

When Interpretability Is Unequally Distributed: Fairness in Hybrid Interpretable Models

Ziba Jabbar Zare¹, Ulrich Aïvodji², Julien Ferry³, Thibaut Vidal³

¹Department of Mechanical, Industrial and Aerospace Engineering, Concordia University, Canada

²École de technologie supérieure, Mila, Canada

³CIRRELT & SCALE-AI Chair in Data-Driven Supply Chains, Department of Mathematics and Industrial Engineering, Polytechnique Montréal, Canada

ziba.jabbarzare@mail.concordia.ca, ulrich.aivodji@etsmtl.ca, julien.ferry@polymtl.ca, thibaut.vidal@polymtl.ca

Abstract

Hybrid interpretable models combine a transparent component with a black-box model by assigning some examples to the former and deferring the rest to the latter. While this design enables flexible tradeoffs between accuracy and interpretability, it also raises a distinct procedural fairness concern: some demographic groups may systematically receive interpretable decisions, while others are disproportionately routed to a black box. We formalize this issue as *Interpretability Coverage Disparity* (ICD), a demographic-parity-style measure applied to the routing decision of hybrid interpretable models. Using tools from predictive multiplicity, we study ICD across four hybrid interpretable learning methods, three standard fairness benchmark datasets, and multiple sensitive attributes. Our experiments reveal substantial ICD in intermediate transparency regimes, where both the interpretable and black-box components are actively used. We further show that simple coverage-disparity constraints can significantly reduce ICD in exact hybrid learning methods, with marginal impact on accuracy and sparsity. In several settings, ICD mitigation also improves standard algorithmic fairness metrics. These results show that hybrid interpretable models should be audited not only for predictive fairness, but also for how they allocate interpretability across individuals and groups.

Code — github.com/vidalt/Fair-Hyb-Interpretability

Paper with appendices — arxiv.org/abs/2605.28626

1 Introduction

Explainability broadly refers to the accessibility and comprehensibility of information about how an AI system operates. It is increasingly emphasized in regulatory frameworks, particularly for high-risk applications such as hiring, healthcare, and criminal justice, where it is essential to support informed decision-making and enable meaningful human oversight. A common strategy to enhance explainability in machine learning is the use of post-hoc explanation techniques. These methods aim to approximate or explain the behavior of complex black-box models after training, often by using surrogate models that provide a simpler and more interpretable representation of the underlying decision process (Guidotti et al. 2019). However, post-hoc explanations

can be unreliable and susceptible to manipulation, raising concerns about their faithfulness (Aïvodji et al. 2019).

An alternative is to design inherently interpretable models, whose structure directly reveals how inputs are mapped to predictions. Examples include decision trees (Breiman et al. 1984), rule lists (Rivest 1987), rule sets (Rijnbeek and Kors 2010), and scoring systems (Ustun and Rudin 2016). While these models are generally easier for humans to understand, they may exhibit reduced predictive performance in complex settings. This limitation has motivated research on *hybrid interpretable models* (Pan, Wang, and Hara 2020), which combine an interpretable component with a black-box component. A gating mechanism assigns each example to either the interpretable component or the black box. By varying the proportion of examples assigned to the interpretable component, often referred to as the model’s *transparency*, hybrid interpretable models can realize different trade-offs between interpretability and predictive performance. However, this design also raises a largely overlooked question: even when hybrid models improve the aggregate accuracy-interpretability trade-off, they may distribute interpretability unevenly across individuals or demographic groups.

A large body of work on fairness in machine learning focuses on fairness criteria defined with respect to a predictive model’s outputs. Existing approaches typically intervene at one of three points in the machine learning pipeline: by modifying the training data through preprocessing, by incorporating fairness constraints during model training, or by adjusting model outputs through post-processing (Kamiran and Calders 2009; Zemel et al. 2013; Hardt, Price, and Srebro 2016; Zhang, Lemoine, and Mitchell 2018). Although it has received considerably less attention, fairness has also been studied with respect to other dimensions of predictive models. In particular, fairness in explanation quality has been examined for several post-hoc explanation methods, including feature-attribution methods (Dai et al. 2022; Balagopalan et al. 2022) and counterfactual explanations (Bell et al. 2024). These works reveal demographic disparities in explanation quality, measured through properties such as sparsity, robustness, time to recourse, and effort to recourse. They show that fairness issues can arise not only from what predictions are made, but also from how decisions are explained or acted upon.

Our work follows this procedural perspective, but studies

a distinct source of disparity arising in hybrid interpretable models. In these models, interpretability is not provided by a separate post-hoc explanation method; it is allocated by the model itself through its gating mechanism. This raises the question of whether different individuals or demographic groups have comparable access to the interpretable component. Our main contributions are as follows:

- We introduce *Interpretability Coverage Disparity* (ICD), a demographic-parity-style disparity measure defined over interpretability assignment in hybrid interpretable models. We also leverage notions from predictive multiplicity to characterize individual-level *Interpretability Coverage Arbitrariness* (ICA) in whether examples are assigned to the interpretable or black-box component.
- We conduct extensive experiments over sets of high-performing hybrid interpretable models learned by four state-of-the-art methods. Our results show that substantial ICA or ICD can arise on real-world datasets, leading to situations in which (i) some individuals are arbitrarily assigned to the black-box component, and (ii) some demographic groups are predominantly routed to the black-box component, while others are predominantly covered by the interpretable component.
- We extend existing hybrid interpretable model learning methods to mitigate ICD, and investigate the impact of the resulting constraints on model sparsity, predictive performance, and algorithmic fairness.

The remainder of the paper is organized as follows. We first review the relevant background and related work in Sections 2 and 3, respectively. We then formalize our proposed group-level and individual-level metrics for quantifying unfairness in interpretability coverage in Section 4. In Section 5, we empirically investigate the extent to which ICD and ICA arise when state-of-the-art hybrid interpretable models are applied to real-world datasets. Next, in Section 6, we address ICD by adapting existing learning methods and assessing the empirical tradeoffs induced by mitigation. Finally, we conclude in Section 7.

2 Technical Background

Notation. Let $\mathcal{X}$ denote the input space and $\mathcal{Y}$ the output space. Given a dataset $S = \{(\mathbf{x}_i, y_i)\}_{i=1}^n \subset \mathcal{X} \times \mathcal{Y}$, the objective of a supervised learning algorithm is to learn a function $h : \mathcal{X} \rightarrow \mathcal{Y}$ from a hypothesis space $\mathcal{H}$ by minimizing an empirical loss, such as the 0/1 loss:

$$\ell_{0/1}(h; S) = \frac{1}{n} \sum_{i=1}^n \mathbb{1}(h(\mathbf{x}_i) \neq y_i).$$

Following prior work on hybrid interpretable models and algorithmic fairness, we focus on binary classification tasks, i.e., $\mathcal{Y} = \{0, 1\}$. This choice allows us to rely on established methods and metrics from these research areas while leaving our main contribution unchanged: ICD concerns whether a decision is handled by the interpretable component, not the predicted outcome itself.

```

if ["Prior-Crimes=0" and "Age>=30"] then 0
else if ["Prior-Crimes>5" and "Age=24-30"] then 1
else if ["Prior-Crimes=1-3" and "Age>=30"] then 0
else
  RandomForest()

```

Figure 1: Example hybrid interpretable model from Ferry, Laberge, and Aïvodji (2024), trained on the COMPAS dataset (for recidivism prediction).

Hybrid Interpretable Models. A hybrid interpretable model is a predictive model defined by a triplet $\langle \Omega, h_s, h_c \rangle$ in which $\Omega \subset \mathcal{X}$ defines the *interpretable region* of the input space, in which examples are handled by a simple component h_s belonging to some interpretable hypothesis class $\mathcal{H}_s$, and examples outside Ω are deferred to a black-box model h_c from a more complex hypothesis class $\mathcal{H}_c$, i.e.,:

$$\forall \mathbf{x} \in \mathcal{X}, \quad \langle \Omega, h_s, h_c \rangle(\mathbf{x}) = \begin{cases} h_s(\mathbf{x}) & \text{if } \mathbf{x} \in \Omega, \\ h_c(\mathbf{x}) & \text{otherwise.} \end{cases}$$

The *transparency* of such a model on a given dataset S is then the proportion of examples covered by the interpretable part:

$$C_\Omega = \frac{|\Omega \cap S|}{n},$$

such that a pure black-box corresponds to $C_\Omega = 0$ while a single interpretable model to $C_\Omega = 1$.

In practice, existing hybrid interpretable models (Pan, Wang, and Hara 2020; Wang and Lin 2021; Ferry, Laberge, and Aïvodji 2024) use a set of rules r as the interpretable component, which simultaneously defines the interpretable region Ω as the support of its rules: any example not captured by at least one rule is deferred to the black-box. We use S_r to denote the subset of S captured by the rules in r , such that the transparency can be written as $C_r = \frac{|S_r|}{n}$. Figure 1 illustrates a hybrid interpretable model whose interpretable component h_s is a rule list, and whose support defines the interpretable region Ω .

Algorithmic Fairness. Undesirable disparities in predictive models may stem from biased data, historical inequalities, or biases introduced during model development (Mehrabi et al. 2022). Several notions have been proposed to quantify and mitigate such disparities. They are commonly divided into *individual fairness*, which requires similar individuals to receive similar predictions (Dwork et al. 2012), and *group fairness*, which requires some form of parity across demographic groups. The latter has received particular attention because it can be operationalized through aggregate statistics over protected groups, and because it echoes legal notions of disparate impact in anti-discrimination law (uni 1978; Feldman et al. 2015; Barocas and Selbst 2016).

In our experiments, we will assess the effect of ICD mitigation on algorithmic fairness using two popular group fairness metrics. Let $\mathcal{P}$ denote a set of protected groups, where each $p \in \mathcal{P}$ defines a subset $S_p \subseteq S$ of examples sharing a

given sensitive attribute value, such as gender, race, or age. The maximum statistical parity (SP) (Dwork et al. 2012) disparity is

$$d_{SP}(h, S) := \max_{(p, p') \in \mathcal{P}^2} \left(\frac{\sum_{i \in S_p} h(\mathbf{x}_i)}{|S_p|} - \frac{\sum_{i \in S_{p'}} h(\mathbf{x}_i)}{|S_{p'}|} \right).$$

It measures the largest difference in positive prediction rates across protected groups. We also consider maximum equal opportunity (EO) (Hardt, Price, and Srebro 2016) disparity, defined as

$$d_{EO}(h, S) := \max_{(p, p') \in \mathcal{P}^2} \left(\frac{\sum_{i \in S_p^+} h(\mathbf{x}_i)}{|S_p^+|} - \frac{\sum_{i \in S_{p'}^+} h(\mathbf{x}_i)}{|S_{p'}^+|} \right),$$

where $S_p^+ = \{i \in S_p : y_i = 1\}$ denotes the set of positive examples in group p . This metric measures the largest difference in true positive rates across protected groups.

Predictive Multiplicity. In many settings, several models can achieve nearly indistinguishable predictive performance while disagreeing on the predictions assigned to individual instances. This coexistence of competing, similarly accurate models is known as *predictive multiplicity*. The set of such near-optimal models is referred to as the *Rashomon set* (Breiman 2001):

$\mathcal{R}(\mathcal{H}, \varepsilon, S) := \{h \in \mathcal{H} : \ell_{0/1}(h; S) \leq (1+\varepsilon) \cdot \ell_{0/1}(h^*; S)\}$, where h^* is a reference model minimizing $\ell_{0/1}$ over S and ε controls the tolerated loss increase relative to h^* . Although models in the Rashomon set have comparable predictive performance, they can differ substantially with respect to other desiderata. Assessing these desiderata across $\mathcal{R}(\mathcal{H}, \varepsilon, S)$ therefore provides a more comprehensive evaluation than focusing on a single model, which may be only one arbitrary representative of this set.

Some *enumeration-based* methods have been proposed to exactly construct $\mathcal{R}(\mathcal{H}, \varepsilon, S)$ for specific and relatively simple hypothesis classes, including rule lists (Mata, Kanamori, and Arimura 2022), rule sets (Ciaperoni, Xiao, and Gionis 2024), and decision trees (Xin et al. 2022; Arslan et al. 2025). However, existing methods for rule-based models tend to scale poorly with the Rashomon parameter ε , the number of candidate rules, and the maximum number of rules. Some *enumeration-free* methods have therefore been proposed to characterize properties of interest within the Rashomon set of a given hypothesis class, without explicitly constructing the full set. These methods, however, are again limited to particular properties and hypothesis classes, such as linear models (Coker, Rudin, and King 2021; Watson-Daniels, Parkes, and Ustun 2023; Fisher, Rudin, and Dominici 2019; Coston, Rambachan, and Chouldechova 2021), additive models (Zhong et al. 2024), or models whose learning problem can be cast as a mathematical programming formulation (Langlade et al. 2025).

For more complex models, including black-box models, simple model-sampling tools have been shown to provide useful approximations of the Rashomon set, notably through bootstrapping (Cooper et al. 2024). Although these methods typically recover only a subset of the full Rashomon set, they still allow meaningful statistics to be computed over collections of *good* models.

3 Related Works

While fairness and explainability have traditionally been studied as separate dimensions of trustworthy AI, recent work has begun to examine their intersection. In particular, an emerging line of research investigates whether explanations themselves are fair, namely, whether their quality, reliability, and informativeness are consistent across individuals and demographic groups. This notion, referred to as *fairness in explanation quality*, was formalized by Dai et al. (2022). Subsequent work has shown that explanation methods can exhibit disparities across subgroups, even when the underlying predictive model is fair. For instance, Balagopalan et al. (2022) analyze subgroup differences in post-hoc explanations, while Dhaini et al. (2025) demonstrate the presence of gender-based disparities in feature attribution methods for language models. Furthermore, Zhao, Wang, and Derr (2022) propose optimization-based approaches to mitigate disparities in explanation quality.

A closely related line of work examines fairness through the lens of *algorithmic recourse*, also known as *recourse disparity*. These works examine whether individuals from different demographic subgroups have equally feasible and actionable paths to overturn an unfavorable prediction. In the context of counterfactual explanations, recourse cost is typically defined as the distance between an input instance and a corresponding counterfactual (or the decision boundary), and recourse fairness is characterized by disparities in average recourse cost across groups (Gupta et al. 2019; Slack et al. 2021; Kavouras et al. 2023). Empirical studies have shown that such disparities can arise even when the predictive performance of the model at hand is comparable across groups, which motivated recent works on mitigating them (Artelt and Hammer 2022; Ezzeddine et al. 2026).

Together, these developments highlight an important distinction between *outcome-based fairness*, which focuses on the fairness of model predictions, and *procedure-oriented fairness*, which concerns the fairness of the mechanisms through which decisions are produced, explained, or acted upon. Fairness in explanation quality and recourse fairness both fall into the latter category, as they evaluate whether the reasoning provided by an explanation method is equally reliable and informative across groups.

In this work, we introduce complementary procedural fairness notions tailored to hybrid interpretable models: *Interpretability Coverage Disparity* (ICD) and *Interpretability Coverage Arbitrariness* (ICA). Unlike prior work on explanation quality or recourse, ICD and ICA do not assess the quality of an explanation produced after a prediction. Instead, they measure how interpretability itself is allocated by the model, through the routing decision that determines whether an example is handled by the interpretable component or deferred to the black-box component. This distinction is important because, in hybrid interpretable models, access to interpretability is not provided by an external explanation method; it is an intrinsic property of the predictive model. ICD therefore shifts attention from whether explanations are equally good to whether different individuals and groups are equally likely to receive interpretable decisions in the first place.

4 Interpretability Coverage Unfairness

We now formalize metrics to quantify unfairness in interpretability coverage for hybrid interpretable models. We first consider disparities at the group level, before turning to arbitrariness at the individual level.

4.1 Group-level Disparity

We first introduce the *Interpretability Coverage* of a protected group $p \in \mathcal{P}$ as the fraction of individuals from group p that fall within the interpretable region Ω :

$$\text{IC}_p := \frac{|\Omega \cap S_p|}{|S_p|}. \quad (1)$$

Higher values of IC_p indicate that individuals from group p are more frequently handled by the interpretable component of the hybrid interpretable model. Note that the overall model *transparency*, defined in Section 2, corresponds to interpretability coverage measured over the entire dataset.

We now define the *Interpretability Coverage Disparity* (ICD) as the maximum pairwise difference in interpretability coverage across protected groups:

$$\text{ICD} := \max_{(p,p') \in \mathcal{P}^2} (\text{IC}_p - \text{IC}_{p'}). \quad (2)$$

Equivalently, ICD is the difference between the largest and smallest group-specific interpretability coverage:

$$\text{ICD} = \max_{p \in \mathcal{P}} \text{IC}_p - \min_{p \in \mathcal{P}} \text{IC}_p. \quad (3)$$

ICD captures the largest gap in access to interpretability between any two protected groups. A value of zero indicates that all groups are covered by the interpretable component at the same rate, whereas larger values indicate that some groups are disproportionately routed to the black-box component. In our experiments, we also report the per-group interpretability coverage when relevant, since ICD alone does not identify which groups are relatively advantaged or disadvantaged, and only reflects the largest observed disparity.

Importantly, ICD depends only on the allocation of individuals to the interpretable and black-box regions, and not on the predictions made by either component of the hybrid model. It is therefore a procedural fairness notion: rather than measuring disparities in predictive outcomes, it quantifies the extent to which individuals from different protected groups have unequal access to interpretable decisions.

For rule-based hybrid interpretable models, this notation can be made more explicit. Let $S_{r,p}$ denote the set of training examples from protected group $p \in \mathcal{P}$ that are classified by the rules in r . We thus have $\text{IC}_p = \frac{|S_{r,p}|}{|S_p|}$ and:

$$\text{ICD} := \max_{(p,p') \in \mathcal{P}^2} \left(\frac{|S_{r,p}|}{|S_p|} - \frac{|S_{r,p'}|}{|S_{p'}|} \right). \quad (4)$$

Although ICD is primarily intended for a given set of protected groups, extending it to settings such as intersectional fairness could be insightful. The main challenge, namely efficiently navigating exponentially many subgroup intersections, is not ICD-specific, so existing fairness tools may be adaptable. We leave this extension to future work.

4.2 Individual-level Arbitrariness

Inspired by previous work on predictive multiplicity, we further characterize the *arbitrariness* of interpretability coverage at the level of an individual example $\mathbf{x}$. Let $\mathcal{H}$ denote the hypothesis class of hybrid interpretable models, and let $\mathcal{R}(\mathcal{H}, \varepsilon, S)$ be its ε -Rashomon set on a dataset S . We first define the interpretability coverage frequency of $\mathbf{x}$ as

$$\text{ICF}(\mathbf{x}) := \frac{|\{\langle \Omega, h_s, h_c \rangle \in \mathcal{R}(\mathcal{H}, \varepsilon, S) : \mathbf{x} \in \Omega\}|}{|\mathcal{R}(\mathcal{H}, \varepsilon, S)|}. \quad (5)$$

This quantity measures how often $\mathbf{x}$ is handled by the interpretable component across near-optimal hybrid interpretable models. We then define the *Interpretability Coverage Arbitrariness* (ICA) of $\mathbf{x}$ as

$$\text{ICA}(\mathbf{x}) := 1 - 2 \left| \text{ICF}(\mathbf{x}) - \frac{1}{2} \right|. \quad (6)$$

Values close to 0 indicate low arbitrariness, since $\mathbf{x}$ is consistently assigned either to the black-box component or to the interpretable component across the Rashomon set. In contrast, values close to 1 indicate high arbitrariness, since $\mathbf{x}$ is assigned to the interpretable region for roughly half of the models and to the black-box region for the other half.

In practice, building the exact Rashomon set is not tractable for black-box models, and even more so for hybrid interpretable models, whose interpretable component training often depends on a black-box trained beforehand. Accordingly, we replace $\mathcal{R}(\mathcal{H}, \varepsilon, S)$ in Definitions (5) and (6) with an approximation $\tilde{\mathcal{R}}(\mathcal{H}, \varepsilon, S)$. This approximation, described in Section 5, does not aim to enumerate all near-optimal models, but rather to provide a large and diverse collection of high-performing models for estimating interpretability-coverage statistics.

5 Experimental Evaluation of ICD and ICA

We now assess to what extent ICD and ICA arise when using state-of-the-art hybrid interpretable models trained over real-world datasets. We first describe the considered setup before presenting the main results.

5.1 Experimental Setup

Datasets. In our experiments, we consider three datasets of varying sizes and prediction tasks. All datasets are pre-processed by discretizing numerical features using quantile-based binning and applying one-hot encoding to categorical variables. To ensure consistency across hybrid model implementations and avoid dependence on the rule generation procedure, we pre-mine a set of candidate rules using the FP-Growth algorithm (Han, Pei, and Yin 2000). The resulting rules have a maximum cardinality of 2 and a minimum support threshold of 1%. Following the setup of Ferry, Laberge, and Aïvodji (2024), we retain a maximum of 300 rules, keeping the ones with the largest support.

All group-level disparities, including ICD and standard algorithmic fairness metrics, are evaluated with respect to three sensitive attributes: Age, Gender, and Race. For each attribute, we consider up to three subgroups corresponding to the most frequent categories in the data, grouping less frequent categories into an “Other” category when necessary:

- **COMPAS** (Angwin et al. 2016): (7,214 samples, 230 mined rules) contains records of criminal defendants from Florida (2013–2014). The prediction task is to determine whether a defendant will recidivate within two years. Demographic subgroups are defined as follows: Age: {18–25, 26–29, ≥ 30 }, Gender: {Male, Female}, Race: {African-American, Caucasian, Other}.
- **UCI Adult Income** (Dua and Graff 2017): (48,842 samples, 300 mined rules) consists of demographic and occupational attributes from the 1994 U.S. Census. The task is to predict whether an individual earns more than \$50K per year. Demographic subgroups are defined as: Age: {Low, Middle, High} (quantile-based bins), Gender: {Male, Female}, Race: {White, Black, Other}.
- **ACS Employment** (Ding et al. 2021): (208,358 samples, 300 mined rules) is derived from the American Community Survey (2014–2018) and aims to predict whether an individual is employed based on socioeconomic features. Demographic subgroups are defined as: Age: {Low, Middle, High} (quantile-based bins), Gender: {Male, Female}, Race: {White alone, Black or African American alone, Other}.

Learning algorithms. We consider four state-of-the-art hybrid interpretable learning methods: Hybrid Rule Set (HyRS), Companion Rule List (CRL), HybridCORELSPre, and HybridCORELSPost. HyRS (Wang and Lin 2021) and CRL (Pan, Wang, and Hara 2020) are heuristic methods based on simulated annealing. Both take as input a pre-trained black-box model and learn a rule-based interpretable component that covers part of the input space, while deferring the remaining instances to the black-box. They optimize trade-offs between accuracy, transparency, and sparsity, but do not provide optimality guarantees and cannot enforce a prescribed transparency level.

By contrast, HybridCORELSPre and HybridCORELSPost (Ferry, Laberge, and Aïvodji 2024) rely on an exact branch-and-bound procedure to explore ordered rule prefixes. Given a pre-trained black-box, HybridCORELSPost optimizes a weighted sum of hybrid-model accuracy and sparsity subject to a hard transparency constraint. HybridCORELSPre follows the same principle, but learns the interpretable prefix before training the black-box component; it therefore optimizes an upper bound on the final hybrid-model accuracy when fitting the prefix.

Hyperparameters. Across all datasets, our experimental setup is designed to explore a broad range of transparency levels. We therefore vary only the hyperparameters that directly control, or strongly influence, interpretability coverage, while keeping the remaining hyperparameters at their default values. For HybridCORELSPre and HybridCORELSPost, transparency is explicitly controlled via a minimum-transparency constraint. We consider 10 values, $\text{min_transp} \in \{0.1, 0.2, \dots, 0.9, 0.95\}$. For HyRS, we vary the hyperparameter β , which governs the trade-off between accuracy and transparency. We consider 10 values logarithmically spaced between 10^{-3} and 10^0 , following Wang and Lin (2021). For CRL, we vary the regulariza-

tion parameter λ , which controls model sparsity and consequently affects transparency. We consider 10 logarithmically spaced values in the range $[10^{-3}, 10^{-1}]$.

The black-box component of each hybrid interpretable model is instantiated as a random forest classifier from the scikit-learn library (Pedregosa et al. 2011), with `min_samples_split` = 10 and `max_depth` = 10.

Bootstrapping setup. To better quantify the extent to which ICD or ICA can arise, we evaluate them over sets of *good models*. Since exactly enumerating the Rashomon set of hybrid models is computationally infeasible, we approximate it using a bootstrap-based procedure. Since this procedure approximates the true Rashomon set, the resulting model distribution provides an estimate of the true distribution, and our results should be interpreted accordingly.

We first split each dataset into training (80%) and test (20%) sets. For each method and hyperparameter configuration, we first train a reference model on the original training set. We then generate 1000 bootstrap samples from the training set and retrain the same configuration on each resampled dataset. This yields 1001 trained models per configuration, including the reference model. Since we consider 10 configurations per method, this results in 10,010 trained models for each method. For each bootstrap run, the black-box component is trained with a different random seed.

The models obtained through this procedure are not necessarily unique. We therefore post-process the set of hybrid interpretable models obtained for each learning algorithm to remove duplicates. In practice, we consider two hybrid interpretable models from the same learning algorithm to be duplicates when their interpretable components are identical, that is, composed of the same rules in the same order, and when their black-box components produce similar predictions on the training set. Finally, the resulting unique models are grouped into four bins according to their training-set transparency values and filtered to keep only those whose training accuracy satisfies the Rashomon parameter ε (relative to the reference model within each bin). All downstream analyses are performed on the held-out test set.

Using the described bootstrap sampling procedure, we obtain large and diverse approximate Rashomon sets across transparency bins for all evaluated hybrid interpretable methods. Additional analyses provided in Figures 8–11 in Appendix B further illustrate how the size of the approximate Rashomon sets evolves as the Rashomon parameter ε increases, highlighting that this approach is able to recover a substantial set of near-optimal models, enabling meaningful analyses of predictive multiplicity and fairness properties over Rashomon sets.

Computational resources. To keep experiments computationally manageable, we impose a 5-minute limit on the training of each interpretable component, run on a single thread, with a maximum memory usage of 8 GB. The training of the black-box component is not limited, but typically takes only a few seconds. All experiments are run on a homogeneous set of cluster nodes equipped with Intel 6972P CPUs at 2.4 GHz.




Figure 2: Distribution of test set ICD across Rashomon sets for all transparency bins ($\varepsilon = 0.01$). Results are shown for all datasets and sensitive attributes (Age, Gender, Race), ordered from left to right columns.

5.2 Results

We now present the main empirical findings through a subset of representative results. Unless otherwise specified, all analyses are performed on approximate Rashomon sets with $\varepsilon = 0.01$. Comprehensive results for all hybrid interpretable methods, datasets, sensitive attributes, and experimental settings are deferred to the Appendix C.

Result 1. Hybrid interpretable models exhibit significant disparities in interpretability coverage across demographic groups. Figure 2 provides a comprehensive overview of the distribution of ICD across transparency bins using box plots. As defined in Section 4, ICD reflects the largest observed subgroup difference in interpretability access for a given model. Each box plot summarizes the distribution of ICD values across Rashomon sets, where the box indicates the interquartile range, the horizontal line represents the median, the triangle denotes the mean, and the whiskers correspond to the minimum and maximum values.

The results reveal substantial interpretability coverage disparities across nearly all experimental settings, confirming that hybrid interpretable models can distribute interpretability unevenly across demographic groups. In several cases, ICD values approach 1.0, indicating extreme disparities where certain demographic groups are systematically




Figure 3: Distribution of test set Interpretability Coverage (IC) for demographic subgroups of the Race attribute on the COMPAS dataset, for the different transparency bins, across Rashomon sets of HyRS ($\varepsilon = 0.01$). The red curve represents the average ICD across models within each bin, and the shaded region indicates the standard deviation.

deferred to the black-box while others are always handled by the interpretable component of the hybrid model.

To further investigate whether the observed disparities consistently favor some demographic subgroups in receiving access to interpretable decisions, Figure 3 presents the distri-




Figure 4: Distribution of test set ICA across Rashomon sets for all transparency bins ($\varepsilon = 0.01$). Results are shown for all datasets.

bution of interpretability coverage (IC) across demographic subgroups of the Race attribute for the HyRS method on the COMPAS dataset. Analogous figures for the remaining hybrid interpretable methods, datasets, and sensitive attributes are provided in Appendix C. The results show that within the same transparency bin, certain subgroups consistently receive higher interpretability coverage than others, indicating that some populations are more frequently routed to the interpretable component of the hybrid model than others. For example, as seen in Figure 3, the African-American subgroup generally exhibits lower interpretability coverage, with smaller maximum and average IC values compared to the other Race groups across transparency bins.

Another important observation from Figures 2 and 3 concerns how ICD evolves as model transparency increases. According to the mechanism of hybrid interpretable models, the lowest and highest transparency bins correspond to regimes in which predictions are largely dominated by a single component. In the first bin, most instances are routed to the black-box model, whereas in the last bin, most instances are assigned to the interpretable component. The intermediate transparency bins (transparency values between 0.25 and 0.75) therefore better characterize the truly hybrid regime, where both components actively contribute to prediction. The largest ICD values are observed in these bins, as the curves tend to have a *bell-like* shape, where ICD increases from low to intermediate transparency bins and then decreases or stabilizes at higher transparency bins.

To statistically validate this behavior, we conduct a non-parametric pairwise testing procedure across transparency bins. For each dataset, method, and sensitive attribute, we compare the distributions of ICD between adjacent bins, namely (Q_1 vs. Q_2), (Q_2 vs. Q_3), and (Q_3 vs. Q_4). We employ the Mann–Whitney U test, which compares two distributions without assuming normality, to test the null hypothesis that the ICD distributions of adjacent bins are not significantly different (Mann and Whitney 1947). Since multiple comparisons are performed within each experimental setting, Holm’s correction is applied to control the family-wise error rate (Holm 1979). For each adjacent comparison, we determine whether ICD exhibits a statistically significant increase, a statistically significant decrease, or no significant change at ($p \leq 0.05$) after correction.

Method	$\varepsilon = 0.01$		$\varepsilon = 0.05$	
	Bell-like	Mixed	Bell-like	Mixed
HybridCORELSPost	89%	11%	78%	22%
HybridCORELSPre	78%	22%	100%	0%
HyRS	89%	11%	100%	0%
CRL	67%	33%	100%	0%

Table 1: Prevalence of statistically significant bell-like ICD patterns for two Rashomon parameters ε .

Table 1 summarizes, for each method and for two Rashomon parameters ($\varepsilon = 0.01$ and $\varepsilon = 0.05$), the proportion of experimental settings exhibiting a *bell-like* pattern in a statistically significant manner, all remaining transition patterns being categorized as *mixed*. Overall, *bell-like* patterns are frequently observed across datasets and methods, supporting the hypothesis that ICD is amplified in intermediate transparency regimes, where hybridization is most pronounced, and decreases when predictions become increasingly dominated by one component or the other. Because the size of the approximate Rashomon set depends directly on ε , smaller tolerance values yield relatively fewer Rashomon models in higher transparency bins, limiting the statistical power of the pairwise comparisons and making bell-like patterns harder to detect reliably. Increasing the tolerance to $\varepsilon = 0.05$ yields larger Rashomon sets and, consequently, more stable statistical comparisons across bins. Consequently, the bell-like trend is strengthened when a richer set of near-optimal models is available for analysis.

Result 2. Hybrid interpretable models exhibit assignment arbitrariness across Rashomon sets. Beyond group-level disparities in interpretability access, Figure 4 displays the distribution of ICA values over all test set examples for the three considered datasets and each transparency bin. The results reveal that hybrid interpretable models also exhibit assignment arbitrariness at the individual level. More precisely, arbitrariness is generally limited in the lowest and highest transparency bins, where predictions are predominantly handled by a single component of the hybrid model. In contrast, substantially higher arbitrariness emerges in intermediate transparency bins, where both interpretable and




Figure 5: Distribution of our proposed metrics across Rashomon sets over transparency bins ($\varepsilon = 0.01$) for HybridCORELSPre and HybridCORELSPost, with or without our proposed ICD mitigation ($\eta = 0.05$). Each subplot corresponds to a mitigation applied to a sensitive attribute (Age, Gender, Race), ordered from left to right. Results are reported for the COMPAS dataset.

black-box components contribute to prediction. In these regimes, a larger proportion of individuals exhibit instability in whether they receive interpretable or black-box decisions across equally accurate models. This behavior is observed across all methods, supporting the hypothesis that hybridization amplifies assignment uncertainty.

6 Mitigating ICD

6.1 Methodology

Modified Learning Algorithms. The experimental analysis in the previous section revealed substantial ICD across demographic groups, methods, and datasets, as well as high variability within Rashomon sets over transparency bins. These observations suggest that near-optimal models with similar predictive performance can nonetheless induce substantially different levels of ICD. This variability further motivates the use of predictive multiplicity to explicitly search for and select fairer models while preserving high predictive performance. To this aim, we mitigate the observed ICD by incorporating maximum ICD constraints into HybridCORELS (Ferry, Laberge, and Aïvodji 2024) branch-and-bound algorithms. In a nutshell, given a set of protected groups $\mathcal{P}$ and a maximum ICD value η , each time a new best prefix is found, we check whether it satisfies the maximum ICD constraint. If it does, the prefix is retained and used to update the current incumbent; otherwise, it is discarded. This integration follows the same principle as the overall transparency constraint used in HybridCORELS. Importantly, the bounds and symmetry-aware data structures used in HybridCORELSPre and HybridCORELSPost remain valid and tight under the new constraint. The corresponding pseudocodes are provided in Appendix A.

Experimental Setup. We run experiments for both HybridCORELSPre and HybridCORELSPost, following the setup introduced in Section 5.1. These experiments have two main objectives. First, we assess whether ICD mitigation is effective, and whether enforcing ICD constraints on the bootstrapped training sets also translates into improved ICD on the test set. Second, we evaluate the effect of this mitigation on other dimensions: we report sparsity, measured as the number of rules in the prefix of each hybrid interpretable model, as well as test-set accuracy and algorithmic fairness, measured using SP and EO. Similar to Section 5.2, all experiments are conducted on approximate Rashomon sets with $\varepsilon = 0.01$ across transparency bins. Unless otherwise specified, mitigation is performed by enforcing maximum ICD constraints during training with $\eta = 0.05$, while all reported metrics are evaluated on the test set to assess the generalization ability of the proposed mitigation approach.

6.2 Results

Comprehensive results for all datasets, sensitive attributes, and experimental settings are provided in Appendix D. We hereafter report the key trends using a subset of these results.

Result 3. ICD mitigation generalizes effectively to unseen test data. Figure 5a illustrates the distribution of test set ICD within Rashomon sets of different transparency bins, on the COMPAS dataset, with or without our proposed ICD mitigation. Across all sensitive attributes, mitigation consistently reduces both the magnitude and variability of ICD distributions. In particular, post-mitigation Rashomon sets exhibit substantially lower median and maximum ICD values, along with narrower distributions across transparency bins, indicating more stable and equitable interpretability alloca-




Figure 6: Distribution of several desiderata across Rashomon sets over transparency bins ($\varepsilon = 0.01$) for HybridCORELSPre and HybridCORELSPost, with or without our proposed ICD mitigation ($\eta = 0.05$). Each subplot corresponds to a mitigation applied to a sensitive attribute (Age, Gender, Race), ordered from left to right. Results are reported for the COMPAS dataset.

tion across demographic groups. These results demonstrate that enforcing ICD constraints during training generalizes effectively to unseen test data.

Result 4. ICD mitigation has no clear impact on assignment arbitrariness. Figure 5b displays the distribution of test set ICA within Rashomon sets of different transparency bins, on the COMPAS dataset, with or without our proposed ICD mitigation. Overall, the ICD mitigation procedure produces no substantial change in the distribution of assignment arbitrariness across datasets and transparency bins: with or without ICD mitigation, the ICA values remain largely similar. This indicates that although ICD mitigation is highly effective at reducing group-level interpretability coverage disparities, it has limited impact on individual-level assignment arbitrariness across Rashomon sets.

Result 5. ICD mitigation frequently improves algorithmic fairness metrics. Figure 6a displays the distribution of test set EO within Rashomon sets of different transparency bins, on the COMPAS dataset, with or without our proposed ICD mitigation. Analogous results for the SP metric are provided in Appendix D. The experiments show frequent improvements in both EO and SP after mitigation, de-

spite these metrics not being directly targeted during training. In many settings, mitigated models exhibit lower median disparities and reduced variability across Rashomon sets, particularly noticeable decreases in both the spread and the upper range of EO values in higher transparency bins. However, the improvement is not uniform across all datasets and bins, and certain settings show relatively limited or inconsistent changes after mitigation. These observations suggest that reducing ICD can often indirectly improve algorithmic fairness metrics, further illustrating how improving procedural fairness might benefit outcome-based fairness.

Result 6. ICD mitigation preserves predictive performance and sparsity. Figure 6b (respectively, 6c) reports the distribution of the number of rules in the learned prefix (respectively, of test accuracy) within Rashomon sets of different transparency bins, on the COMPAS dataset, with or without our proposed ICD mitigation. Corresponding results for the other datasets are provided in Appendix D. Since the number of rules quantifies the *sparsity* of the interpretable component, this analysis evaluates whether ICD mitigation increases model complexity. The results show that enforcing ICD constraints leads to only minor increases in the number of rules, with sparsity remaining unchanged in many




Figure 7: Test set ICD, accuracy and EO, and model sparsity for HybridCORELSPre and HybridCORELSPost with ICD mitigation across Rashomon sets over all transparency values in $[0, 1]$ ($\epsilon = 0.01$). For different levels of mitigation induced by the maximum ICD constraint (η), curves report the average and standard deviation across Rashomon models. Results are provided for the COMPAS dataset and Gender as sensitive attribute.

settings. Enforcing ICD constraints also has only a limited impact on predictive performance. Although some settings exhibit slight decreases in test accuracy after mitigation, the overall accuracy distributions remain highly comparable to those without mitigation across all transparency bins. In some configurations, enforcing ICD constraints can improve test accuracy, suggesting a slight regularization effect.

Result 7. Varying the enforced ICD constraint confirms the previous observations. To further analyze the effect of ICD mitigation strength, Figure 7 illustrates how test set ICD, accuracy, EO, as well as sparsity evolve as the maximum ICD constraint is gradually relaxed from a highly restrictive setting ($\eta = 0.01$) to the unconstrained case ($\eta = 1$), corresponding to no mitigation. This analysis is conducted on the COMPAS dataset with mitigation applied to the Gender attribute. We consider all models within the approximate Rashomon set satisfying the desired performance threshold ($\epsilon = 0.01$) without partitioning them into transparency bins. Appendix D provides bin-level analyses.

In terms of predictive performance, test accuracy remains remarkably stable across the full range of mitigation strengths, with no substantial degradation even under the strictest maximum ICD constraints. As expected, relaxing the maximum ICD constraint consistently increases both the average and the variability of ICD across Rashomon models. Combined with relatively unchanged predictive performance across mitigation levels, these results suggest that models with substantially lower ICD can be selected from the Rashomon set without sacrificing predictive accuracy.

For algorithmic fairness metrics, measured through EO and SP, the aggregated effect across all Rashomon models appears relatively modest. However, the bin-level analyses in Appendix D (Figures 24 and 25) show that the largest improvements occur in higher-transparency bins, where interpretability coverage is greatest and the mitigation constraint most strongly affects routing. For instance, in the highest-transparency bin, tightening η from 0.25 to 0.05 reduces the average EO of HybridCORELSPost from about 0.25 to about 0.18, and its average SP from about 0.22 to about 0.16.

Finally, the sparsity analysis shows only a slight increase in the number of rules under the strictest mitigation settings, consistent with previous observations. Overall, these findings provide evidence that ICD can often be substan-

tially reduced with limited changes in predictive performance and sparsity, while frequently yielding improvements in outcome-based fairness metrics.

7 Conclusion

In this work, we formalized ICD and ICA as procedural fairness concerns arising at the group and individual levels in hybrid interpretable models, where some individuals receive decisions from an interpretable component while others are deferred to a black box. Our empirical results show that state-of-the-art hybrid interpretable methods can exhibit substantial ICD across protected groups when applied to real-world datasets. We also showed that predictive multiplicity provides a practical path to mitigation: within sets of similarly accurate models, one can often identify models with substantially lower ICD, with little impact on predictive performance or sparsity and, in some cases, indirect improvements in standard algorithmic fairness metrics. Further work will assess how consistently these trends extend across datasets, and evaluate, through user studies, whether greater interpretability coverage for unprivileged groups systematically translates into improved human understanding.

Future work also includes investigating which characteristics make an example more or less likely to be affected by ICA, as well as designing mechanisms to mitigate such arbitrariness. Because ICA is defined with respect to a set of models, it cannot be directly evaluated during the training of a single model, making its mitigation a challenging open problem. Mechanisms inspired by the stability literature (Hara and Yoshida 2023) may help address it.

Understanding the upstream causes of ICD is another important research direction, as it may arise from complex, dataset-dependent correlations among features, ground-truth labels, and sensitive attributes. Investigating its downstream effects on other dimensions of trustworthy machine learning is also important. For instance, ICD might cause disparate privacy risks, since the examples deferred to the black-box are vulnerable to inference attacks (Rigaki and Garcia 2023) against both the interpretable part they went through, and the black-box component ultimately handling them. ICD may also interact with robustness, as examples deferred to the black-box could be harder to audit for adversarial vulnerability (Chaubey et al. 2020), especially when robustness must be assessed locally within a decision region.

Acknowledgments

This research was enabled by support provided by Calcul Québec and the Digital Research Alliance of Canada, as well as funding from the SCALE-AI Chair in Data-Driven Supply Chains and the *Fonds de recherche du Québec – Nature et technologies (FRQNT)* through a Team Research Project (327090). Ziba Jabbar Zare is also supported by the *Fonds de recherche du Québec – Nature et technologies (FRQNT)* through a B2X Doctoral scholarship (350463). The authors would like to thank the anonymous reviewers for their valuable suggestions.

References

1978. Uniform Guidelines on Employee Selection Procedures. 29 CFR Part 1607. Equal Employment Opportunity Commission.
- Aïvodji, U.; Arai, H.; Fortineau, O.; Gambs, S.; Hara, S.; and Tapp, A. 2019. Fairwashing: The Risk of Rationalization. In *Proceedings of the 36th International Conference on Machine Learning, ICML*, 161–170.
- Angwin, J.; Larson, J.; Mattu, S.; and Kirchner, L. 2016. Machine bias: There’s software used across the country to predict future criminals. And it’s biased against blacks. ProPublica (2016). *ProPublica*, May, 23.
- Arslan, E.; van der Linden, J.; Hoogendoorn, S.; Rinaldi, M.; and Demirović, E. 2025. SORTeD Rashomon Sets of Sparse Decision Trees: Anytime Enumeration. In *Proceedings of the 39th Annual Conference on Neural Information Processing Systems, NeurIPS*, 44717–44755.
- Artelt, A.; and Hammer, B. 2022. “Explain it in the Same Way!” - Model-Agnostic Group Fairness of Counterfactual Explanations. arXiv:2211.14858.
- Balagopalan, A.; Zhang, H.; Hamidieh, K.; Hartvigsen, T.; Rudzicz, F.; and Ghassemi, M. 2022. The Road to Explainability Is Paved with Bias: Measuring the Fairness of Explanations. In *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency, FAccT*, 1194–1206.
- Barocas, S.; and Selbst, A. D. 2016. Big Data’s Disparate Impact. *California Law Review*, 104: 671–732.
- Bell, A.; Fonseca, J.; Abrate, C.; Bonchi, F.; and Stoyanovich, J. 2024. Fairness in algorithmic recourse through the lens of substantive equality of opportunity. arXiv:2401.16088.
- Breiman, L. 2001. Statistical modeling: The two cultures (with comments and a rejoinder by the author). *Statistical science*, 16(3): 199–231.
- Breiman, L.; Friedman, J. H.; Olshen, R. A.; and Stone, C. J. 1984. *Classification and Regression Trees*. Wadsworth.
- Chaubey, A.; Agrawal, N.; Barnwal, K.; Guliani, K. K.; and Mehta, P. 2020. Universal Adversarial Perturbations: A Survey. arXiv:2005.08087.
- Ciaperoni, M.; Xiao, H.; and Gionis, A. 2024. Efficient Exploration of the Rashomon Set of Rule-Set Models. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD*, 478–489.
- Coker, B.; Rudin, C.; and King, G. 2021. A theory of statistical inference for ensuring the robustness of scientific results. *Management Science*, 67(10): 6174–6197.
- Cooper, A. F.; Lee, K.; Choksi, M. Z.; Barocas, S.; Sa, C. D.; Grimmelmann, J.; Kleinberg, J. M.; Sen, S.; and Zhang, B. 2024. Arbitrariness and Social Prediction: The Confounding Role of Variance in Fair Classification. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence, AAAI*, 22004–22012.
- Coston, A.; Rambachan, A.; and Chouldechova, A. 2021. Characterizing fairness over the set of good models under selective labels. In *Proceedings of the 38th International Conference on Machine Learning, ICML*, 2144–2155.
- Dai, J.; Upadhyay, S.; Aïvodji, U.; Bach, S. H.; and Lakkaraju, H. 2022. Fairness via Explanation Quality: Evaluating Disparities in the Quality of Post Hoc Explanations. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society, AIES*, 203–214.
- Dhaini, M.; Erdogan, E.; Feldhus, N.; and Kasneci, G. 2025. Gender Bias in Explainability: Investigating Performance Disparity in Post-hoc Methods. In *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency, FAccT*, 3006–3029.
- Ding, F.; Hardt, M.; Miller, J.; and Schmidt, L. 2021. Retiring Adult: New Datasets for Fair Machine Learning. In *Proceedings of the 35th Annual Conference on Neural Information Processing Systems, NeurIPS*, 6478–6490.
- Dua, D.; and Graff, C. 2017. UCI Machine Learning Repository.
- Dwork, C.; Hardt, M.; Pitassi, T.; Reingold, O.; and Zemel, R. 2012. Fairness Through Awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference, ITCIS*, 214–226.
- Ezzeddine, F.; Ammar, O.; Giordano, S.; and Ayoub, O. 2026. Fair Recourse for All: Ensuring Individual and Group Fairness in Counterfactual Explanations. arXiv:2601.20449.
- Feldman, M.; Friedler, S. A.; Moeller, J.; Scheidegger, C.; and Venkatasubramanian, S. 2015. Certifying and Removing Disparate Impact. In *Proceedings of the 21st ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD*, 259–268.
- Ferry, J.; Laberge, G.; and Aïvodji, U. 2024. Learning Hybrid Interpretable Models: Theory, Taxonomy, and Methods. *Transactions on Machine Learning Research*.
- Fisher, A.; Rudin, C.; and Dominici, F. 2019. All Models are Wrong, but Many are Useful: Learning a Variable’s Importance by Studying an Entire Class of Prediction Models Simultaneously. *Journal of Machine Learning Research*, 20(177): 1–81.
- Guidotti, R.; Monreale, A.; Ruggieri, S.; Turini, F.; Pedreschi, D.; and Giannotti, F. 2019. A Survey Of Methods For Explaining Black Box Models. *ACM Computing Surveys*, 51(5): 93:1–93:42.
- Gupta, V.; Nokhiz, P.; Roy, C. D.; and Venkatasubramanian, S. 2019. Equalizing Recourse across Groups. arXiv:1909.03166.

- Han, J.; Pei, J.; and Yin, Y. 2000. Mining Frequent Patterns without Candidate Generation. *SIGMOD Rec.*, 29(2): 1–12.
- Hara, S.; and Yoshida, Y. 2023. Average sensitivity of decision tree learning. In *The Eleventh International Conference on Learning Representations, ICLR*.
- Hardt, M.; Price, E.; and Srebro, N. 2016. Equality of Opportunity in Supervised Learning. In *Proceedings of the 30th Annual Conference on Neural Information Processing Systems, NIPS*, 3315–3323.
- Holm, S. 1979. A Simple Sequentially Rejective Multiple Test Procedure. *Scandinavian Journal of Statistics*, 6(2): 65–70.
- Kamiran, F.; and Calders, T. 2009. Classifying without discriminating. In *Proceedings of the 2nd International Conference on Computer, Control and Communication*, 1–6.
- Kavouras, L.; Tsopelas, K.; Giannopoulos, G.; Sacharidis, D.; Psaroudaki, E.; Theologitis, N.; Rontogiannis, D.; Fotakis, D.; and Emiris, I. 2023. Fairness Aware Counterfactuals for Subgroups. In *Proceedings of the 37th Annual Conference on Neural Information Processing Systems, NeurIPS*, 58246–58276.
- Langlade, L.; Ferry, J.; Laberge, G.; and Vidal, T. 2025. Fairness and Sparsity within Rashomon Sets: Enumeration-Free Exploration and Characterization. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society, AIES*, 1536–1547.
- Mann, H. B.; and Whitney, D. R. 1947. On a Test of Whether One of Two Random Variables Is Stochastically Larger than the Other. *The Annals of Mathematical Statistics*, 18(1): 50–60.
- Mata, K.; Kanamori, K.; and Arimura, H. 2022. Computing the Collection of Good Models for Rule Lists. In *Proceedings of the 18th International Conference on Machine Learning and Data Mining, MLDM*.
- Mehrabani, N.; Morstatter, F.; Saxena, N.; Lerman, K.; and Galstyan, A. 2022. A Survey on Bias and Fairness in Machine Learning. *ACM Computing Surveys*, 54(6): 115:1–115:35.
- Pan, D.; Wang, T.; and Hara, S. 2020. Interpretable Companions for Black-Box Models. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics, AISTATS*, 2444–2454.
- Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; Vanderplas, J.; Passos, A.; Cournapeau, D.; Brucher, M.; Perrot, M.; and Duchesnay, E. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12: 2825–2830.
- Rigaki, M.; and Garcia, S. 2023. A Survey of Privacy Attacks in Machine Learning. *ACM Computing Surveys*, 56(4).
- Rijnbeek, P. R.; and Kors, J. A. 2010. Finding a Short and Accurate Decision Rule in Disjunctive Normal Form by Exhaustive Search. *Machine Learning*, 80(1): 33–62.
- Rivest, R. L. 1987. Learning Decision Lists. *Machine Learning*, 2(3): 229–246.
- Slack, D.; Hilgard, S.; Lakkaraju, H.; and Singh, S. 2021. Counterfactual Explanations Can Be Manipulated. In *Proceedings of the 35th Annual Conference on Neural Information Processing Systems, NeurIPS*, 62–75.
- Ustun, B.; and Rudin, C. 2016. Supersparse Linear Integer Models for Optimized Medical Scoring Systems. *Machine Learning*, 102(3): 349–391.
- Wang, T.; and Lin, Q. 2021. Hybrid Predictive Models: When an Interpretable Model Collaborates with a Black-box Model. *Journal of Machine Learning Research*, 22(137): 1–38.
- Watson-Daniels, J.; Parkes, D. C.; and Ustun, B. 2023. Predictive Multiplicity in Probabilistic Classification. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence, AAAI*, 10306–10314.
- Xin, R.; Zhong, C.; Chen, Z.; Takagi, T.; Seltzer, M.; and Rudin, C. 2022. Exploring the whole rashomon set of sparse decision trees. In *Proceedings of the 36th Annual Conference on Neural Information Processing Systems, NeurIPS*, 14071–14084.
- Zemel, R.; Wu, Y.; Swersky, K.; Pitassi, T.; and Dwork, C. 2013. Learning Fair Representations. In *Proceedings of the 30th International Conference on Machine Learning, ICML*, 325–333.
- Zhang, B. H.; Lemoine, B.; and Mitchell, M. 2018. Mitigating unwanted biases with adversarial learning. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society, AIES*, 335–340.
- Zhao, Y.; Wang, Y.; and Derr, T. 2022. Fairness and Explainability: Bridging the Gap Towards Fair Model Explanations. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence, AAAI*, 11363–11371.
- Zhong, C.; Chen, Z.; Liu, J.; Seltzer, M.; and Rudin, C. 2024. Exploring and interacting with the set of good sparse generalized additive models. *Proceedings of the 38th Annual Conference on Neural Information Processing Systems, NeurIPS*, 56673–56699.