses on assigning a set of performance and objective weights $\langle w_0, w_1, w_2, \dots \rangle$ based on domain knowledge, human preferences, and insights taken from the high-dimensional Pareto analysis.

Domain Knowledge and Human Preferences Fairness objective weights can be derived from discussions about the relative importance of fairness objectives. To facilitate these discussions, we draw an analogy between fairness metrics and their related social objectives (Section 3.1), making it easier for domain experts and stakeholders to express their preferences. It is important to note that stakeholders and domain experts may have different priorities, which can be accounted for during objective weighting. For instance, a hiring agency might prioritize individual fairness to maintain a reputation for non-discriminatory decisions, while a hiring discrimination expert might focus on group fairness to ensure workplace diversity.

Another crucial discussion point is the identification of potential fairness metric conflict-related risks, which are the potential downstream consequences when one fairness metric is satisfied by violating another (e.g., self-fulfilling prophecy). Key questions to consider include: “Will ML-assisted decisions be used to train future algorithms?” and “Are individuals from underprivileged groups being thoughtfully selected?” It is important to keep in mind that the identified risks may not be exhaustive or guaranteed to occur, motivating the need for exploratory analysis of the high-dimensional Pareto Fronts.

Interpretability via Pareto Analysis MaO optimization provides interpretability through post-hoc analysis of the objective landscape. ManyFairHPO offers interpretable metrics and visualizations to identify fairness metric conflicts, visualize them, and evaluate their social consequences. To quantify the strength of fairness metric conflicts, we introduce the notion of fairness metric *contrast*:

Definition 4.1 (Contrast) *The contrast of fairness metric f_i with respect to fairness metric f_j is defined as the difference in normalized hypervolume when optimizing for f_j and f_i , respectively:*

$$C(f_i, f_j) := \mathcal{H}_{f_j}(\mathcal{P}(\mathcal{Y}_j)) - \mathcal{H}_{f_j}(\mathcal{P}(\mathcal{Y}_i)) \quad (3)$$

It’s worth noting that the contrast metric is not symmetric (Appendix Section A.2). A large positive value of $C(f_i, f_j)$ indicates a severe conflict, where optimizing for f_i fails to optimize for f_j . A value close to zero suggests a weak conflict, and a negative value points to a *negative* conflict, where optimizing for a different fairness metric is better than optimizing for one directly.

To help practitioners understand fairness metric trade-offs, we provide visualizations related to contrast, such as 3-dimensional ternary plots, which reveal whether conflicts occur at specific accuracy levels. We also recommend comparing the predictive behavior of models selected from similar accuracy levels on two conflicting Pareto Fronts to assess whether the conflict results in a downstream risk.

4.3 Fairness-Aware Model Selection

The final stage of ManyFairHPO involves selecting a model that effectively captures social objectives and mitigates conflict-related risks. Insights from the previous stage can be used to assign performance and fairness objective weights. For a pair of conflicting fairness metrics (f_i, f_j) , where satisfying f_i violates f_j and leads to unwanted downstream consequences, a higher weight can be assigned to f_j to avoid this risk (Figure 1).

Given a high-dimensional Pareto Front $P(\mathcal{Y}_{MULTI})$ and a vector of weights $\langle w_0, w_1, \dots \rangle$, the MaO model selection task is simplified to:

$$\lambda^* = \operatorname{argmin}_{\lambda} \sum w_i \cdot f_i(\lambda) \quad (4)$$

For simplicity, we define Equation 4 as a linear combination, but more complex functions can be used to incorporate the utility and marginal gains of metrics.

4.4 Experimental Setup

Our experimental scope comprises of hyperparameter search spaces from HPOBench (Eggenberger et al. 2021) for Random Forest (RF), XGBoost (XGB), and Multi-Layer Perceptron (NN) models (Pedregosa et al. 2011) (Appendix Table 2) applied to fairness datasets: Bank Marketing, German Credit, Adult Census Income, COMPAS Criminal Recidivism, and Law School Admissions (Vanschoren et al. 2014) (Appendix Table 3). We use commonly employed fairness metrics from IBM’s `ai360` library (Bellamy et al. 2019): Demographic Statistical Parity (DDSP), Equalized Odds (DEOD), Equality of Opportunity (DEOP), and Inverse Distance (INVD) (Appendix Table 1).

We apply NSGA-III to optimize for a five-dimensional Pareto Front $P(\mathcal{Y}_{MULTI})$ of F_1 -Score and all four fairness metrics, resulting in 15 MaO experiments repeated with 10 random seeds (150 total runs). For comparison, we apply BiO optimization (NSGA-II) to optimize for F_1 -Score and a single fairness metric, resulting in 60 BiO experiments repeated with 10 random seeds (600 total runs). We also compare our results with Exponentiated Gradient Reduction (EGR) by in-processing the most accurate model for each search space and dataset across 10 independent trials. The source code for experiments and analysis as well as supplementary material is available at <https://github.com/jr2021/ManyFairHPO-AIES>.

Research Questions To systematically evaluate the effectiveness of ManyFairHPO and demonstrate its practical application, we propose the following research questions:

RQ1: Fairness-Aware Many-Objective Optimization. What are the advantages of ManyFairHPO as an approach to bias mitigation compared to BiO optimization and state-of-the-art bias-mitigation techniques? How well does it explore fairness-accuracy trade-offs and fairness metric conflicts? What interpretable insights can be gained into fairness objective landscapes that can be used to select and prioritize fairness metrics, as well as identify and assess conflict-related risks?

RQ2: *Law School Admissions Case Study.* How can ManyFairHPO be applied in a real-world scenario with complex fairness objectives and diverse stakeholder interests? We explore a Law School Admissions case study to demonstrate how ManyFairHPO can relate social objectives and downstream risks to a prioritized set of fairness metrics, informing the selection of a model that incorporates multiple conflicting performance and fairness objectives.

5 Results

5.1 ManyFairHPO for Bias-Mitigation

In this section, we systematically assess the effectiveness of the MaO problem formulation (ManyFairHPO) from a technical perspective. We demonstrate that ManyFairHPO is competitive from a bias-mitigation standpoint, achieving similar fairness and accuracy values compared to BiO optimization and outperforming the state-of-the-art bias-mitigation technique EGR. Moreover, we showcase ManyFairHPO’s novel ability to explore and balance trade-offs between fairness metrics themselves.

Fairness-Accuracy Trade-offs We first verify that optimizing for multiple fairness metrics together using ManyFairHPO is at least as effective as optimizing for the same metrics separately using BiO optimization.

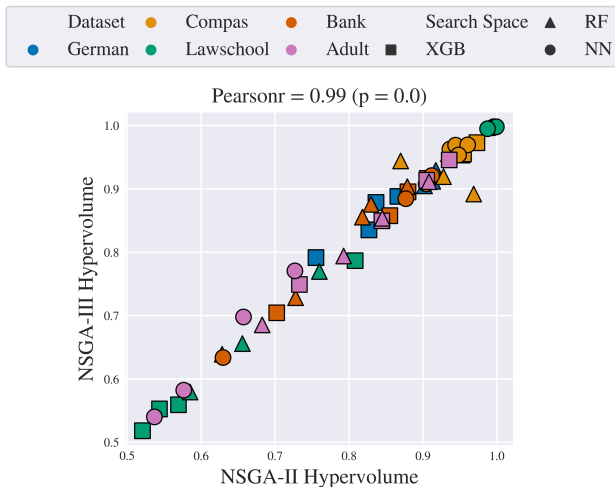


Figure 3: *ManyFairHPO vs. Bi-Objective Optimization:* Comparison of hypervolume achieved by ManyFairHPO and bi-objective optimization across datasets and search spaces. NSGA-III matches the hypervolume achieved by NSGA-II with a correlation of 0.991, indicating that similar fairness-accuracy Pareto Fronts are achieved when optimizing for fairness metrics together compared to optimizing for them separately.

To empirically validate this, we compare the fairness-accuracy hypervolume (normalized) achieved by the MaO optimizer NSGA-III compared to its BiO counterpart NSGA-II across different BiO combinations, datasets, and hyperparameter search spaces (Figure 3). We observe a

Dataset	Objectives	Search Space		
		XGB	RF	NN
German	F1-DDSP	10/0/0	10/0/0	10/0/0
	F1-DEOD	10/0/0	10/0/0	10/0/0
Compas	F1-DDSP	10/0/0	0/10/0	10/0/0
	F1-DEOD	5/5/0	0/0/10	10/0/0
Lawschool	F1-DDSP	0/10/0	0/0/10	10/0/0
	F1-DEOD	0/0/10	0/0/10	7/3/0
Bank	F1-DDSP	10/0/0	10/0/0	10/0/0
	F1-DEOD	10/0/0	10/0/0	10/0/0

Table 2: *ManyFairHPO vs. In-Processing (EGR):* Number of wins (ManyFairHPO dominates EGR), ties (non-dominated solutions), and losses (EGR dominates ManyFairHPO) across different datasets and search spaces. ManyFairHPO dominates EGR in 9/12 cases for DDSP and 7/12 cases for DEOD.

strong and significant correlation of 0.991 ($p = 0.0$) between the hypervolume achieved from optimizing for multiple fairness metrics together as opposed to optimizing for them separately.

Next, we compare the Pareto Fronts obtained from ManyFairHPO to EGR, a state-of-the-art bias-mitigation technique that reduces the FairML task to a series of cost-sensitive classification problems (Agarwal et al. 2018). To provide a fair comparison, we in-process the most accurate hyperparameter configurations with EGR for 10 independent trials across all search spaces and models with respect to group fairness metrics DDSP and DEOD.¹ We summarize the wins, ties, and losses (W/T/L) of ManyFairHPO against EGR in Table 2, where ManyFairHPO dominates EGR in the majority of trials in 9/12 cases for DDSP and 7/12 cases for DEOD. We also visualize these outcomes in Appendix Figure 2. The dominance of ManyFairHPO demonstrates that MaO optimization is capable of efficiently exploring multiple fairness-accuracy trade-offs, achieving comparable Pareto Fronts to BiO optimization and outperforming the state-of-the-art bias-mitigation technique EGR in most cases.

Fairness Metric Trade-offs Beyond achieving competitive fairness-accuracy trade-offs, a key novelty of ManyFairHPO is its ability to explore trade-offs between fairness metrics themselves. To identify which fairness metrics conflict on which search spaces and datasets, we calculate the contrast between fairness metrics $C(f_i, f_j)$ as described in Definition 4.1 and provide a summary of observed conflicts in Figure 4. A dark red cell in row j , column i , indicates that optimizing for F_1 -Score and fairness metric f_i fails to optimize for F_1 -Score and fairness metric f_j .

A key insight from this analysis is the strong data dependence of fairness metric conflicts, motivating the need for

¹Our results do not include a comparison on the Adult dataset because single evaluations of EGR exceeded our 24-hour time budget.

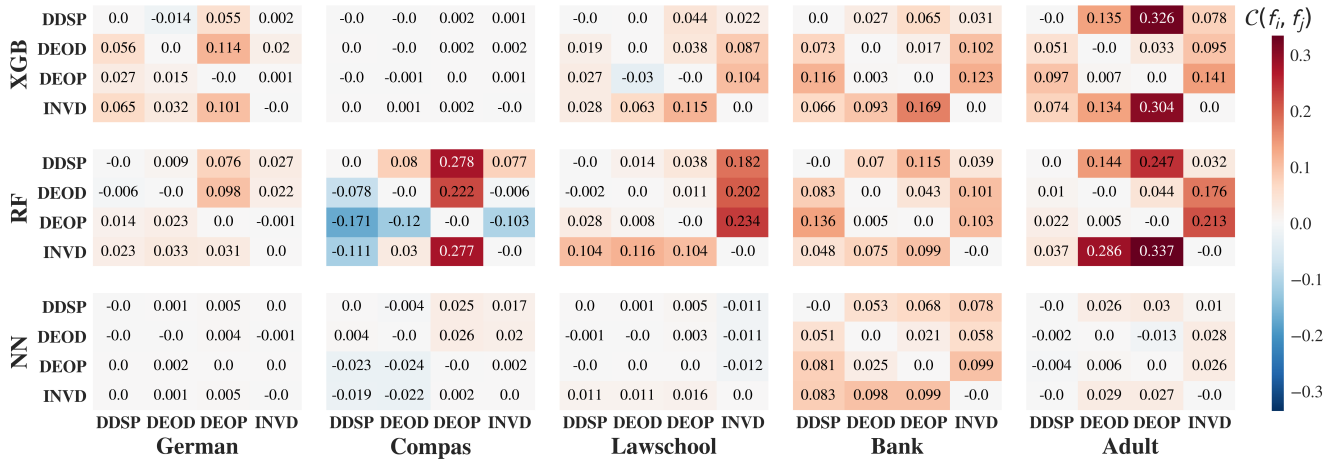


Figure 4: *Fairness Metric Contrast*: Heatmap visualization of fairness metric conflicts across different datasets and ML models, measured by the degree to which optimizing for one fairness-accuracy trade-off fails to optimize for another. Light red cells indicate more severe conflicts. The results highlight the data and model dependence of fairness metric conflicts.

MaO exploration of fairness metric trade-offs on real-world problems. We observe similar conflict patterns on datasets with semantically similar target variables, such as the Bank and Adult datasets, whose targets (Bank Marketing Subscription and Income) both serve as proxies for financial status. This suggests a relationship between FairML problem characteristics and the presence of certain conflicts, motivating future work on which conflicts and risks to be aware of for different classes of tasks. We also observe a slight model dependence, with lower error models displaying stronger conflicts, indicating that ManyFairHPO should be applied over multiple hyperparameter search spaces to reveal the true fairness objective landscape. Finally, we observe several negative conflicts on the RF-Compas experiment, where a fairness metric is better optimized indirectly than directly, a result which suggests a possible interaction effect between fairness metrics in hyperparameter search spaces. We further explain this result in Appendix A.1.

To illustrate ManyFairHPO’s novel capability in balancing fairness metric conflicts, we focus on a conflict observed between INVD and DEOP ($C = 0.337$) from the RF-Adult experiment (Figure 4 top-right). In Figure 5, we provide a ternary plot of this conflict projecting the 3-dimensional normalized locations of the Pareto Fronts $\mathcal{P}(\mathcal{Y}_{DEOP})$ and $\mathcal{P}(\mathcal{Y}_{INVD})$ in 2-dimensional space.

The first observation we make is that INVD and DEOP are only not in conflict at the high and low error extrema of the objective space, where the model selections are either unusable due to high error (top corner) or perform poorly in terms of INVD and DEOP (bottom edge). In the top 20% of most accurate models, the Pareto Fronts $\mathcal{P}(\mathcal{Y}_{DEOP})$ and $\mathcal{P}(\mathcal{Y}_{INVD})$ diverge, meaning that for a given accuracy level in this range, a selection from the $\mathcal{P}(\mathcal{Y}_{DEOP})$ Pareto Front would result in a significant trade-off in INVD. In contrast, the MaO Pareto Front $\mathcal{P}(\mathcal{Y}_{MULTI})$ provides a large number of model selections that interpolate between these choices,

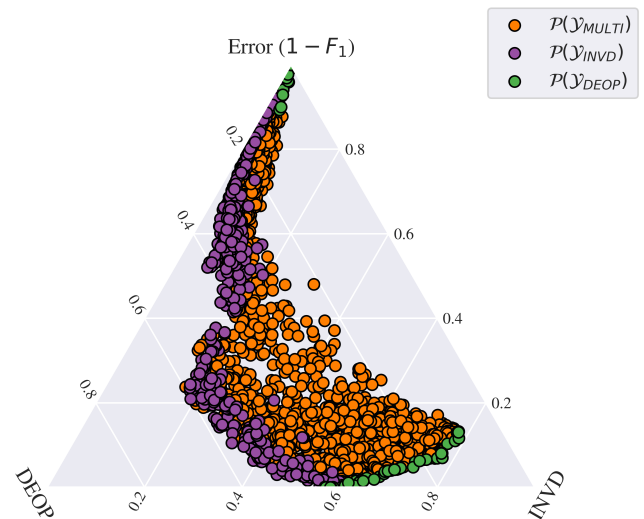


Figure 5: *Conflict Interpolation (RF-Adult)*: Ternary plot of the high-dimensional Pareto Front (orange) in the presence of a fairness metric conflict between INVD (purple) and DEOP (green). Metric values get better further away from each corner. The many-objective Pareto Front provides a selection of models that interpolate between fairness metric conflicts.

allowing practitioners to trade-off fairness metrics with a finer grain.

In summary, these results showcase ManyFairHPO’s ability to not only find competitive fairness-accuracy trade-offs compared to BiO optimization and state-of-the-art bias-mitigation techniques but also to explore and balance trade-offs between fairness metrics themselves. This is a novel contribution to the FairML literature.

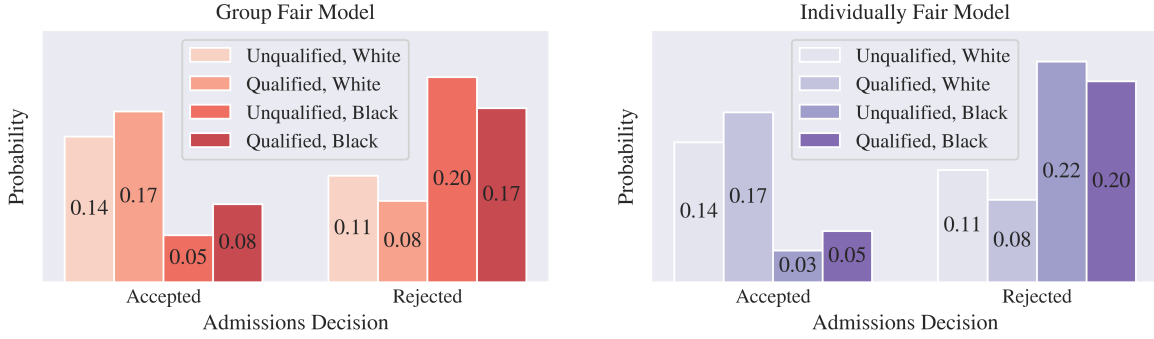


Figure 6: *Risk Assessment (RF-Lawschool)*: Explanation of the conflict between individually-fair and group-fair models discovered selected from the RF-Lawschool experiments. The individually fair model has a higher rejection rate, resulting in fewer similarly unqualified Black and White students receiving different outcomes (individual fairness). However, this strategy also increases between-group acceptance rates (group unfairness).




Figure 7: *Model Selection (RF-Law School)*: Model selection on the Law School Admissions problem under objective weights $\mathbf{w} = \langle 0.5, 0.2, 0.3 \rangle$ for F_1 -Score, DDSP, and INVD. Single objective scalarization leads to a modeling decision (green cross) with strong accuracy ($F_1 = 0.7$) that successfully mitigates the risk of self-fulfilling prophecy.

5.2 Law School Admission Case Study

To illustrate the practical application of ManyFairHPO in navigating intricate fairness landscapes, we present a case study using the Law School Admissions dataset from the 1991-1997 LSAC National Longitudinal Bar Passage Study (Wightman 1998).

Problem Description and Stakeholder Analysis This dataset, containing information from nearly 30,000 Law School applicants, is commonly used in the fairness community to simulate acceptance decisions based on predicted first-year average (FYA) scores. Given the high ethnic disparity identified in the initial LSAC study, we demonstrate how ManyFairHPO can be used to balance the complex and

conflicting fairness objectives relevant to this scenario.

As outlined in Section 4, the first step of ManyFairHPO is MaO optimization. We apply NSGA-III to optimize for performance (F_1) as well as fairness metrics DDSP, DEOP/D, and INVD. Before examining the fairness metric conflicts, we introduce a hypothetical scenario involving two stakeholders with different preferences toward higher-level objectives.

The University, the first stakeholder, aims to make admissions decisions that maintain a high academic standard while attracting a large applicant pool. Consequently, they are hesitant towards Affirmative Action strategies that may discourage applicants from well-represented groups. The second stakeholder, the University’s Equality, Diversity, and Inclusion (EDI) committee, prioritizes increasing diversity and breaking down systemic inequalities in both the admissions process and the University as a whole. After discussing their objectives, the University and EDI committee agree on initial weights of $\mathbf{w} = \langle 0.5, 0.5 \rangle$ for F_1 -Score and DDSP, respectively.

Conflict Evaluation and Assessment With the initial objective weights established based on the University and EDI committee’s preferences, we demonstrate how ManyFairHPO’s interpretability can assist practitioners in identifying, understanding, and assessing fairness metric conflict-related risks.

Figure 4 reveals a conflict between the individual fairness metric INVD and group fairness metrics DEOP ($C = 0.234$), DEOD ($C = 0.202$), and DDSP ($C = 0.182$). Notably, the group fairness metrics do not conflict with each other, implying that the assigned weight $w_{DDSP} = 0.25$ for DDSP captures the EDI committee’s preference for all three group fairness metrics. Additionally, the conflict is slightly asymmetric, meaning that selecting a Pareto optimal model for INVD results in a smaller violation of the EDI committee’s preferred group fairness metrics compared to the reverse scenario.

Figure 7 further illustrates this conflict, showing that $\mathcal{P}(\mathcal{Y}_{INVD})$ fails to approximate group fairness Pareto

Fronts, particularly in high to moderate error, low group unfairness regions of the $(F_1, DDS\!P)$ objective spaces (40-50%).

To exemplify the potential downstream social consequences of this fairness metric conflict, we compare the predictive behavior of two models: one selected from the medium error, moderate individual unfairness region of $\mathcal{P}(\mathcal{Y}_{INVD})$, and another chosen at a similar error level from $\mathcal{P}(\mathcal{Y}_{DDS\!P})$. The objective space locations of these models are also visualized in Figure 7 as “Group Fair” and “Indiv. Fair”.

Next, we perform risk-assessment on the two selected models, providing a detailed summary of these models’ predictive behavior in Figure 6. Both models exhibit relatively low group and individual unfairness (but high error) by increasing the predicted qualification rate for Black students from $P(\text{Qualified}|\text{Black}) = 0.01$ in the data to $P(\text{Accept}|\text{Black}) = 0.08$ for the individually fair model and $P(\text{Accept}|\text{Black}) = 0.13$ for the group-fair model. The individually fair model rejects 3% more qualified and 2% more unqualified Black applicants compared to the group-fair model. This improves INVD, as 2% fewer similarly unqualified Black and White applicants receive different outcomes. However, the group-fair model accepts 5% more Black applicants overall, reducing the disparity between group acceptance rates (DDS\!P) from 23% in the individually-fair model to 18%. While the group-fair model achieves a more diverse accepted class, aligning with the EDI committee’s objectives, it does so by accepting 2% more unqualified Black applicants than the individually-fair model.

Implementing such an admissions strategy may lead to the downstream consequence of *self-fulfilling prophecy* if the underprivileged students admitted through this approach struggle academically, which concerns both the University and the EDI committee.

Fairness-Aware Model Selection Having identified the fairness metric conflict between INVD and group fairness metrics, the University and EDI committee update the weights from $\mathbf{w} = \langle 0.5, 0.5 \rangle$ to include a preference for INVD to mitigate the risk of SFP. The new weights for F_1 , DDS\!P, and INVD become $\mathbf{w} = \langle 0.5, 0.2, 0.3 \rangle$. Using these objective weights, we perform model selection via single-objective scalarization (Figure 7) to perform a model selection (green cross) that achieves a moderate level of accuracy while balancing DDS\!P and INVD to mitigate the risk of SFP. This result satisfies both the University’s and EDI committee’s objectives.

This case study shows the effectiveness of ManyFairHPO in capturing and navigating the complex fairness objective landscape of the Law School Admissions problem, considering the preferences and concerns of multiple stakeholders. By providing interpretable insights into fairness metric conflicts and their potential downstream consequences, ManyFairHPO enables the University and EDI committee to make informed decisions that align with their priorities and mitigate risks.

6 Conclusion

In this study, we address the criticisms of FairML approaches that attempt to solve the fairness problem *algorithmically* by motivating the consideration of fairness metrics as conflicting objectives. Our main contribution is ManyFairHPO, a human-in-the-loop and interpretable MaO optimization framework for fairness that enables practitioners to thoroughly navigate complex fairness objective landscapes. ManyFairHPO is designed to encourage socially responsible fairness modeling decisions that encapsulate multiple fairness objectives and avoid fairness metric conflict-related risks.

Our empirical evaluation demonstrates that ManyFairHPO achieves similar Pareto fronts compared to the typical BiO optimization and even outperforms the state-of-the-art bias-mitigation technique Exponentiated Gradient Reduction (EGR). Moreover, ManyFairHPO exhibits a unique ability to explore and balance trade-offs between fairness metrics themselves, providing interpretable insights into problems-specific fairness objective landscapes. We further illustrate the practical application of ManyFairHPO through a Law School Admissions case study, guiding practitioners from fairness metric selection and risk identification to MaO model selection. The case study showcases ManyFairHPO’s effectiveness in navigating complex fairness objective landscapes, considering the preferences and concerns of multiple stakeholders, and mitigating potential downstream consequences.

One key assumption we make is that stakeholder compromise is possible in the given use case, and that a compromising solution can be found by ManyFairHPO. Through our comparison to BiO HPO and EGR we show that ManyFairHPO is largely effective at approximating fairness-accuracy and fairness-fairness Pareto Fronts. However, the ManyFairHPO’s approximation largely depends on the quality of the hyperparameter search space. As mentioned in Section 3, future extensions or applications of ManyFairHPO could see benefit from including bias mitigation approaches and their hyperparameters in the search space as seen in (Wu and Wang 2021). From a socio-technical perspective, we note that the success of ManyFairHPO as a socio-technical framework also rises and falls with honest human interaction, including preferences from a diverse set of stakeholders as well as thorough analysis and assessment of fairness metric conflicts and their related social consequences. While we believe ManyFairHPO is a valuable approach, it could have adverse impacts if mitigating one conflict-related risk leads to another unforeseen social consequence.

This work also paves the way for socio-technical discussions on how fairness metrics can be used in balance to meet complex social objectives and requirements while mitigating downstream conflict-related risks, rephrasing the question from “which fairness metric should I use?” towards “which prioritization of metrics would balance social objectives and mitigate risks?” We also encourage socio-technical studies to identify further conflict-related risks and form connections between problem characteristics and the presence, severity, and significance of fairness metric conflicts.

Acknowledgments

The authors of this work would like to thank the reviewers and the organization of the Artificial Intelligence, Ethics, and Society Conference (AIES '24) for the opportunity to share our work and receive feedback. We also would like to thank Dr. Janek Thomas for his contribution through his supervision of the original Master's Project which this work stemmed from. We would finally like to thank the Zuse School ELIZA Master's Scholarship Program for their financial and professional support of our main author.

References

- Agarwal, A.; Beygelzimer, A.; Dudík, M.; Langford, J.; and Wallach, H. 2018. A reductions approach to fair classification. In *International conference on machine learning*, 60–69. PMLR, PMLR.
- Angwin, J.; Larson, J.; Mattu, S.; and Kirchner, L. 2016. Machine Bias. *ProPublica*, May, (201barocas6): 139–159.
- Barocas, S.; Hardt, M.; and Narayanan, A. 2023. *Fairness and machine learning: Limitations and opportunities*. Cambridge, MA, USA: MIT Press.
- Bellamy, R. K.; Dey, K.; Hind, M.; Hoffman, S. C.; Houde, S.; Kannan, K.; Lohia, P.; Martino, J.; Mehta, S.; Mojsilović, A.; et al. 2019. AI Fairness 360: An extensible toolkit for detecting and mitigating algorithmic bias.
- Berk, R.; Heidari, H.; Jabbari, S.; Joseph, M.; Kearns, M.; Morgenstern, J.; Neel, S.; and Roth, A. 2017. A convex framework for fair regression. *arXiv:1706.02409*.
- Binns, R. 2018. Fairness in machine learning: Lessons from political philosophy. In *Conference on fairness, accountability and transparency*, 149–159. PMLR.
- Bischl, B.; Binder, M.; Lang, M.; Pielok, T.; Richter, J.; Coors, S.; Thomas, J.; Ullmann, T.; Becker, M.; Boulesteix, A.-L.; et al. 2023. Hyperparameter optimization: Foundations, algorithms, best practices, and open challenges. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 13(2): e1484.
- Chen, Y.; Huang, A.; Wang, Z.; Antonoglou, I.; Schrittwieser, J.; Silver, D.; and de Freitas, N. 2018. Bayesian Optimization in AlphaGo. *arXiv:1812.06855 [cs.LG]*.
- Chouldechova, A. 2016. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *arXiv:1610.07524*.
- Cooper, A. F.; and Abrams, E. 2021. Emergent Unfairness in Algorithmic Fairness-Accuracy Trade-Off Research. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '21, 46–54. New York, NY, USA: Association for Computing Machinery. ISBN 9781450384735.
- de Zwart, H. 2022. Dutch Institute for Human Rights: Use of anti-cheating software can be algorithmic discrimination (i.e. racist).
- Deb, K.; and Jain, H. 2014. An Evolutionary Many-Objective Optimization Algorithm Using Reference-Point-Based Nondominated Sorting Approach, Part I: Solving Problems With Box Constraints. *IEEE Transactions on Evolutionary Computation*, 18(4): 577–601.
- Deb, K.; Pratap, A.; Agarwal, S.; and Meyarivan, T. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE transactions on evolutionary computation*, 6(2): 182–197.
- Dhiman, G.; and Kumar, V. 2019. KnRVEA: A hybrid evolutionary algorithm based on knee points and reference vector adaptation strategies for many-objective optimization. *Applied Intelligence*, 49: 2434–2460.
- Dooley*, S.; Sukthanker*, R. S.; Dickerson, J. P.; White, C.; Hutter, F.; and Goldblum, M. 2023. Rethinking Bias Mitigation: Fairer Architectures Make for Fairer Face Recognition. In *Thirty-seventh Conference on Neural Information Processing Systems (NeurIPS 2023)*.
- Dwork, C.; Hardt, M.; Pitassi, T.; Reingold, O.; and Zemel, R. 2012. Fairness through awareness. 214–226.
- Eggensperger, K.; Müller, P.; Mallik, N.; Feurer, M.; Sass, R.; Klein, A.; Awad, N.; Lindauer, M.; and Hutter, F. 2021. HPOBench: A Collection of Reproducible Multi-Fidelity Benchmark Problems for HPO. In Vanschoren, J.; and Yeung, S., eds., *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*. Curran Associates.
- Eiben, A. E.; and Smith, J. E. 2015. *Introduction to evolutionary computing*. Springer.
- Hansen, N.; Auger, A.; Brockhoff, D.; and Tušar, T. 2022. Anytime performance assessment in blackbox optimization benchmarking. *IEEE Transactions on Evolutionary Computation*, 26(6): 1293–1305.
- Hoffmann, A. L. 2019. Where fairness fails: data, algorithms, and the limits of antidiscrimination discourse. *Information, Communication & Society*, 22(7): 900–915.
- Hsu, B.; Mazumder, R.; Nandy, P.; and Basu, K. 2022. Pushing the limits of fairness impossibility: Who's the fairest of them all? volume 35, 32749–32761.
- Islam, R.; Pan, S.; and Foulds, J. R. 2021. Can We Obtain Fairness For Free? In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '21, 586–596. New York, NY, USA: Association for Computing Machinery. ISBN 9781450384735.
- Jordan, S. M.; Cohen, D.; and Thomas, P. S. 2018. Using cumulative distribution based performance analysis to benchmark models. In *NeurIPS 2018 Workshop on Critiquing and Correcting Trends in Machine Learning*.
- Kamishima, T.; Akaho, S.; Asoh, H.; and Sakuma, J. 2012. Fairness-aware classifier with prejudice remover regularizer. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2012, Bristol, UK, September 24-28, 2012. Proceedings, Part II* 23, 35–50. Springer.
- Karl, F.; Pielok, T.; Moosbauer, J.; Pfisterer, F.; Coors, S.; Binder, M.; Schneider, L.; Thomas, J.; Richter, J.; Lang, M.; et al. 2022. Multi-Objective Hyperparameter Optimization—An Overview. *arXiv preprint arXiv:2206.07438*.
- Kleinberg, J.; Mullainathan, S.; and Raghavan, M. 2016. Inherent Trade-Offs in the Fair Determination of Risk Scores. *arXiv:1609.05807*.

Martinez, N.; Bertran, M.; and Sapiro, G. 2020. Minimax Pareto Fairness: A Multi Objective Perspective. In III, H. D.; and Singh, A., eds., *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 6755–6764. PMLR.

Miconi, T. 2017. The impossibility of” fairness”: a generalized impossibility result for decisions. *arXiv preprint arXiv:1707.01195*.

Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; et al. 2011. Scikit-learn: Machine learning in Python. *the Journal of machine Learning research*, 12: 2825–2830.

Perrone, V.; Donini, M.; Zafar, M. B.; Schmucker, R.; Kenthapadi, K.; and Archambeau, C. 2021. Fair bayesian optimization. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, 854–863.

Pessach, D.; and Shmueli, E. 2022. A review on fairness in machine learning. *ACM Computing Surveys (CSUR)*, 55(3): 1–44.

Robertson, J.; Stinson, C.; and Hu, T. 2022. A Bio-Inspired Framework for Machine Bias Interpretation. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, 588–598.

Ruf, B.; and Detyniecki, M. 2021. Towards the Right Kind of Fairness in AI. .

Schmucker, R.; Donini, M.; Perrone, V.; and Archambeau, C. 2020. Multi-objective multi-fidelity hyperparameter optimization with application to fairness.

Schöffer, J. 2023. On the Interplay of Transparency and Fairness in AI-Informed Decision-Making.

Selbst, A. D.; Boyd, D.; Friedler, S. A.; Venkatasubramanian, S.; and Vertesi, J. 2019. Fairness and abstraction in sociotechnical systems. In *Proceedings of the conference on fairness, accountability, and transparency*, 59–68.

Sharma, S.; and Kumar, V. 2022. A comprehensive review on multi-objective optimization techniques: Past, present and future. *Archives of Computational Methods in Engineering*, 29(7): 5605–5633.

Vanschoren, J.; van Rijn, J.; Bischl, B.; and Torgo, L. 2014. OpenML: Networked Science in Machine Learning. *SIGKDD Explor. Newsl.*, 15(2): 49–60.

Weerts, H.; Pfisterer, F.; Feurer, M.; Eggenberger, K.; Bergman, E.; Awad, N.; Vanschoren, J.; Pechenizkiy, M.; Bischl, B.; and Hutter, F. 2023. Can Fairness be Automated? Guidelines and Opportunities for Fairness-aware AutoML. *arXiv preprint arXiv:2303.08485*.

Wightman, L. F. 1998. LSAC National Longitudinal Bar Passage Study. LSAC Research Report Series. In *Non-Journal*, 111.

Wu, Q.; and Wang, C. 2021. FairAutoML: Embracing Unfairness Mitigation in AutoML. *arXiv preprint arXiv:2111.06495*.