



HAL
open science

Learning Optimal Fair Scoring Systems for Multi-Class Classification

Julien Rouzot, Julien Ferry, Marie-José Huguet

► **To cite this version:**

Julien Rouzot, Julien Ferry, Marie-José Huguet. Learning Optimal Fair Scoring Systems for Multi-Class Classification. ICTAI 2022 - The 34th IEEE International Conference on Tools with Artificial Intelligence, Oct 2022, Virtual, United States. hal-03782735

HAL Id: hal-03782735

<https://hal.science/hal-03782735v1>

Submitted on 21 Sep 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Learning Optimal Fair Scoring Systems for Multi-Class Classification

Julien Rouzot
LAAS-CNRS,
Université de Toulouse, CNRS
Toulouse, France
jrouzot@laas.fr

Julien Ferry
LAAS-CNRS,
Université de Toulouse, CNRS
Toulouse, France
jferry@laas.fr

Marie-José Huguet
LAAS-CNRS, Université de Toulouse,
CNRS, INSA
Toulouse, France
huguet@laas.fr

Abstract—Machine Learning models are increasingly used for decision making, in particular in high-stakes applications such as credit scoring, medicine or recidivism prediction. However, there are growing concerns about these models with respect to their lack of interpretability and the undesirable biases they can generate or reproduce. While the concepts of interpretability and fairness have been extensively studied by the scientific community in recent years, few works have tackled the general multi-class classification problem under fairness constraints, and none of them proposes to generate fair and interpretable models for multi-class classification. In this paper, we use Mixed-Integer Linear Programming (MILP) techniques to produce inherently interpretable scoring systems under sparsity and fairness constraints, for the general multi-class classification setup. Our work generalizes the SLIM (*Supersparse Linear Integer Models*) framework that was proposed by Rudin and Ustun to learn optimal scoring systems for binary classification. The use of MILP techniques allows for an easy integration of diverse operational constraints (such as, but not restricted to, fairness or sparsity), but also for the building of certifiably optimal models (or sub-optimal models with bounded optimality gap).

Index Terms—Machine Learning, Fairness, Interpretability, Mixed-Integer Linear Programming, Scoring System, Multi-class Classification

I. INTRODUCTION

Machine Learning (ML) models are increasingly used for decision making, in particular in high-stakes applications such as credit scoring [1], medicine [2] or recidivism prediction [3]. In the context of supervised learning, models such as Deep Neural Networks are used commonly for their capacity to accurately analyse and capture complex correlations. However, such tools have been criticized in the literature for their lack of interpretability and the undesirable biases they can reproduce or introduce [4], [5].

Despite their capabilities to handle large amounts of data, most *black-box* (i.e., non-interpretable) models should be avoided in high-stakes decision settings [5]. In a nutshell, the interpretability of a ML model can be defined as “the ability to explain or to present in understandable terms to a human” [6]. This definition is rather abstract, and the quality of any explanation mechanism depends on the task at hand, the context and on the type of user receiving the explanation.

From a technical perspective, two main approaches have emerged in the literature to enhance the comprehensibility of

ML models [7]. On the one hand, *post-hoc explanations* [8] can be used *a posteriori* to explain (globally or locally) the decisions of a black-box model. While such explanations can be useful in non-sensitive settings, they usually come without trustworthiness guarantees, in particular with respect to how they really reflect the underlying reasoning of the model, and may be manipulated [5], [9]. On the other hand, *transparent-box design* methods build models that are interpretable by nature, such as rule-based models, tree-based models or scoring systems [10]. To ensure interpretability, *sparsity* is often encouraged to keep the models’ sizes reasonable. In addition, because of their apparent simplicity (which is necessary to keep them understandable), *interpretable-by-design* models are often criticized in the literature for their presumed lack of performances (e.g., in terms of accuracy). However, it has been demonstrated that they often perform as well as black-box models for a wide range of real-world applications [5].

The problem of algorithmic bias, in which a trained ML model uses irrelevant or discriminatory features for decision making, has been identified and widely studied [11]. To tackle it, different families of fairness notions have been proposed, namely individual fairness, causal fairness and statistical fairness [12]. In this paper, we consider statistical fairness notions, which have been extensively studied in the last decade.

Many works have been proposed in recent years, attempting to define interpretability notions [6], [7], fairness metrics [11], [12] and developing new frameworks that address these important challenges [13]. However, most of these works have focused on binary classification and few solutions have been proposed to handle the more general multi-class setup [14], [15]. Furthermore, to the best of our knowledge, none of the works addressing fairness in multi-class classification also considered interpretability. In this paper, we address this issue by proposing a framework to generate fair and interpretable models for multi-class classification. Our approach produces certifiably optimal models, which is crucial as lack of optimality can have societal implications [16].

More precisely, our contributions are as follows:

- 1) We review and summarize in a unified notation common multi-class fairness metrics proposed in the literature, and build on them to formulate our own metrics, introducing the flexible notion of *sensitive labels*.

- 2) We design and implement a new framework based on Mixed Integer Linear Programming to generate optimal sparse scoring systems for multi-class classification.
- 3) We integrate fairness constraints in this framework to generate optimal sparse and fair scoring systems for multi-class classification. The resulting method, named FAIRScoringSystems, is available online¹.
- 4) We empirically evaluate the effectiveness of the proposed approach to learn interpretable sparse models achieving good trade-offs between accuracy and fairness in multi-class classification problems.

This paper is organized as follows. First, in Section II, we provide the necessary background on multi-class classification and fairness in ML. Then, in Section III, we present the SLIM framework [3], [17] that generates optimal scoring systems for binary classification and provide a summary of the multi-class fairness metrics proposed in the literature in a unified notation. Afterwards, in Section IV, we detail our contributions for fair scoring systems for multi-class classification. Finally, we present and analyse our experimental results in Section V before concluding in Section VI.

II. BACKGROUND

In this section, we first introduce the notations that are used throughout the paper, and describe the ML classification task. Then, we present the relevant performance metrics. Finally, we define the notion of statistical fairness for binary classification.

A. Classification

We use the following notations throughout this paper. Let \mathcal{E} be a dataset containing N samples: $\mathcal{E} = \{e_i, i \in \{1 \dots N\}\}$. Each sample $e_i \in \mathcal{E}$ is defined by the values of D binary features $(x_{i,1}, \dots, x_{i,D})$ and one label $y_i \in \mathcal{K}$, with $|\mathcal{K}| = K$. If $y_i = k$, we say that e_i belongs to class k . Finally, let \hat{y}_i denotes the label predicted by a trained ML model for sample e_i .

In ML, a classification task refers to a predictive problem in which, given some input features $(x_{i,1}, \dots, x_{i,D})$, a *model* aims to predict the label associated to this input. Supervised learning methods process a labelled dataset and exploit the correlations learnt from the data to produce a model. Given a sufficient amount of training data, the objective is to accurately predict (through \hat{y}_i) the label y_i of a new sample [18]. A common classification problem is binary classification, where $K = 2$ (e.g., is an email “spam” or “not spam”). In this paper, we will focus on the more general multi-class classification, in which the output of the classification algorithm can take any value in a set \mathcal{K} of possible classes.

Two main approaches can be used to aggregate binary classification models to tackle multi-class classification tasks [19]. The *one-vs-all* strategy aims to fit one binary classifier per class, distinguishing the class’ examples from those of the other classes. The *all-vs-all* strategy (also called *one-vs-one*) constructs one classifier for each pair of classes. The later often leads to a lack of interpretability as the *all-vs-all* strategy

produces much more classifiers, so we consider the *one-vs-all* approach in our framework.

B. Performance Metrics

To evaluate the performances of a multi-class classification model, its *confusion matrix* is commonly used. In this $K \times K$ matrix, each row corresponds to an actual label and each column represents a predicted label. Each cell (k_a, k_b) contains the number of samples belonging to the class k_a in the dataset and predicted in the class k_b by the model. Hence, the diagonal of this matrix represents the samples that have been correctly classified. Based on these values, a popular metric for assessing a model’s performance is its predictive *accuracy* [24], which evaluates the proportion of correctly classified samples:

$$\text{Accuracy} = \frac{\sum_{k=1}^K TP_k}{N} \quad (1)$$

in which TP_k represents the number of *True Positive* examples for a given class k : $TP_k = |\{e_i \mid y_i = k \wedge \hat{y}_i = k\}|$.

However, when dealing with imbalanced datasets, the *balanced accuracy* metric is often preferred. In a nutshell, it computes the model’s average accuracy over the different classes:

$$\text{Balanced Accuracy} = \frac{\sum_{k=1}^K \frac{TP_k}{N_k}}{K} \quad (2)$$

in which N_k is the number of samples belonging to class k : $N_k = |\{e_i \mid y_i = k\}|$.

C. Fairness

Statistical fairness (also called group fairness) notions aim to correct a prediction bias between different subsets of a dataset called *protected groups*. These protected groups usually differ by the value of one or several protected feature(s) (e.g., age, gender or ethnicity) [25]. In this paper, we let $\mathcal{A} \subset \mathcal{E}$ be a protected group, and the purpose of fairness notions is to ensure that the learnt classifier behaves similarly between individuals from protected group \mathcal{A} and from the rest of the population $(\mathcal{E} \setminus \mathcal{A})$. To realize this, fairness metrics are used to measure the unfairness of a classifier based on a certain statistical criterion. While different fairness metrics have been introduced, some of them are incompatible [12], meaning that they cannot be optimized altogether.

We first introduce widely used fairness metrics in the binary classification setup ($\mathcal{K} = \{0, 1\}$) in Table I. In a nutshell, statistical parity ensures that individuals from the two groups have the same probabilities to be positively predicted. Overall misclassification rate equalizes the probabilities of being incorrectly classified. Predictive equality equalizes the False Positive Rates (FPR), while equal opportunity equalizes the True Positive Rates (TPR). Finally, equalized odds can be interpreted as the conditional independence between the prediction of the classifier and the protected feature given the ground truth (and is, in this setup, the conjunction of predictive equality and equal opportunity).

¹<https://gitlab.laas.fr/roc/julien-rouzot/fairscoringsystems0>

TABLE I: Standard fairness metrics for binary classification.

Metric	Mathematical expression
Statistical Parity [20]	$P(\hat{y}_i = 1 e_i \in \mathcal{A}) = P(\hat{y}_i = 1 e_i \notin \mathcal{A})$
Overall Misclassification Rate [21]	$P(\hat{y}_i \neq y_i e_i \in \mathcal{A}) = P(\hat{y}_i \neq y_i e_i \notin \mathcal{A})$
Predictive Equality [22]	$P(\hat{y}_i = 1 e_i \in \mathcal{A}, y_i = 0) = P(\hat{y}_i = 1 e_i \notin \mathcal{A}, y_i = 0)$
Equal Opportunity [23]	$P(\hat{y}_i = 1 e_i \in \mathcal{A}, y_i = 1) = P(\hat{y}_i = 1 e_i \notin \mathcal{A}, y_i = 1)$
Equalized Odds [23]	$P(\hat{y}_i = 1 e_i \in \mathcal{A}, y_i) = P(\hat{y}_i = 1 e_i \notin \mathcal{A}, y_i)$

As exact fairness is often too restrictive, a common relaxation consists in quantifying the difference of some statistical measure among the different protected groups. The model at hand is then coined as *fair* if its measured fairness violation (unfairness) is lower than an *unfairness tolerance* ϵ . In practice, probabilities are estimated using empirical rates.

III. RELATED WORK

In this section, we first introduce the SLIM framework [3], [17] that was proposed to learn optimal scoring systems for binary classification. Then, we review existing works on fairness in multi-class classification before summarizing the existing metrics in a unified notation.

A. Learning Optimal Scoring Systems for Binary Classification with SLIM

Rudin and Ustun [3], [17] have developed the *Supersparse Linear Integer Model* (SLIM) framework, which produces optimal (in terms of accuracy and sparsity) scoring systems for binary classification. Scoring systems are considered interpretable-by-design models and are deployed in real-world applications such as medicine [17] or crime prediction [26].

As illustrated in Figure 1, a scoring system can be represented as a table in which each row associates a Boolean condition over the datasets’ features to a number of points (which can be negative). If a new sample satisfies the given condition, the associated number of points is added to the score of the sample. If the final score is above a given threshold, the sample is predicted to belong to the positive class, otherwise the model predicts the negative one.

SLIM is based on a Mixed Integer Linear Program (MILP), which can be solved using any off-the-shelf solver. While fairness requirements are not originally part of the SLIM framework, it is possible to incorporate (linearized) fairness constraints (as well as other operational constraints [17]), thanks to the declarative nature of the method.

PREDICT MUSHROOM IS POISONOUS IF SCORE > 3			
1.	<i>spore_print_color = green</i>	4 points
2.	<i>stalk_sur_face.above_ring = grooves</i>	2 points	+
3.	<i>population = clustered</i>	2 points	+
4.	<i>gill_size = broad</i>	-2 points	+
5.	<i>odor ∈ {none, almond, anise}</i>	-4 points	+
ADD POINTS FROM ROWS 1-5		SCORE	=

Fig. 1: Example scoring system for mushroom toxicity detection generated by SLIM [17].

B. Statistical Fairness in Multi-Class Classification

While most of the fairness literature focuses on the binary classification setup, some recent works have also studied fairness requirements for multi-class classification. Table II summarizes the proposed fairness metrics using a unified notation. While most of the fairness metrics for multi-class classification are straightforward extensions of the binary metrics defined in Table I, equalized odds can be interpreted in different ways in the multi-class setup. In particular, multi-class equalized odds (also called term-by-term equality of odds [27]) aims at equalizing the probabilities of being predicted $k \in \mathcal{K}$ given true label $k' \in \mathcal{K}$, for all pairs $(k, k') \in \mathcal{K}^2$, which requires K^2 fairness constraints. In contrast, classwise equality of odds [27] considers all false positive classifications altogether, which only requires $2 \cdot K$ constraints.

Similarly to the binary classification case, fairness-enhancing techniques for multi-class classification can be categorized into three different categories, depending on which stage of the ML pipeline they intervene [28]:

- *Pre-processing* methods modify the input dataset to remove undesired correlations with the sensitive feature(s) and/or to satisfy some statistical criteria [29], [30].
- *Post-processing* methods modify the output of the learning algorithm, such as a trained classifier, to satisfy some fairness criteria [27], [31], [32].
- Our framework belongs to the last category, which corresponds to *in-processing* (also called *algorithmic modification*) approaches. These methods tackle the unfair biases directly during the learning process [21], [31], [33], [34], which generally leads to the best fairness/utility tradeoffs [11].

To the best of our knowledge, fairness and interpretability have not been tackled altogether in any work for multi-class classification.

IV. LEARNING OPTIMAL FAIR AND SPARSE SCORING SYSTEMS FOR MULTI-CLASS CLASSIFICATION

In this section, we present our novel approach for learning optimal scoring systems for multi-class classification under fairness and sparsity constraints. First, we describe the considered fairness metrics for multi-class classification, introducing the novel concept of *sensitive labels*. Then, we extend binary scoring systems (as produced by SLIM) to the multi-class setting using the *one-vs-all* approach. Finally, we introduce FAIRScoringSystems, a flexible MILP framework to generate fair and interpretable models for multi-class classification.

TABLE II: Standard fairness metrics for multi-class classification.

Metric	Mathematical expression
Statistical Parity (or Demographic Parity) [30]–[32], [34]	$P(\hat{y}_i = k e_i \in \mathcal{A}) = P(\hat{y}_i = k e_i \notin \mathcal{A}) \quad \forall k \in \mathcal{K}$
Equal Opportunity (or Overall Accuracy) [31], [32]	$P(\hat{y}_i = k e_i \in \mathcal{A}, y_i = k) = P(\hat{y}_i = k e_i \notin \mathcal{A}, y_i = k) \quad \forall k \in \mathcal{K}$
Overall Misclassification Rate (or Max Loss Discrepancy) [21]	$P(\hat{y}_i \neq k e_i \in \mathcal{A}, y_i = k) = P(\hat{y}_i \neq k e_i \notin \mathcal{A}, y_i = k) \quad \forall k \in \mathcal{K}$
Equalized Odds [27], [29], [32]	$P(\hat{y}_i = k e_i \in \mathcal{A}, y_i) = P(\hat{y}_i = k e_i \notin \mathcal{A}, y_i) \quad \forall k \in \mathcal{K}$
Classwise Equality of Odds [27]	$P(\hat{y}_i = k y_i \neq k, e_i \in \mathcal{A}) = P(\hat{y}_i = k y_i \neq k, e_i \notin \mathcal{A})$ AND $P(\hat{y}_i = k y_i = k, e_i \in \mathcal{A}) = P(\hat{y}_i = k y_i = k, e_i \notin \mathcal{A}) \quad \forall k \in \mathcal{K}$

A. Proposed Fairness Metrics for Multi-Class Classification

Applying binary fairness metrics (as defined in Table I) on each label $k \in \mathcal{K}$ is not always necessary for all use cases and may sometimes lead to poor performances when unnecessary constraints are applied. We propose a more generic setup in which the set of labels \mathcal{K} is partitioned into a subset of *sensitive labels* $\mathcal{K}_s \subseteq \mathcal{K}$ and a subset of *unsensitive labels* ($\mathcal{K} \setminus \mathcal{K}_s$). Intuitively, sensitive labels correspond to outcomes yielding great impact on individuals’ lives for the decision-making process in which the scoring system is applied.

Afterwards, our proposed fairness metrics for multi-class classification only apply on the sensitive labels. More precisely, given a subset of sensitive labels \mathcal{K}_s and a protected group \mathcal{A} , we apply the binary metrics defined in Table I on each sensitive label $k \in \mathcal{K}_s$. The resulting metrics are defined in Table III. Note that some standard fairness metrics introduced in Section III-B correspond to the special case in which $\mathcal{K}_s = \mathcal{K}$. In particular, our multi-class Equalized Odds metric corresponds to the classwise equality of odds metric [27]. As Overall Misclassification Rate is equivalent to Equal Opportunity in our setup (balancing the True Positive Rate for class k is equivalent to balancing its False Negative Rate), we do not consider this metric for our framework.

B. Scoring Systems for Multi-Class Classification

We extend binary scoring systems (as originally produced by SLIM) to multi-class classification using the *one-vs-all* paradigm [19]. More precisely, one scoring system is generated for each label $k \in \mathcal{K}$ of the dataset, whose purpose is to detect examples belonging to class k . To classify a new sample e_i , each scoring system is applied and the class corresponding to the scoring system with the highest score is predicted.

Figure 2 provides an example of multi-class scoring system generated for the *customer* dataset. This dataset contains 4 different labels ($\mathcal{K} = \{A, B, C, D\}$) corresponding to customers’ categories and one scoring system is generated for each one of them. Given a new sample e_i , each scoring system is applied and the predicted label is the one whose scoring system yields the highest score.

C. FAIRScoringSystems

Hereafter, we describe FAIRScoringSystems, a Mixed Integer Linear Program (MILP) extending the SLIM framework [3], [17] (introduced in Section III-A) to the multi-class classification setup. More precisely, FAIRScoringSystems generates optimal multi-class scoring systems, maximizing accu-

Score for	A	Score for	B
starts with	9 pts	starts with	-9 pts
Graduated	9 pts	Ever_Married	-9 pts
Work experience <= 1	-9 pts	Graduated	9 pts
Cat == Cat_6	-9 pts	Work experience <= 1	6 pts

Score for	C	Score for	D
starts with	2 pts	starts with	1 pts
Graduated	9 pts	Spending_Score == Low	9 pts
Cat == Cat_6	9 pts	Work experience <= 1	-4 pts
		Cat == Cat_6	9 pts

Fig. 2: Example multi-class scoring system generated by FAIRScoringSystems for customer segmentation with sparsity (number of lines in each scoring system less or equal to 4) and fairness (equal opportunity violation less or equal to 0.01) constraints. This classifier exhibits an accuracy of 0.4495 and an equal opportunity violation of 0.0057 ($\mathcal{A} = \{e_i | x_{i,\text{sex}} = \text{female} = 1\}$, $\mathcal{K}_s = \{A\}$). “starts with” can be seen as a virtual attribute, which is True for all examples. It introduces a constant bias inherent to each individual scoring system.

racy (4) or balanced-accuracy (5), given fairness and sparsity constraints. Due to the declarative nature of the approach, such constraints can be tuned by the user and additional operational constraints can easily be handled. The proposed MILP formulation is detailed hereafter.

1) Variables:

- $\lambda_{k \in \{1..K\}, j \in \{1..D\}} \in \mathcal{L}$: matrix (K, D) of integer variables representing the value of each coefficient $j \in \{1..D\}$ for each scoring system $k \in \{1..K\}$. These *decision variables* fully describe our multi-class scoring system. Each coefficient $\lambda_{k,j}$ belongs to a user-defined set of integers \mathcal{L} , lower-bounded (respectively upper-bounded) by \mathcal{L}_{min} (respectively \mathcal{L}_{max}).
- $z_{i \in \{1..N\}} \in \{0, 1\}$: binary loss variables, each representing the misclassification of a training sample $e_{i \in \{1..N\}}$.
- $\alpha_{k \in \{1..K\}, j \in \{1..D\}} \in \{0, 1\}$: matrix (K, D) of binary variables indicating if each coefficient $\lambda_{k,j}$ is non-zero. As null coefficients do not appear in the final scoring system, the sparsity of the model directly depends on the number of non-zero coefficients.
- $pos_{i \in \{1..N\}, k \in \{1..K\}} \in \{0, 1\}$: matrix (N, K) of binary variables modelling the (one-hot encoded) prediction \hat{y}_i

TABLE III: Proposed fairness metrics for multi-class classification, using the notion of *sensitive labels*

Metric	Mathematical expression
Multi-class Statistical Parity (SP)	$P(\hat{y}_i = k e_i \in \mathcal{A}) = P(\hat{y}_i = k e_i \notin \mathcal{A}) \quad \forall k \in \mathcal{K}_s$
Multi-class Predictive Equality (PE)	$P(\hat{y}_i = k e_i \in \mathcal{A}, y_i \neq k) = P(\hat{y}_i = k e_i \notin \mathcal{A}, y_i \neq k) \quad \forall k \in \mathcal{K}_s$
Multi-class Equal Opportunity (EO)	$P(\hat{y}_i = k e_i \in \mathcal{A}, y_i = k) = P(\hat{y}_i = k e_i \notin \mathcal{A}, y_i = k) \quad \forall k \in \mathcal{K}_s$
Multi-class Equalized Odds (EOD)	Predictive Equality AND Equal Opportunity

of our classifier for each sample $e_{i \in \{1..N\}}$ of the training set, as this information is explicitly required to formulate fairness constraints.

2) *Objective*: The objective of our MILP, given in (3), is to minimize the misclassification error L along with a regularisation term γ_{reg} encouraging sparsity.

$$\min L + \gamma_{reg} \quad (3)$$

The misclassification error L can either be the empirical loss (4) or the balanced empirical loss (5). This is equivalent to maximizing the accuracy (1) or the balanced accuracy (2).

$$L = \frac{1}{N} \cdot \sum_{i=1}^N z_i \quad (4)$$

$$\text{or } L = \sum_{k=1}^K \left(\frac{1}{N_k} \sum_{i|y_i=k} z_i \right) \quad (5)$$

The regularisation term γ_{reg} , defined in (6), is weighted by a constant to ensure that the model will choose a sparser model (as the sum of the non-zero coefficients is minimized) only if it does not affect the loss. Thus, γ_{reg} must be strictly smaller than the additional loss implied by a single misclassification. As $\max \sum_{k=1}^K \sum_{j=1}^D \alpha_{k,j} = D \cdot K$ and as a single misclassification adds $\frac{1}{\min_{k \in \mathcal{K}}(N_k)}$ to the balanced empirical loss (with $\min_{k \in \mathcal{K}}(N_k)$ the number of samples in the most underrepresented class), we have (in the case of the empirical loss, $\min_{k \in \mathcal{K}}(N_k)$ is replaced by N):

$$\gamma_{reg} = \frac{1}{\min_{k \in \mathcal{K}}(N_k) \cdot D \cdot K + 1} \cdot \sum_{k=1}^K \sum_{j=1}^D \alpha_{k,j}. \quad (6)$$

3) *Constraints*:

Classification-related constraints:

$$-M_i z_i \leq s_{i,k} - s_{i,k'} \quad i \in 1..N \quad y_i = k \\ (k, k') \in (1..K)^2 \quad k \neq k' \quad (7)$$

$$s_{i,k} = \sum_{j=1}^D \lambda_{k,j} \cdot x_{i,j} - \gamma k \quad i \in 1..N \quad k \in 1..K \quad (8)$$

Sparsity-related constraints:

$$\mathcal{L}_{max} \cdot \alpha_{k,j} \geq \lambda_{k,j} \quad k \in 1..K \quad j \in 1..D \quad (9)$$

$$\mathcal{L}_{min} \cdot \alpha_{k,j} \leq -\lambda_{k,j} \quad k \in 1..K \quad j \in 1..D \quad (10)$$

$$\sum_{j=1}^D \alpha_{k,j} \leq \epsilon_s \quad k \in 1..K \quad (11)$$

Predictions-modeling constraints:

$$-M_i(1 - pos_{i,k}) \leq s_{i,k} - s_{i,k'} \\ i \in 1..N \quad (k, k') \in (1..K)^2 \quad k \neq k' \quad (12)$$

$$\sum_{k=1}^K pos_{i,k} = 1 \quad i \in 1..N \quad (13)$$

Statistical parity constraints:

$$\left| \frac{1}{|\mathcal{A}|} \sum_{e_i \in \mathcal{A}} pos_{i,k} - \frac{1}{|\mathcal{E} \setminus \mathcal{A}|} \sum_{e_i \notin \mathcal{A}} pos_{i,k} \right| \leq \epsilon_{sp} \quad k \in \mathcal{K}_s \quad (14)$$

Predictive Equality constraints:

$$\left| \frac{1}{|\mathcal{A}|} \sum_{e_i \in \mathcal{A}, y_i \neq k} pos_{i,k} - \frac{1}{|\mathcal{E} \setminus \mathcal{A}|} \sum_{e_i \notin \mathcal{A}, y_i \neq k} pos_{i,k} \right| \leq \epsilon_{pe} \quad k \in \mathcal{K}_s \quad (15)$$

Equal opportunity constraints:

$$\left| \frac{1}{|\mathcal{A}|} \sum_{e_i \in \mathcal{A}, y_i = k} pos_{i,k} - \frac{1}{|\mathcal{E} \setminus \mathcal{A}|} \sum_{e_i \notin \mathcal{A}, y_i = k} pos_{i,k} \right| \leq \epsilon_{eo} \quad k \in \mathcal{K}_s \quad (16)$$

Equalized odds constraints: Equalized odds is formulated as the conjunction of the predictive equality (15) and the equal opportunity (16) constraints, with $\epsilon_{pe} = \epsilon_{eod}$ et $\epsilon_{eo} = \epsilon_{eod}$.

Constraints (7) and (8) allow to model the variable z_i representing the classification errors. Constraint (8) keeps track of the score $s_{i,k}$ of each sample for each label's individual scoring system $k \in \{1..K\}$. A Big-M constraint (7) allows to set z_i to 1 if the associated sample is misclassified ($s_{i,k} < s_{i,k'} \quad y_i = k$). The value of M_i must be greater than $\max(s_{i,k} - s_{i,k'})$. A penalty γk is added to apply the lexicographical order in case the scores are the same for two different labels (arbitrary break of ties). As the ground truth $y_i = k$ is known, $\hat{y}_i = y_i$ (e_i is correctly classified) implies that $s_{i,k} > s_{i,k'} \quad \forall k' \neq k$.

To model the variables $\alpha_{k,j}$, we use constraints (9) and (10) that will set $\alpha_{k,j}$ to 1 if its associated coefficient $\lambda_{k,j} \neq 0$. Otherwise, $\alpha_{k,j}$ is set to 0 as $\mathcal{L}_{max} \geq \lambda_{k,j}$ and $\mathcal{L}_{min} \leq -\lambda_{k,j}$. The constraint (11) is applied to each label's scoring system to impose a sparsity constraint by limiting the maximal number of non-zero coefficients. The sparsity limit ϵ_s is a parameter.

The variable $pos_{i,k}$ represents the (one-hot encoded) decision \hat{y}_i of the multi-class scoring system for each sample $e_{i \in \{1..N\}}$ of the training set (required to express our fairness constraints). To model this variable, we set $pos_{i,k}$ to 0 if the

associated score is lower than at least one other score (12) and let $pos_{i,k}$ free otherwise. Constraint (13) ensures that the variable $pos_{i,k}$ associated with the highest score is set to 1.

The fairness constraints are then defined according to the user choice. In this paper, we focus on the most popular fairness metrics, but more operational constraints can be easily implemented. In line with state-of-the-art methods, our constraints bound the fairness violation (for protected group \mathcal{A}) by some unfairness tolerance ϵ_f set by the user, in which f represents the chosen fairness metric. In practice, each fairness constraint is implemented as a conjunction of two linear constraints (to get rid of the non-linear absolute value).

V. EXPERIMENTS

In this section, we empirically evaluate FAIRScoringSystems using one synthetic and two real-world datasets. First, we introduce the experimental setup before evaluating the resulting trade-offs between accuracy (or balanced accuracy) and sparsity. Finally, we assess the effectiveness of the method to produce interesting trade-offs between accuracy (or balanced accuracy) and fairness under different sparsity constraints.

A. Setup

We compare the performances of FAIRScoringSystems with two baseline methods: FAIR (fair baseline) and SVM (accurate baseline). More precisely, FAIR corresponds to a constant classifier predicting the majority label. It is the trivial model with highest accuracy, exhibiting perfect fairness (unfairness 0) for all our considered metrics. In contrast, SVM is a standard one-vs-all multi-class linear kernel SVM (without fairness constraints) using the scikit-learn² implementation. The value of the regularization hyper-parameter C is optimized using a standard grid-search technique with a validation set. While often considered as a black-box model [10], SVM offers an intuitive accuracy baseline because it consists in an aggregation of linear models, similarly to our proposed multi-class scoring systems. All method are evaluated using one synthetic and two real-world datasets:

- *Synthetic dataset* ($N = 800$, $D = 6$, $\mathcal{K} = \{L1, L2, L3\}$, $\mathcal{A} = \{e_i \mid x_{i,A1} = 1\}$, $\mathcal{K}_s = \{L1\}$): The features are generated with a random distribution between 40% and 60% of 1, and 3 labels L1, L2 and L3 are computed from the features (L1 depends on features 1-2-3, L2 on features 3-4-5, L3 on features 5-6-1) with a random noise³. By construction, this dataset is biased towards labels L1 and L3 for the sensitive feature A1.
- *wine*⁴ ($N = 6497$, $D = 25$, $\mathcal{K} = \{\text{good, medium, bad}\}$, $\mathcal{A} = \{e_i \mid x_{i,\text{colour}} = \text{red}\}$, $\mathcal{K}_s = \{\text{medium}\}$): This dataset associates chemical characteristics of wines to their quality (good, medium or bad). For the sake of illustration, the considered sensitive feature is the colour of the wine (red or white).

- *customer*⁵ ($N = 6665$, $D = 29$, $\mathcal{K} = \{A, B, C, D\}$, $\mathcal{A} = \{e_i \mid x_{i,\text{sex}} = \text{female} = 1\}$, $\mathcal{K}_s = \{A, B, C, D\}$): This dataset deals with customer segmentation, according to their profile. There are four anonymous customer categories (A, B, C and D), and the sensitive feature is gender (male or female).

Following standard procedures, both real-world datasets are binarized and potential inconsistencies are removed as preprocessing. As the wine dataset is highly imbalanced (77% of the samples are labelled “medium”), we use balanced accuracy as the performance metric for this dataset, while we use accuracy for the other datasets.

FAIRScoringSystems’ MILP is implemented and solved using the IBM ILOG CPLEX 20.1.0.0 solver⁶ via the *DOcplex*⁷ Python Modeling API and its default configuration. For all experiments, we set the integer coefficients’ range $\mathcal{L} = [-9, 9]$. Restricting coefficients to single-digit numbers enhances interpretability while still allowing sufficient expressivity.

The experiments are conducted on a computing grid over a set of homogeneous nodes using Intel Xeon E5-2695 v4 @ 2.10GHz CPU. Each experiment runs on a single CPU core for a fixed time limit of one hour and an allocated memory of 20 GB. Results are averaged using 5-folds cross-validation.

B. Accuracy/Sparsity Trade-offs

We first evaluate the accuracy of the produced multi-class scoring systems as a function of the sparsity constraint (*i.e.*, maximum number of lines for each individual scoring system). Figure 3 (respectively, 4) displays the performances on the training set (respectively, test set) for the three datasets with sparsity limits ranging from 1 to 10.

Figures 3 and 4 show that while very tight sparsity constraints can have a considerable impact on the models’ accuracy, the use of *reasonably* sparse interpretable models results in good performances. For the synthetic dataset, all experiments produced certifiably optimal scoring systems, while for the wine dataset, optimality gaps always lied below 0.1%.

C. Fairness/Accuracy Trade-offs

We then evaluate our framework on the multi-class fairness metrics introduced in section IV-A, for a wide range of unfairness tolerances (up to 20 values of ϵ ranging non linearly between 1% and the SVM unfairness, with a higher density near 1%), and different sparsity constraints (3, 5 and 10). Due to the limited space available, we only report results for the multi-class statistical parity metric. Nonetheless, the results obtained for the other metrics summarized in Table III show similar trends and are available on our repository.

Figure 5 (respectively, Figure 6) provides the Pareto frontiers between accuracy and fairness violation for the multi-class statistical parity metric on the training set (respectively, test set). The results show that FAIRScoringSystems is able

²<https://scikit-learn.org>

³<https://gitlab.laas.fr/roc/julien-rouzot/fairscoringssystemsv0>

⁴<https://www.kaggle.com/datasets/rajyellow46/wine-quality>

⁵<https://www.kaggle.com/datasets/vetirah/customer>

⁶<https://www.ibm.com/docs/en/icos/20.1.0>

⁷<http://ibmdecisionoptimization.github.io/docplex-doc/>

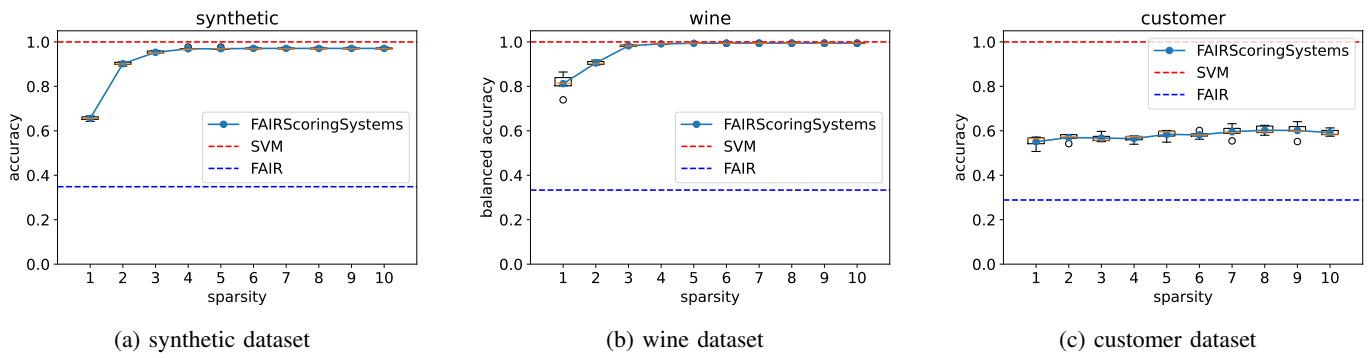


Fig. 3: Accuracy (train set) as a function of the sparsity constraint (maximum number of rules in each scoring system).

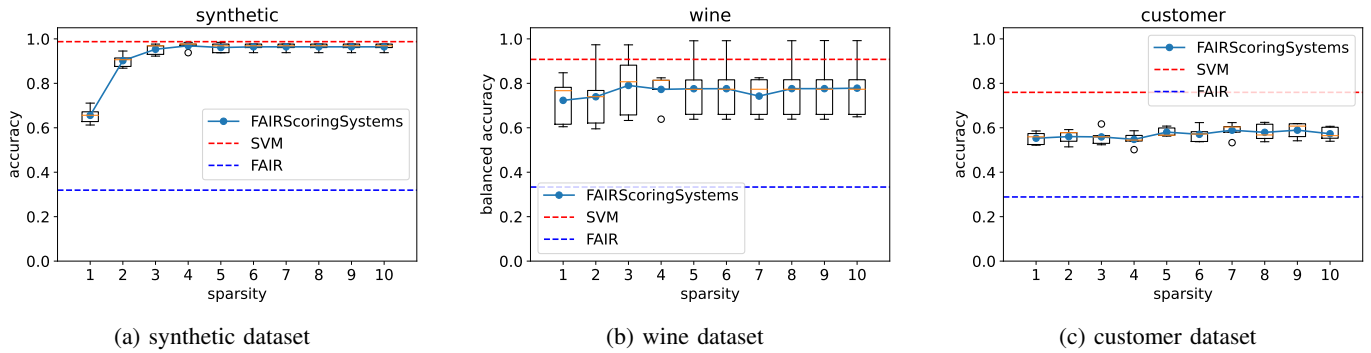


Fig. 4: Accuracy (test set) as a function of the sparsity constraint (maximum number of rules in each scoring system).

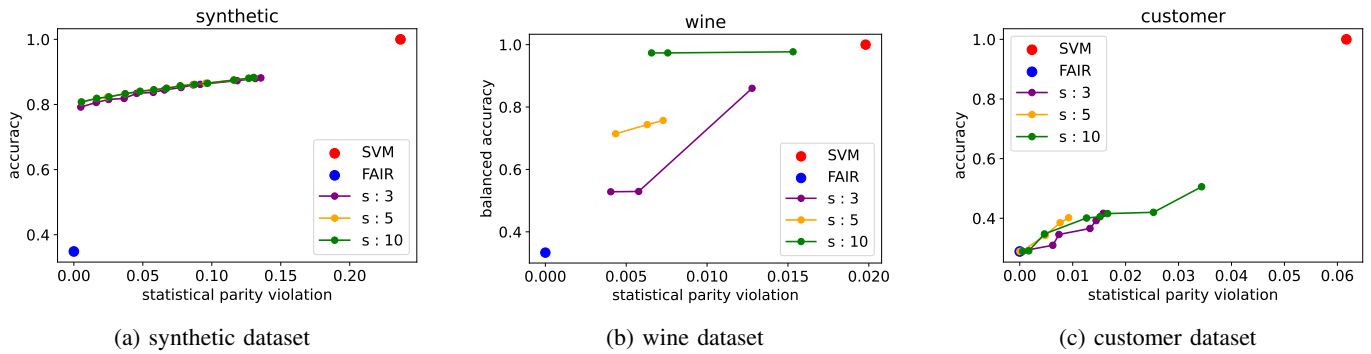


Fig. 5: Accuracy (train set) as a function of the unfairness value (statistical parity), for different sparsity constraints (3, 5, 10).

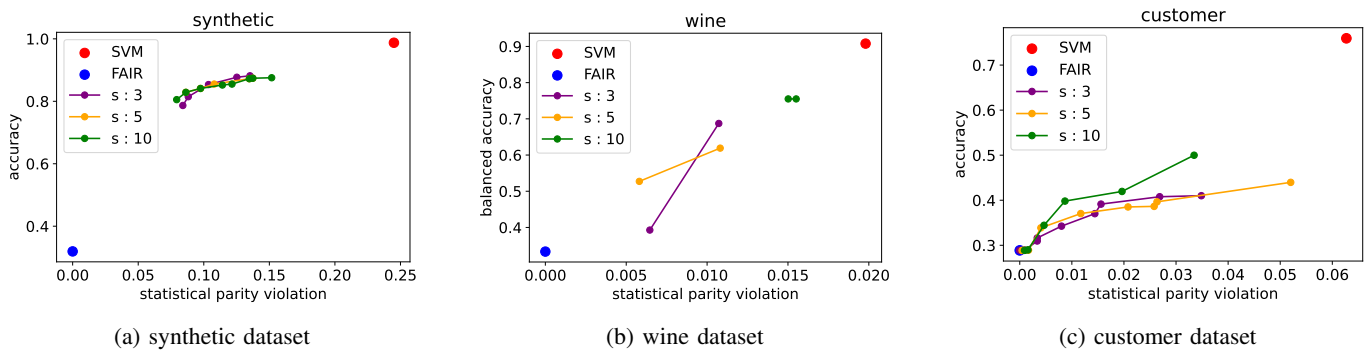


Fig. 6: Accuracy (test set) as a function of the unfairness value (statistical parity), for different sparsity constraints (3, 5, 10).

to produce interpretable models exhibiting interesting trade-offs between accuracy and fairness. For reasonable sparsity constraints, the generated models allow for substantial fairness violation reductions (compared to SVM) while still exhibiting good classification performances. Given the accuracy and fairness values, along with the interpretable models themselves, a domain-expert may then pick the most relevant model among the produced frontier.

VI. CONCLUSION

In this paper, we have proposed FAIRScoringSystems, a Mixed Integer Linear Programming framework producing optimal scoring systems, which are inherently interpretable models, under fairness and sparsity constraints for multi-class classification. To the best of our knowledge, this is the first work tackling these different requirements altogether. Our experimental evaluation demonstrates that FAIRScoringSystems is able to generate interesting trade-offs between accuracy, fairness and sparsity on both synthetic and real-world multi-class classification datasets of various shapes.

Future work includes improving our framework’s scalability, although it is already able to learn well-performing models for real-size datasets. While reaching and proving optimality for difficult datasets (*i.e.*, non-linearly separable, with high numbers of samples and features) is computationally challenging, our method can still be used to produce well-performing models. Using off-the-shelf solvers, our framework works in an *any-time* fashion: even if the learning process is stopped before proving optimality, a valid solution (*i.e.*, a multi-class scoring system satisfying sparsity and fairness constraints) can be returned. This implies that it can be used to learn fair and accurate interpretable models given limited time budgets. The produced models additionally come with a quality guarantee provided by optimality gaps, which upper-bound the distance to the actual optimal solution.

REFERENCES

- [1] X. Dastile, T. Celik, and M. Potsane, “Statistical and machine learning models in credit scoring: A systematic literature survey,” *Applied Soft Computing*, vol. 91, p. 106263, 2020.
- [2] B. J. Erickson, P. Korfiatis, Z. Akkus, and T. L. Kline, “Machine learning for medical imaging,” *Radiographics*, vol. 37, no. 2, p. 505, 2017.
- [3] C. Rudin and B. Ustun, “Optimized scoring systems: Toward trust in machine learning for healthcare and criminal justice,” *Interfaces*, vol. 48, no. 5, pp. 449–466, 2018.
- [4] J. Angwin, J. Larson, S. Mattu, and L. Kirchner, “Machine bias,” in *Ethics of Data and Analytics*. Auerbach Publications, 2016, pp. 254–264.
- [5] C. Rudin, “Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead,” *Nature Machine Intelligence*, vol. 1, no. 5, pp. 206–215, 2019.
- [6] F. Doshi-Velez and B. Kim, “Towards a rigorous science of interpretable machine learning,” *arXiv preprint arXiv:1702.08608*, 2017.
- [7] Z. C. Lipton, “The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery,” *Queue*, vol. 16, no. 3, pp. 31–57, 2018.
- [8] R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, and D. Pedreschi, “A survey of methods for explaining black box models,” *ACM computing surveys (CSUR)*, vol. 51, no. 5, pp. 1–42, 2018.
- [9] U. Aïvodji, H. Arai, O. Fortineau, S. Gambs, S. Hara, and A. Tapp, “Fairwashing: the risk of rationalization,” in *International Conference on Machine Learning*. PMLR, 2019, pp. 161–170.

- [10] A. A. Freitas, “Comprehensible classification models: a position paper,” *ACM SIGKDD explorations newsletter*, vol. 15, no. 1, pp. 1–10, 2014.
- [11] S. Barocas, M. Hardt, and A. Narayanan, *Fairness and Machine Learning*. fairmlbook.org, 2019, <http://www.fairmlbook.org>.
- [12] S. Verma and J. Rubin, “Fairness definitions explained,” in *2018 IEEE/ACM international workshop on software fairness (fairware)*. IEEE, 2018, pp. 1–7.
- [13] M. S. A. Lee and J. Singh, “The landscape and gaps in open source fairness toolkits,” in *Proceedings of the 2021 CHI conference on human factors in computing systems*, 2021, pp. 1–13.
- [14] S. Caton and C. Haas, “Fairness in machine learning: A survey,” *arXiv preprint arXiv:2010.04053*, 2020.
- [15] M. B. Zafar, I. Valera, M. G. Rogniguez, and K. P. Gummadi, “Fairness constraints: Mechanisms for fair classification,” in *Artificial intelligence and statistics*. PMLR, 2017, pp. 962–970.
- [16] E. Angelino, N. Larus-Stone, D. Alabi, M. I. Seltzer, and C. Rudin, “Learning certifiably optimal rule lists for categorical data,” *J. Mach. Learn. Res.*, vol. 18, pp. 234:1–234:78, 2017.
- [17] B. Ustun and C. Rudin, “Supersparse linear integer models for optimized medical scoring systems,” *Mach. Learn.*, vol. 102, no. 3, pp. 349–391, 2016.
- [18] R. Saravanan and P. Sujatha, “A state of art techniques on machine learning algorithms: a perspective of supervised learning approaches in data classification,” in *2018 Second International Conference on Intelligent Computing and Control Systems (ICICCS)*. IEEE, 2018, pp. 945–949.
- [19] M. Aly, “Survey on multiclass classification methods,” *Neural Netw.*, vol. 19, no. 1-9, p. 2, 2005.
- [20] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel, “Fairness through awareness,” in *Proceedings of the 3rd innovations in theoretical computer science conference*, 2012, pp. 214–226.
- [21] Q. Ye and W. Xie, “Unbiased subdata selection for fair classification: A unified framework and scalable algorithms,” *arXiv preprint arXiv:2012.12356*, 2020.
- [22] A. Chouldechova, “Fair prediction with disparate impact: A study of bias in recidivism prediction instruments,” *Big data*, vol. 5, no. 2, pp. 153–163, 2017.
- [23] M. Hardt, E. Price, and N. Srebro, “Equality of opportunity in supervised learning,” *Advances in neural information processing systems*, vol. 29, 2016.
- [24] M. Grandini, E. Bagli, and G. Visani, “Metrics for multi-class classification: an overview,” *arXiv preprint arXiv:2008.05756*, 2020.
- [25] F. Ding, M. Hardt, J. Miller, and L. Schmidt, “Retiring adult: New datasets for fair machine learning,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 6478–6490, 2021.
- [26] J. T. Andrade *et al.*, *Handbook of violence risk assessment and treatment: New approaches for mental health professionals*. Springer Publishing Company, 2009.
- [27] P. Putzel and S. Lee, “Blackbox post-processing for multiclass fairness,” *arXiv preprint arXiv:2201.04461*, 2022.
- [28] R. K. Bellamy, K. Dey, M. Hind, S. C. Hoffman, S. Houde, K. Kannan, P. Lohia, J. Martino, S. Mehta, A. Mojsilović *et al.*, “Ai fairness 360: An extensible toolkit for detecting and mitigating algorithmic bias,” *IBM Journal of Research and Development*, vol. 63, no. 4/5, pp. 4–1, 2019.
- [29] Y. Romano, S. Bates, and E. Candes, “Achieving equalized odds by resampling sensitive attributes,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 361–371, 2020.
- [30] I. Alabdulmohsin, J. Schrouff, and O. Koyejo, “A reduction to binary approach for debiasing multiclass datasets,” *arXiv preprint arXiv:2205.15860*, 2022.
- [31] F. Yang, M. Cisse, and S. Koyejo, “Fairness with overlapping groups: a probabilistic perspective,” *Advances in neural information processing systems*, vol. 33, pp. 4067–4078, 2020.
- [32] W. Alghamdi, H. Hsu, H. Jeong, H. Wang, P. W. Michalak, S. Asoodeh, and F. P. Calmon, “Beyond adult and compas: Fairness in multi-class prediction,” *arXiv preprint arXiv:2206.07801*, 2022.
- [33] S. Hossain, A. Mladenovic, and N. Shah, “Designing fairly fair classifiers via economic fairness notions,” in *Proceedings of The Web Conference 2020*, 2020, pp. 1559–1569.
- [34] C. Denis, R. Elie, M. Hebiri, and F. Hu, “Fairness guarantee in multi-class classification,” *arXiv preprint arXiv:2109.13642*, 2021.