
FACT: A Diagnostic for Group Fairness Trade-offs

Joon Sik Kim^{1,2} Jiahao Chen³ Ameet Talwalkar^{1,4}

Abstract

Group fairness, a class of fairness notions that measure how different groups of individuals are treated differently according to their protected attributes, has been shown to conflict with one another, often with a necessary cost in loss of model’s predictive performance. We propose a general diagnostic that enables systematic characterization of these trade-offs in group fairness. We observe that the majority of group fairness notions can be expressed via the fairness–confusion tensor, which is the confusion matrix split according to the protected attribute values. We frame several optimization problems that directly optimize both accuracy and fairness objectives over the elements of this tensor, which yield a general perspective for understanding multiple trade-offs including group fairness incompatibilities. It also suggests an alternate post-processing method for designing fair classifiers. On synthetic and real datasets, we demonstrate the use cases of our diagnostic, particularly on understanding the trade-off landscape between accuracy and fairness.

1. Introduction

As machine learning continues to be more widely used for applications with societal impact such as credit decisioning, predictive policing, and employment applicant screening, practitioners face regulatory, ethical, and legal challenges to prove whether or not their models are fair (Crawford et al., 2019). To provide quantitative tests of model fairness, the practitioners further need to choose between multiple definitions of fairness that exist in the machine learning literature (Calders et al., 2009; Žliobaitė, 2015; Narayanan, 2018). Among them is a class of definitions called *group*

fairness, which measures how a group of individuals with certain protected attributes are treated differently from other groups. This notion is widely studied as a concept of *disparate impact* in the legal context, and one specific instance of this notion was enforced as a law for fair employment process back in 1978 (Biddle, 2006). From a technical point of view however, several notions of group fairness have been shown to conflict with one another (Kleinberg et al., 2017; Chouldechova, 2017), sometimes with a necessary cost in loss of accuracy (Liu et al., 2019). Such considerations complicate the practical development and assessment of machine learning models designed to satisfy group fairness, as the conditions under which these trade-offs must necessarily occur can be too abstract to understand. Previous works on these trade-offs have been presented in ad hoc and definition-specific manner, which further calls for a more general perspective addressing the trade-offs in practice.

As an example, suppose an engineer is responsible for training a loan prediction model from a large user dataset, subject to mandatory group fairness requirements shaped by regulatory concerns. She has many choices for how to train this fair model, with fairness enforced before (Kamiran et al., 2010; Zemel et al., 2013; Madras et al., 2018; Samadi et al., 2018; Song et al., 2019; Tan et al., 2020), during (Zafar et al., 2015; 2017), or after (Dwork et al., 2012; Feldman et al., 2015; Hardt et al., 2016) training. However, she must resort to trial and error to determine which of these myriad approaches, if any, will produce a compliant model with sufficient performance¹ to satisfy business needs. It may even turn out that despite her best efforts, the fairness constraints set by the regulators are actually impossible to satisfy to begin with, due to limitations intrinsic to the prediction task and data at hand. If there was a tool to understand the potential trade-offs exhibited by the model, even before training, it would be easier for multiple parties to effectively reconcile the conflicting components in designing fair classifiers.

Motivated by such practical considerations, we propose the *FACT (FAirness-Confusion Tensor) diagnostic* for exploring the trade-offs involving group fairness: the diagnostic provides a general framework under which the practitioners can understand both fairness–accuracy trade-offs and fairness–

¹Machine Learning Department, Carnegie Mellon University, Pittsburgh, USA ²Work paritally done during an internship at JP Morgan ³JP Morgan AI Research, New York, USA ⁴Determined AI, San Francisco, USA. Correspondence to: Joon Sik Kim <joonkim@cmu.edu>.

¹In this work, *performance* refers to classical metrics derived from the confusion matrix, e.g., accuracy, precision and fairness notions are not part of it.

performance trade-offs. At the core of our diagnostic lies the *fairness–confusion tensor*, which is the confusion matrix divided along an additional axis for protected attributes. The FACT diagnostic first expresses the majority of group fairness notions as linear/quadratic functions of the elements of this tensor. The simplicity of these functions makes it easy for them to be naturally integrated into a class of optimization problems over the elements of the tensor (not over the model parameters), which we call *performance–fairness optimality problem* (PFOP). It essentially considers the geometry of valid fairness–confusion tensors that satisfy a specified set of performance and/or fairness conditions.

By noting that many settings involve only linear notions of fairness, in this work we focus on *least-squares accuracy–fairness optimality problem* (LAFOP) and *model-specific least-squares accuracy–fairness optimality problem* (MS-LAFOP), which are specific instantiations of PFOP, each representative of model-agnostic and model-specific scenarios. In particular, for the model-agnostic case, the diagnostic allows for a comparative analysis of the *relative* difficulty of learning a classifier under additional group fairness constraints imposed. This difficulty is interpreted with respect to the Bayes error, which is the inherent difficulty of the fairness-unconstrained learning problem, hence a natural reference point.

Our contributions are:

1. to demonstrate how fairness–confusion tensor characterizes the majority of group fairness definitions in the literature as linear or quadratic functions, whose simplicity can be leveraged to formulate optimization problems suited for trade-off analysis,
2. to formulate the FACT diagnostic as a PFOP, LAFOP, and MS-LAFOP over the fairness–confusion tensor, enabling both model-agnostic and model-specific analysis of fairness trade-offs,
3. to provide a general understanding of group fairness incompatibility, which simplifies the existing results in the literature and extends them to new types,
4. to demonstrate the use of the FACT diagnostic on synthetic and real datasets, e.g. how it can be used for diagnosis of relative influence of the fairness notions on performance and other fairness conditions, and how it can be used as a post-processing method for designing fair classifiers.

2. Related Work

Fairness–confusion tensor is not a completely new notion – several work has implicitly mentioned it, mostly disregarding it as a simple computational tool that eases the

computation on an implementation level (Bellamy et al., 2018; Celis et al., 2019). It is also a natural object considered in several post-processing methods in fairness (Hardt et al., 2016; Pleiss et al., 2017), a group of algorithms that fine-tune a trained model to mitigate the unfairness while keeping the performance change minimal. Here we take a closer look at the fairness–confusion tensor itself and study how this object naturally brings together several notions of group fairness, simplifying and generalizing the analysis of inherent trade-offs within.

Quantitative definitions of group fairness exist in many different variations (Narayanan, 2018; Kleinberg et al., 2017; Chouldechova, 2017; Dwork et al., 2012; Hardt et al., 2016; Calders & Verwer, 2010; Berk et al., 2018) but few work exists to categorize these notions with a broader perspective encompassing the trade-off schemes. Verma & Rubin (2018) categorized the existing group fairness definitions based on entries and rates derived from the fairness–confusion tensor but did not explore any trade-offs and incompatibilities within. Our work extends this effort and provides a versatile geometric formalism to study the trade-offs.

Fairness–performance trade-offs have been studied in many specific cases (Calders et al., 2009; Žliobaitė, 2015; Kamiran et al., 2010; Feldman et al., 2015; Menon & Williamson, 2018; Liu et al., 2019; Zhao & Gordon, 2019), for limited definitions of fairness, performance, and models. To our knowledge, these trade-offs have not been studied in the general way we present below. Zafar et al. (2015; 2017) presented an optimization-based analysis of the trade-offs, albeit over the parameter space of a particular model.

Fairness–fairness trade-offs describe the incompatibility of multiple notions of group fairness (Kleinberg et al., 2017; Chouldechova, 2017; Pleiss et al., 2017; Berk et al., 2018) without some strong assumptions about the data and the model. Previous incompatibility results have been presented mostly in ad hoc and definition-specific manner, which our diagnostic addresses with a more general perspective for understanding incompatibilities. We show a general incompatibility result involving Calibration fairness condition, which naturally implies the result in Kleinberg et al. (2017) along with many other new ones. To the best of our knowledge, our work is the first to provide a systematic approach to diagnose both fairness–fairness and fairness–performance trade-offs together for group fairness under the same formalism.

3. The Fairness–confusion Tensor

Our key insight is that the elements of the fairness–confusion tensor encode all the information needed to study many notions of performance and group fairness. The

Name of fairness	Definition and linear system	Terms in fairness–confusion tensor
Demographic parity (DP)	$\Pr(\hat{y} = 1 \mathbf{a} = 1) = \Pr(\hat{y} = 1 \mathbf{a} = 0)$ $\mathbf{A}_{\text{DP}} = \frac{1}{N} (N_0 \ 0 \ N_0 \ 0 \ -N_1 \ 0 \ -N_1 \ 0)$	
Equality of opportunity (EOp)(Hardt et al., 2016)	$\Pr(\hat{y} = 1 y = 1, \mathbf{a} = 1) = \Pr(\hat{y} = 1 y = 1, \mathbf{a} = 0)$ $\mathbf{A}_{\text{EOP}} = \frac{1}{N} (M_0 \ 0 \ 0 \ 0 \ -M_1 \ 0 \ 0 \ 0)$	
Predictive equality (PE)(Chouldechova, 2017)	$\Pr(\hat{y} = 1 y = 0, \mathbf{a} = 1) = \Pr(\hat{y} = 1 y = 0, \mathbf{a} = 0)$ $\mathbf{A}_{\text{PE}} = \frac{1}{N} (0 \ 0 \ N_0 - M_0 \ 0 \ 0 \ 0 \ -N_1 + M_1 \ 0)$	
Equalized odds (EOd)(Hardt et al., 2016)	EOp \wedge PE	
Equal false negative rate (EFNR) ²	$\Pr(\hat{y} = 0 y = 1, \mathbf{a} = 1) = \Pr(\hat{y} = 0 y = 1, \mathbf{a} = 0)$ $\mathbf{A}_{\text{EFNR}} = \frac{1}{N} (0 \ M_0 \ 0 \ 0 \ 0 \ -M_1 \ 0 \ 0)$	
Calibration within groups (CG)(Kleinberg et al., 2017)	$\Pr(y = 1 P_\theta(\mathbf{x}) = s, \mathbf{a} = 1) = \Pr(y = 1 P_\theta(\mathbf{x}) = s, \mathbf{a} = 0) = s$ $\mathbf{A}_{\text{CG}} = \begin{pmatrix} 1 - v_1 & 0 & -v_1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 - v_0 & 0 & -v_0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 - v_1 & 0 & -v_1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 - v_0 & -v_0 \end{pmatrix}$	
Positive class balance (PCB)(Kleinberg et al., 2017)	$\mathbb{E}(P_\theta y = 1, \mathbf{a} = 1) = \mathbb{E}(P_\theta y = 1, \mathbf{a} = 0)$ $\mathbf{A}_{\text{PCB}} = \min_a (M_a) \begin{pmatrix} \frac{v_1}{M_1} & \frac{v_0}{M_1} & 0 & 0 & -\frac{v_1}{M_0} & -\frac{v_0}{M_0} & 0 & 0 \end{pmatrix}$	
Negative class balance (NCB)(Kleinberg et al., 2017)	$\mathbb{E}(P_\theta y = 0, \mathbf{a} = 1) = \mathbb{E}(P_\theta y = 0, \mathbf{a} = 0)$ $\mathbf{A}_{\text{NCB}} = \min_a (N_a - M_a) \begin{pmatrix} 0 & 0 & \frac{v_1}{N_1 - M_1} & \frac{v_0}{N_1 - M_1} & 0 & 0 & -\frac{v_1}{N_0 - M_0} & -\frac{v_0}{N_0 - M_0} \end{pmatrix}$	
Relaxed Equalized Odds (REOD)(Pleiss et al., 2017)	$\alpha_0 FPR_0 + \beta_0 FNR_0 = \alpha_1 FPR_1 + \beta_1 FNR_1$ $\mathbf{A}_{\text{REOD}} = \begin{pmatrix} \frac{\beta_1}{M_1} & \frac{\alpha_1}{N_1 - M_1} & 0 & 0 & \frac{\beta_0}{M_0} & \frac{\alpha_0}{N_0 - M_0} & 0 & 0 \end{pmatrix} / N$	
Predictive parity (PP)(Chouldechova, 2017)	$\Pr(y = 1 \hat{y} = 1, \mathbf{a} = 1) = \Pr(y = 1 \hat{y} = 1, \mathbf{a} = 0)$ $\frac{1}{2} \mathbf{z}^T \mathbf{B}_{\text{PP}} \mathbf{z} = (TP_1 F P_0 - TP_0 F P_1) / N^2$	
Equal false omission rate (EFOR) ¹	$\Pr(y = 1 \hat{y} = 0, \mathbf{a} = 1) = \Pr(y = 1 \hat{y} = 0, \mathbf{a} = 0)$ $\frac{1}{2} \mathbf{z}^T \mathbf{B}_{\text{EFOR}} \mathbf{z} = (TN_1 F N_0 - TN_0 F N_1) / N^2$	
Conditional accuracy equality (CA)(Berk et al., 2018)	PP \wedge EFOR	

¹ To our knowledge, EFOR has not been described in literature in isolation, but is used in the definition of conditional accuracy equality (CA)(Berk et al., 2018).
² Defined implicitly in (Chouldechova, 2017).

Table 1. Some common group fairness definitions and corresponding abbreviations used throughout the paper in terms of linear functions $\phi(\mathbf{z}) = \mathbf{A}\mathbf{z}$ or quadratic functions $\phi(\mathbf{z}) = \frac{1}{2} \mathbf{z}^T \mathbf{B}\mathbf{z}$ that appear in the performance–fairness optimality problem (5). There are two groups separated by the horizontal line: those that are specified by linear functions (above), or quadratic functions (below). The graphical notation is described in Section 3. P_θ is the probability produced by a model (parameterized by θ) of $\hat{y} = 1$. The fairness functions ϕ are uniquely defined only up to a normalization factor and overall sign.

fairness–confusion tensor is simply the stack of confusion matrices for each protected attribute a , as shown in Table 2. We focus on the simplest case, with one binary protected attribute $a \in \{0, 1\}$, and a binary classifier $\hat{y} \in \{0, 1\}$ for a binary prediction label $y \in \{0, 1\}$.²

$a = 1$	$y = 1$	$y = 0$	$a = 0$	$y = 1$	$y = 0$
$\hat{y} = 1$	TP ₁	FP ₁	$\hat{y} = 1$	TP ₀	FP ₀
$\hat{y} = 0$	FN ₁	TN ₁	$\hat{y} = 0$	FN ₀	TN ₀

Table 2. The fairness–confusion tensor, showing the two planes corresponding to the confusion matrix for each of the favored ($a = 1$) and disfavored groups ($a = 0$).

Let us denote the elements of the fairness–confusion tensor as TP_a, FP_a, FN_a, TN_a , each element with subscripts indicating a , N be the number of data points, $N_a = TP_a + FN_a + FP_a + TN_a$ be the number of data points in each group $a \in \{0, 1\}$, and $M_a = TP_a + FN_a$ be the number of positive-class instances ($y = 1$) for each group. Assume N, N_a and M_a are known constants. Unraveling the fairness–confusion tensor into an 8-dimensional vector, we write it as

$$\mathbf{z} = (TP_1, FN_1, FP_1, TN_1, TP_0, FN_0, FP_0, TN_0)^T / N,$$

²The arguments generalize to multiple and non-binary protected attributes with high-dimensional tensors.

normalized and constrained to lie on $\mathcal{K} = \{\mathbf{z} \geq 0 : \mathbf{A}_{\text{const}} \mathbf{z} = \mathbf{b}_{\text{const}}, \|\mathbf{z}\|_1 = 1\}$, where $\mathbf{A}_{\text{const}}$ and $\mathbf{b}_{\text{const}}$ encode marginal sum constraints of the dataset (e.g., $TP_a + FN_a = M_a$) in matrix notations:

$$\mathbf{A}_{\text{const}} = \begin{pmatrix} 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \end{pmatrix},$$

$$\mathbf{b}_{\text{const}} = (N_1, M_1, N_0, M_0)^T / N.$$

We show below that some typical notions of group fairness can be reformulated as simple functions of \mathbf{z} , namely as a form of $\phi(\mathbf{z}) = 0$.

Demographic parity (DP) states that each protected group should receive positive prediction at an equal rate: $\Pr(\hat{y} = 1 | \mathbf{a} = 1) = \Pr(\hat{y} = 1 | \mathbf{a} = 0)$, which is equivalent to $(TP_1 + FP_1) / N_1 = (TP_0 + FP_0) / N_0$, or also the linear system $\phi(\mathbf{z}) = \mathbf{A}_{\text{DP}} \mathbf{z} = 0$, where

$$\mathbf{A}_{\text{DP}} = (N_0 \ 0 \ N_0 \ 0 \ -N_1 \ 0 \ -N_1 \ 0) / N. \quad (1)$$

The choice of normalization, $1/N$, ensures that the matrix coefficients are in $[0, 1]$. We will refer to these matrices \mathbf{A} that encode information about the fairness conditions as fairness matrices.

Predictive parity (PP) (Chouldechova, 2017) states that the likelihood of being in the positive class given the positive prediction is the same for each group: $\Pr(y = 1|\hat{y} = 1, \mathbf{a} = 1) = \Pr(y = 1|\hat{y} = 1, \mathbf{a} = 0)$, which is equivalent to $\frac{TP_1}{TP_1+FP_1} = \frac{TP_0}{TP_0+FP_0} \iff \frac{TP_1}{FP_0} = \frac{FP_1}{TP_0}$. Unlike for DP, the marginal sum constraints do not relate TP_a and FP_a , so this notion of fairness is *not* linear in the fairness–confusion tensor. PP actually can be expressed using a *quadratic* form:

$$\phi(\mathbf{z}) = \frac{1}{2} \mathbf{z}^T \mathbf{B}_{PP} \mathbf{z} = 0, \quad \mathbf{B}_{PP} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}. \quad (2)$$

Calibration within groups (CG) (Kleinberg et al., 2017), when specialized to binary classifiers and binary protected classes, can be written as the system of equations $FN_a = v_0(FN_a + TN_a)$; $TP_a = v_1(TP_a + FP_a)$, where the v_i s are scores satisfying $0 \leq v_0 < v_1 \leq 1$ and have no implicit dependence on any entries of the fairness–confusion tensor. We can rewrite this this condition explicitly as the matrix equation $\phi(\mathbf{z}) = \mathbf{A}_{CG} \mathbf{z} = 0$ with a fairness matrix

$$\mathbf{A}_{CG} = \begin{pmatrix} 1 - v_1 & 0 & -v_1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 - v_0 & 0 & -v_0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 - v_1 & 0 & -v_1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 - v_0 & 0 & -v_0 \end{pmatrix}. \quad (3)$$

Equalized odds (EOd) (Hardt et al., 2016) states that true-positive rates and false-positive rates are the same for both groups, which can be expressed as a linear system $\phi(\mathbf{z}) = \mathbf{A}_{EOd} \mathbf{z} = 0$ with a fairness matrix

$$\mathbf{A}_{EOd} = \frac{1}{N} \begin{pmatrix} M_0 & 0 & 0 & 0 & -M_1 & 0 & 0 & 0 \\ 0 & 0 & N_0 - M_0 & 0 & 0 & 0 & -N_1 + M_1 & 0 \end{pmatrix} \quad (4)$$

where each row respectively corresponds to conditions for Equality of Opportunity (EOP) (Hardt et al., 2016) and Predictive Equality (PE) (Chouldechova, 2017). Likewise, vertically stacking multiple fairness matrices results in a fairness matrix corresponding to the conjunction of different fairness notions.

In Table 1 we generalize this formulation to a wide majority of group fairness definitions in the literature, along with their abbreviations used throughout the paper. We find that most of the definitions take either linear or quadratic form with respect to \mathbf{z} . We further introduce a graphical notation to help visualize which components of the fairness–confusion tensor participate in the fairness definition. Depict the fairness–confusion tensor as $\begin{bmatrix} \square & \square \\ \square & \square \end{bmatrix}$, with the left matrix for the favored class ($a = 1$) and the right matrix for the

disfavored class ($a = 0$). Since each component of \mathbf{z} corresponds to some element of the fairness–confusion tensor, we shade each component that appears in the equation. Blue shading denotes the favored class, while red shading denotes the disfavored class. We further distinguish two kinds of dependencies. Components that have a nonzero coefficient in the matrix are shaded fully. However, the values of these coefficients themselves can depend on other components, albeit implicitly, and we shade these implicit components in a lighter shade. Putting this all together, we can represent DP in (1) graphically as $\begin{bmatrix} \text{blue} & \text{red} \\ \text{blue} & \text{red} \end{bmatrix}$, EOd as $\begin{bmatrix} \text{blue} & \text{red} \\ \text{blue} & \text{red} \end{bmatrix} \wedge \begin{bmatrix} \text{blue} & \text{red} \\ \text{blue} & \text{red} \end{bmatrix}$, PP as $\left(\begin{bmatrix} \text{blue} & \text{red} \\ \text{blue} & \text{red} \end{bmatrix} \right)^{(2)}$, with the superscript denoting the quadratic order of the term. As shown in the third column of Table 1, all group fairness notions can be effectively described in this notation.

4. Optimization over the Fairness–confusion Tensor

The fairness–confusion tensor \mathbf{z} allows for a succinct linear and quadratic characterization of group fairness definitions in the literature. We naturally consider the following family of optimization problems over $\mathbf{z} \in \mathcal{K}$, where the objective function is constructed so that the solution reflects trade-offs between fairness and performance.

Definition 1. Let $f^{(i)} : \mathcal{K} \rightarrow [0, 1]$ be performance metrics (indexed by i) with best performance 0 and worst performance 1, $\phi^{(j)}(\mathbf{z})$ be fairness functions (indexed by j) with μ_i, λ_j be real constants with $\mu_0 = 1$. Then, the performance–fairness optimality problem (PFOP) is a class of optimization problem of form:

$$\arg \min_{\mathbf{z} \in \mathcal{K}} \sum_{i \geq 0} \mu_i f^{(i)}(\mathbf{z}) + \sum_{j \geq 0} \lambda_j \phi^{(j)}(\mathbf{z}) \quad (5)$$

PFOP is a general optimization problem containing two groups of terms; the first quantifying performance loss; the second quantifying unfairness. The restriction $\mathbf{z} \in \mathcal{K}$ is necessary to ensure that \mathbf{z} is a valid fairness–confusion tensor that obeys the requisite marginal sums. In our discussion below, it will be convenient to consider solutions with explicit bounds on their optimality.

Definition 2. Let $\epsilon \geq 0$ and $\delta \geq 0$. Then, a (ϵ, δ) -solution to the PFOP is a \mathbf{z} that satisfies (5) such that $\sum_j \lambda_j \phi^{(j)}(\mathbf{z}) \leq \epsilon$ and $\sum_i \mu_i f^{(i)}(\mathbf{z}) \leq \delta$.

The parameters ϵ and δ represent the sum total of deviation from perfect fairness and perfect predictive performance respectively. Unless otherwise stated, the rest of the paper is dedicated to analyzing one of the simplest instantiations of PFOP, defined below.

Definition 3. The least-squares accuracy–fairness optimality problem (LAFOP) is a PFOP with accuracy (or classification error rate) as the performance function $f^{(0)}$, and

$K \geq 1$ fairness constraints in the form of a fairness matrix \mathbf{A} (each row indexed by j), with

$$\begin{aligned} \phi^{(j)}(\mathbf{z}) &= (\mathbf{A}_{j,*}\mathbf{z})^2, \quad j = 0, \dots, K-1 \\ f^{(0)}(\mathbf{z}) &= (\mathbf{c} \cdot \mathbf{z})^2, \\ \mathbf{c} &= (0, 1, 1, 0, 0, 1, 1, 0)^T, \\ \lambda &= \lambda_0 = \dots = \lambda_{K-1}. \end{aligned} \quad (6)$$

In other words, LAFOP is the problem

$$\arg \min_{\mathbf{z} \in \mathcal{K}} (\mathbf{c} \cdot \mathbf{z})^2 + \lambda \|\mathbf{A}\mathbf{z}\|_2^2, \quad (7)$$

where $\mathbf{c} \cdot \mathbf{z}$ encodes the usual notion of classification error, and \mathbf{A} encodes K linear fairness functions stacked together as the regularizer. A single hyperparameter λ specifies the relative importance of satisfying the fairness constraints while optimizing classification performance, with $\lambda = 0$ considering only performance and disabling all fairness constraints, and $\lambda = \infty$ imposing fairness constraints without regard to accuracy.

LAFOP is a convex optimization problem which is simple to analyze. Despite its simplicity, LAFOP encompasses many situations involving linear notions of fairness, allowing us to reason about multiple fairness constraints as well as fairness–accuracy trade-offs under versatile scenarios.

4.1. Reduction to a post-processing method for fair classification

PFOP and LAFOP do not assume anything about the model, therefore are designed to be model-agnostic. In this section we highlight the versatility of LAFOP by showing that adding a model-specific constraint on LAFOP reduces it to a post-processing algorithm for fair classification.

Post-processing method, in particular for EOd as introduced in [Hardt et al. \(2016\)](#), solves the following optimization problem for \tilde{Y} , which is a post-processed, supposedly fair, classifier, given \hat{Y} , a vanilla classifier:

$$\begin{aligned} \min_{\tilde{Y}} \mathbb{E}l(\tilde{Y}, Y) \text{ such that } \gamma_0(\tilde{Y}) &= \gamma_1(\tilde{Y}) \\ \text{and } \gamma_0(\tilde{Y}) &\in P_0(\hat{Y}), \gamma_1(\tilde{Y}) \in P_1(\hat{Y}) \end{aligned} \quad (8)$$

where $\gamma_a(\tilde{Y})$ represents EOd constraints for \tilde{Y} as a tuple of (FPR_a, TPR_a) , and $P_a(\hat{Y})$ is a model-specific set of feasible γ_a values, defined as $P_a(\hat{Y}) = \text{convhull}\{(0, 0), \gamma_a(\hat{Y}), \gamma_a(1 - \hat{Y}), (1, 1)\}$. All the components of (8) can be rewritten in terms of $\tilde{\mathbf{z}}$ and $\hat{\mathbf{z}}$, the fairness–confusion tensors corresponding to the classifiers \hat{Y} and \tilde{Y} respectively. This yields a LAFOP over $\tilde{\mathbf{z}}$ with additional model-specific constraints derived from $\hat{\mathbf{z}}$ on the solution space. More formally, we have the following optimization problem for post-processing:

Definition 4. Given a classifier to be post-processed and its corresponding fairness–confusion tensor $\hat{\mathbf{z}}$, the model-specific LAFOP (MS-LAFOP) for EOd is the variant of LAFOP with model-specific constraints on the solution space as the following:

$$\arg \min_{\tilde{\mathbf{z}} \in \hat{\mathcal{K}}} (\mathbf{c} \cdot \tilde{\mathbf{z}})^2 + \lambda \|\mathbf{A}_{\text{EOd}}\tilde{\mathbf{z}}\|_2^2, \text{ where} \quad (9)$$

where

$$\begin{aligned} \hat{\mathcal{K}} &= \{\tilde{\mathbf{z}} \geq 0 : \mathbf{A}_{\text{const}}\tilde{\mathbf{z}} = \mathbf{b}_{\text{const}}, \|\tilde{\mathbf{z}}\|_1 = 1, \\ &\beta_a(\tilde{\mathbf{z}}) \in \text{convhull}\{(0, 0), \beta_a(\hat{\mathbf{z}}), \beta_a(1 - \hat{\mathbf{z}}), (1, 1)\} \forall a\} \end{aligned}$$

with β_a expressing (FPR_a, TPR_a) tuples computed from the corresponding fairness–confusion tensor of group a .

From the solution of MS-LAFOP, it is possible to compute mixing rates for post-processing the given classifier. We note that MS-LAFOP can be extended to other group fairness notions as long as the model-specific constraints are accordingly set up for them. For more details, refer to [Appendix G.3](#).

5. Incompatible Group Fairness Definitions

In this section, we show how LAFOP yields a more general view of understanding group fairness incompatibility results. As $\lambda \rightarrow \infty$, for linear fairness functions $\phi^{(i)}(\mathbf{z}) = \mathbf{A}^{(i)}\mathbf{z}$, LAFOP becomes equivalent to solving the following linear system of equations:

$$\begin{pmatrix} \mathbf{A}^{(0)} \\ \vdots \\ \mathbf{A}^{(K-1)} \\ \mathbf{A}_{\text{const}} \end{pmatrix} \mathbf{z} = \begin{pmatrix} 0 \\ \vdots \\ 0 \\ \mathbf{b}_{\text{const}} \end{pmatrix}, \mathbf{z} \geq 0, \quad (10)$$

Notice the compatibility of fairness conditions encoded by these K fairness matrices $\mathbf{A}^{(i)}$ is equivalent to having infinitely many solutions to the above linear system. We formally define (in)compatibility of fairness notions below based on this observation.

Definition 5. Let $\Phi = \{\phi^{(i)}\}_{i=0}^{K-1}$ be a set of linear fairness functions, encoded in a fairness matrix \mathbf{A} (of which each row corresponds to $\phi^{(i)}$), and let ρ be the number of solutions for the system in (10). If $\rho = 0$, then Φ is said to be incompatible. Otherwise, Φ is compatible. When Φ is incompatible, some additional set of constraints on the dataset or the model is required for it to be compatible.

This means that in general, incompatibility results among the group fairness definitions can be proven simply by asking if and when solutions exist to their corresponding linear system of form (10).

Sets of fairness definitions	Necessary conditions
{CG, PP, DP, and any of EOp, PE, PCB, NCB, EFOR}	$M_0 = M_1$ and $N_0 = N_1$
{CG, DP, and any of EOp, PE, PCB, NCB, EFOR}	EBR only
{CG,EOp}, {CG,PCB}, {CG,EOp,PCB}, {CG,EFOR,EOp}, {CG,EFOR,PCB}, {CG,EFOR,EOp,PCB}	$v_0 = 0$ or EBR
{CG,PE}, {CG,NCB}, {CG,EOp,NCB}, {CG,EFOR,PE}, {CG,EFOR,NCB}, {CG,EFOR,EOp,NCB}	$v_1 = 1$ or EBR
{CG,EoD} (Pleiss et al., 2017), {CG, PCB, NCB} (Kleinberg et al., 2017), {CG,EoD,PCB,NCB}, {CG,EFOR,EoD}, {CG,EFOR,PCB,NCB}, {CG,EFOR,EoD,PCB,NCB}	$(v_0 = 0 \text{ and } v_1 = 1)$ or EBR

Table 3. Some sets of fairness definitions containing Calibration(CG), which are incompatible in the sense of Definition 5 (left-column), together with their necessary conditions to be compatible (right column). EBR is the equal base rate condition, $M_0/N_0 = M_1/N_1$. For other abbreviations, refer to Table 1. These are all special cases of Theorem 1, while not exhaustive.

5.1. The incompatibility involving CG

We introduce a general incompatibility result involving CG that leads to many other new results as well as the one from Kleinberg et al. (2017).

Theorem 1. *Let $B = 2$ be the number of bins in the definition of calibration within groups fairness (CG) (Kleinberg et al., 2017), and v_0, v_1 be the scores, with $0 \leq v_0 < v_1 \leq 1$, and $K > 1$ with $\phi^{(0)}(\mathbf{z}) = \mathbf{A}_{CG}\mathbf{z}$. Then, the corresponding (10) has the only solution*

$$z_0 = \frac{1}{N(v_1 - v_0)} \begin{pmatrix} v_1(M_1 - N_1v_0) \\ v_0(-M_1 + N_1v_1) \\ (1 - v_1)(M_1 - N_1v_0) \\ (1 - v_0)(-M_1 + N_1v_1) \\ v_1(M_0 - N_0v_0) \\ v_0(-M_0 + N_0v_1) \\ (1 - v_1)(M_0 - N_0v_0) \\ (1 - v_0)(-M_0 + N_0v_1) \end{pmatrix}, \quad (11)$$

and only when

$$0 \leq v_0 \leq \min_a \left(\frac{M_a}{N_a} \right) \leq \max_a \left(\frac{M_a}{N_a} \right) \leq v_1 \leq 1. \quad (12)$$

Otherwise, no solution exists.

Theorem 1 yields other extended results regarding the incompatibility of CG and other notions of fairness. As one canonical instance, simply substituting z_0 in (11) to the linear system of the form in (10) with PCB and NCB fairness matrices yields the following corollary, which is equivalent to the result presented in Kleinberg et al. (2017) (proof is in Appendix B).

Corollary 1 (Re-derivation of (Kleinberg et al., 2017)). *Consider a classifier that satisfies CG, PCB and NCB fairness simultaneously. Then, at least one of the following statements is true:*

1. the data have equal base rates for each class a , i.e. $M_0/N_0 = M_1/N_1$, or
2. the classifier has perfect prediction, i.e. $v_0 = 0$ and $v_1 = 1$.

Similar approach can be applied to derive incompatibilities of CG with other linear and quadratic notions of fairness as below (proofs in Appendix C, Appendix D).

Corollary 2. *(Linear notion of fairness: DP) Consider a classifier that satisfies CG and DP fairness simultaneously. Then, the data have equal base rates for each group a .*

Corollary 3. *(Quadratic notion of fairness: PP) Consider a classifier that satisfies CG and PP fairness simultaneously. Then, at least one of the following is true:*

1. $v_0 = (M_1 - M_0)/(N_1 - N_0)$.
2. $v_1 = 1$.

From Theorem 1 and its corollaries, we curate the extended incompatibility results involving CG in Table 3 along with conditions for compatibility. To our knowledge, all cases other than the bottom row of the table are new.

5.2. The incompatibility of {PE, EFNR, PP}

Using the same logic as the previous section, we re-derive an incompatibility result in Chouldechova (2017) and provide more precise necessary conditions for compatibility. For details of the proof, refer to Appendix E.

Theorem 2 (Restatement of Chouldechova (2017)). *Consider a classifier that satisfies {PE, EFNR, PP}. Then, at least one of these statements must be true:*

1. The classifier has no true positives.
2. The classifier has no false positives.
3. Each protected class has the same base rate.

Theorem 2 systematically shows that equal false positive rates, equal false negative rates, and predictive parity are compatible only under specific data/model-dependent circumstances, that were otherwise not clear in the original statements in Chouldechova (2017).

6. Experiments

In this section we show how the FACT diagnostic can practically show the relative impact of several notions of fairness

on accuracy on synthetic and real datasets³. First we introduce FACT Pareto frontiers which characterize a model’s achievable accuracy for a given set of fairness conditions, as a tool for understanding the trade-offs and contextualizing some recent works in fair classification (Section 6.2). We then explore a model-agnostic assessment of multiple fairness conditions via LAFOP (Section 6.3, Appendix G.2), as well as a model-specific assessment of post-processing methods in fair classification via MS-LAFOP (Section 6.4, Appendix G.3).

6.1. Datasets

We study a synthetic dataset similar to that in Zafar et al. (2015), consisting of two-dimensional features along with a single binary protected attribute that is either sampled from an independent Bernoulli distribution (“unbiased” variant, denoted $\mathbf{S(U)}$), or sampled dependent on the features (“biased” variant, denoted $\mathbf{S(B)}$). The synthetic dataset consists of two-dimensional data $\mathbf{x} = (x_0, x_1)$ that follow the Gaussian distributions

$$\begin{aligned} \mathbf{x}|y=1 &\sim \mathcal{N}\left(\begin{pmatrix} 2 \\ 2 \end{pmatrix}, \begin{pmatrix} 5 & 1 \\ 1 & 5 \end{pmatrix}\right) \\ \mathbf{x}|y=0 &\sim \mathcal{N}\left(\begin{pmatrix} -2 \\ -2 \end{pmatrix}, \begin{pmatrix} 10 & 1 \\ 1 & 3 \end{pmatrix}\right). \end{aligned} \quad (13)$$

For the $\mathbf{S(U)}$ dataset, the protected attribute value is independent of \mathbf{x} and y , and is instead distributed according to the Bernoulli distribution $a \sim \mathcal{B}(\frac{1}{2})$. This notion of fairness was described in (Calders et al., 2009).

For the $\mathbf{S(B)}$ dataset, the protected attribute value is assigned as $a|\mathbf{x} = \text{sgn}(x_0)$, which corresponds to a situation when some features (but not all) encode a protected attribute.

We also study the UCI Adult dataset (Dua & Graff, 2017), a census dataset used for income classification tasks where we consider sex as the protected attribute of interest.

6.2. FACT Pareto frontiers

With LAFOP and MS-LAFOP, one can naturally consider a FACT Pareto frontier of accuracy and fairness by plotting (ϵ, δ) values of the (ϵ, δ) -solutions. In this section, we want to highlight the use of this frontier in the context of several published results in the literature as well as its implications.

The FACT Pareto frontier can be computed both in model-agnostic (MA) and model-specific (MS) scenarios by solving LAFOP and MS-LAFOP respectively, and Figure 1 shows such example on the Adult dataset for EOD fairness. We also consider three fair classification models: **FGP** (Tan et al., 2020), **Op.** (Zafar et al., 2015), and **Eq.Odd.** (Hardt et al., 2016), individually representing three different ap-




Figure 1. Model-agnostic (MA) and model-specific (MS) FACT Pareto frontiers of equalized odds on the Adult dataset. Three fair models (FGP, Eq.Odd., Op.) are shown in context by varying the strength of the fairness condition imposed, along with some baseline models (LR, SVM, RF, ConstantPrediction). The MA frontier should be interpreted relative to the Bayes error because it is oblivious to it — $\delta = 0$ means that the upper bound of the accuracy is the accuracy of the Bayes classifier, not 1. The MS frontier on the other hand provides realistic more bounds.

proaches one can take in training fair models (imposing fairness before, during, or after training). Some baseline models (logistic regression, SVM, random forest) are also plotted for reference, and a perfectly fair classifier (ConstantPredict: predicting all instances to be negative) on the bottom right corner is considered as an edge-case.

It is important to note that the MA FACT Pareto frontier should be interpreted as characterizing the model’s achievable accuracy *relative* to the Bayes error (i.e., the degree to which the added fairness constraints adversely impact the Bayes error), which in this case is empirically estimated at around 0.12 from a wide range of ML models that have been tested on the Adult dataset (Chakrabarty & Biswas, 2018). This relatively less realizable bound calls for a model-specific counterpart, the MS FACT Pareto frontier, which limits the frontier to be derived from a given pre-trained classifier. As shown in Figure 1, it indeed provides a more reasonable frontier for the models considered.

Placing different types of classifiers on the frontier, it is easy to visually grasp strengths and weaknesses of each models. FGP seems to outperform all other models in terms of the trade-off, while Op and EqOdd suffer more from early accuracy drops. The frontier further informs that for any model trained, only for fairness gaps below 10^{-2} will the accuracy start to suffer. Such understanding of the trade-offs will be helpful in anticipating practical limitations of models to be trained, as well as in comparing multiple models to determine which is better-suited for different situations.

In the rest of the following sections and figures, for the model-agnostic analysis, δ should be interpreted in reference to the Bayes error, i.e $\delta = 0$ means that the upper bound of the best-achievable accuracy is the accuracy of the Bayes

³Code available: github.com/wnstlr/FACT




Figure 2. Model-agnostic FACT Pareto frontier for different groups of fairness notions (colored and grouped according to their convergence value as $\epsilon \rightarrow 0$) for three datasets (Section 6.1). The bottom two groups of fairness notions are incompatible (black, red), hence the halted trajectories before reaching smaller values of ϵ . Similar convergence behaviors within the fairness groups in blue reflect the dominance of $\{\text{EOd}, \text{DP}\}$ – any additional fairness notions added on top of these have no impact on the convergence value. Best viewed in color.

classifier, not 1.

6.3. Model-agnostic scenario with multiple fairness conditions

We are now interested in how a *group* of fairness conditions simultaneously affect accuracy. This can be assessed by looking at the shape of the MA FACT Pareto frontier of LAFOP with multiple fairness constraints, particularly δ values of (ϵ, δ) -solutions when ϵ is varied to be zero (or very close to it) on multiple fairness notions. Figure 2 shows this in two different ways: (i) (ϵ, δ) -solutions obtained when fairness conditions are imposed as hard inequality constraints instead of as regularizers, i.e. solving $\arg \min_{\mathbf{z} \in \mathcal{K}} (\mathbf{c} \cdot \mathbf{z})^2$ s.t. $\|\mathbf{A}\mathbf{z}\|_2^2 \leq \epsilon$ (solid line), and (ii) (ϵ, δ) -solutions obtained from the LAFOP (7) while varying λ s (crosses). Different groups of fairness notions are colored according to their convergence behaviors.

Similar trajectories and convergence of the curves allow us to identify fairness notions that come “for free” given some others, in terms of additional accuracy drops. In other words, the Pareto frontiers are effective at demonstrating the relative strength of the fairness notions within a group. For instance, under $\{\text{EOd}, \text{DP}\}$ (third group, blue) the best attainable accuracy drops by over 60 percent for S(U) and S(B), but we also observe that adding CB, PE, and/or PCB on top of them causes no additional accuracy drop – $\{\text{EOd}, \text{DP}\}$ essentially determines δ for the entire group of fairness notions in blue.

The MA FACT Pareto frontiers for multiple fairness conditions also show not only the existing incompatibility of the fairness notions, but also how much relaxation is required for them to be approximately compatible. The halted trajectories before hitting much smaller ϵ for the bottom two groups in black and red clearly verify this. Because the S(U) dataset has a smaller base rate gap between the groups compared to the Adult or the S(B) dataset by design, the incompatibility in S(U) becomes only visible at a much

smaller ϵ value.

Taking a more macroscopic perspective, the MA FACT Pareto frontiers also show which dataset allows overall better trade-off scheme compared to the others. Because the S(U) dataset was designed to be less biased compared to the S(B) dataset, it exhibits significantly smaller drop in overall accuracy, particularly for the green group involving DP. The way S(U) was designed aligns with this observation, as the sensitive attributes were randomly sampled independently from the features. However, EOd and DP together (in blue) drives down the accuracy just like the biased counterpart, which demonstrates how conservative EOd fairness is for these datasets.

More observations and experiments are presented in Appendix G.2. It is possible to further extend these analyses to an arbitrary number of fairness constraints imposed on LAFOP, as well as to other performance metrics like precision or recall as seem fit.

6.4. Model-specific scenario with post-processing methods

While the MA FACT Pareto frontier shows a broader trade-off landscape for any classifiers, model-specific analysis using MS-LAFOP in (9) can be helpful in practice with more reasonable MS Pareto frontiers. Also after solving the MS-LAFOP, its solution can be used to compute the mixing rates for post-processing any given classifier just like done in [Hardt et al. \(2016\)](#). For more details, refer to Appendix G.3.

Figure 3 shows the MS FACT Pareto frontier of EOd computed from MS-LAFOP for the Adult dataset (it is a zoomed-in version of the MA FACT Pareto frontier in Figure 1). We also plot two types of post-processed classifiers: EOd-solutions using the algorithm in [Hardt et al. \(2016\)](#) (circles), and FACT-solutions using MS-LAFOP (stars). EOd solutions undergo steeper trade-off while the FACT-solutions




Figure 3. Model-specific FACT Pareto frontier of EOD on Adult dataset. Compared to the model-agnostic frontier, it yields a more realizable bounds on the trade-off between fairness and accuracy. Post-processed solutions for the given classifiers (crosses) using the algorithm in (Hardt et al., 2016) (circles, EOD-solution) and FACT (stars, FACT-solution) are also shown. The FACT-solutions suffer significantly less from the trade-off, yielding competitive accuracy to the original classifiers while achieving smaller fairness gaps compared to the EOD-solutions.

are able to find a better configuration with smaller fairness gaps, retaining a competitive accuracy level to the original classifier (cross).

7. Conclusions

The FACT diagnostic facilitates systematic reasoning about different kinds of trade-offs involving arbitrarily many notions of performance and group fairness notions, which all can be expressed as functions of the fairness–confusion tensor. In our formalism, the majority of group fairness definitions in the literature are in fact linear or quadratic thus are easy to be imposed as constraints to the PFOP. The FACT diagnostic further benefits from elementary linear algebra and convex optimization to provide a unified perspective of viewing fairness–fairness trade-offs and fairness–performance trade-offs. We have also empirically demonstrated the practical use of the FACT diagnostic in several scenarios. Many of the presented results require only linear fairness functions and accuracy, as in the LAFOP/MS-LAFOP setting. Nevertheless, it is easy to extend this to quadratic fairness functions with more varied performance metrics depending on different use cases. We also briefly introduce a small theoretical result regarding fairness–accuracy trade-offs using the FACT diagnostic in Appendix F, which deserves further analysis.

Acknowledgements

We thank Valerie Chen, Jeremy Cohen, Amanda Coston, Mikhail Khodak, Jeffrey Li, Liam Li, Gregory Plumb, Nick Roberts, and Samuel Yeom for helpful feedback and discussions. This work was supported in part by DARPA FA875017C0141, the National Science Foundation grants

IIS1705121 and IIS1838017, an Okawa Grant, a Google Faculty Award, an Amazon Web Services Award, a JP Morgan A.I. Research Faculty Award, and a Carnegie Bosch Institute Research Award. JSK acknowledges support from Kwanjeong Educational Fellowship. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of DARPA, the National Science Foundation, or any other funding agency.

This paper was prepared for information purposes by the Artificial Intelligence Research group of JPMorgan Chase & Co and its affiliates (“JP Morgan”), and is not a product of the Research Department of JP Morgan. JP Morgan makes no representation and warranty whatsoever and disclaims all liability, for the completeness, accuracy or reliability of the information contained herein. This document is not intended as investment research or investment advice, or a recommendation, offer or solicitation for the purchase or sale of any security, financial instrument, financial product or service, or to be used in any way for evaluating the merits of participating in any transaction, and shall not constitute a solicitation under any jurisdiction or to any person, if such solicitation under such jurisdiction or to such person would be unlawful. © 2020 JPMorgan Chase & Co. All rights reserved.

References

- Bellamy, R. K., Dey, K., Hind, M., Hoffman, S. C., Houde, S., Kannan, K., Lohia, P., Martino, J., Mehta, S., Mojsilovic, A., et al. Ai fairness 360: An extensible toolkit for detecting, understanding, and mitigating unwanted algorithmic bias. *arXiv preprint arXiv:1810.01943*, 2018.
- Berk, R., Heidari, H., Jabbari, S., Kearns, M., and Roth, A. Fairness in criminal justice risk assessments: The state of the art. *Sociological Methods & Research*, 2018. doi: 10.1177/0049124118782533.
- Biddle, D. *Adverse impact and test validation: A practitioner’s guide to valid and defensible employment testing*. Gower Publishing, Ltd., 2006.
- Calders, T. and Verwer, S. Three naive bayes approaches for discrimination-free classification. *Data Mining and Knowledge Discovery*, 21(2), September 2010. doi: 10.1007/s10618-010-0190-x.
- Calders, T., Kamiran, F., and Pechenizkiy, M. Building classifiers with independency constraints. In *2009 IEEE International Conference on Data Mining Workshops*, 2009. doi: 10.1109/ICDMW.2009.83.
- Celis, L. E., Huang, L., Keswani, V., and Vishnoi, N. K. Classification with fairness constraints: A meta-algorithm

- with provable guarantees. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 2019.
- Chakrabarty, N. and Biswas, S. A statistical approach to adult census income level prediction. In *2018 International Conference on Advances in Computing, Communication Control and Networking (ICACCCN)*. IEEE, 2018.
- Chouldechova, A. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big Data*, 5(2), 2017.
- Crawford, K., Dobbe, R., Dryer, T., Fried, G., Green, B., Kaziunas, E., Kak, A., Mathur, V., McElroy, E., Sánchez, A. N., Raji, D., Rankin, J. L., Richardson, R., Schultz, J., West, S. M., and Whittaker, M. Ai now 2019 report, 2019. URL https://ainowinstitute.org/AI_Now_2019_Report.html.
- Dantzig, G. B. Linear programming and extensions. Technical Report R-366-PR, RAND Corporation, Santa Monica, California, 1963. URL <https://www.rand.org/content/dam/rand/pubs/reports/2007/R366part1.pdf>.
- Dua, D. and Graff, C. UCI machine learning repository. University of California, Irvine, School of Information and Computer Sciences, 2017. URL <http://archive.ics.uci.edu/ml>.
- Dwork, C., Hardt, M., Pitassi, T., Reingold, O., and Zemel, R. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, 2012. doi: 10.1145/2090236.2090255.
- Feldman, M., Friedler, S. A., Moeller, J., Scheidegger, C., and Venkatasubramanian, S. Certifying and removing disparate impact. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '15, 2015. doi: 10.1145/2783258.2783311.
- Hardt, M., Price, E., and Srebro, N. Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems*, 2016.
- Jones, E., Oliphant, T., Peterson, P., et al. SciPy: Open source scientific tools for Python, 2001. URL <http://www.scipy.org/>.
- Kamiran, F., Calders, T., and Pechenizkiy, M. Discrimination aware decision tree learning. In *2010 IEEE International Conference on Data Mining*, 2010. doi: 10.1109/ICDM.2010.50.
- Kleinberg, J., Mullainathan, S., and Raghavan, M. Inherent trade-offs in the fair determination of risk scores. In *Proceedings of the 8th Innovations in Theoretical Computer Science Conference*, 2017. doi: 10.4230/LIPIcs.ITCS.2017.43.
- Kraft, D. A software package for sequential quadratic programming. Technical Report DFVLR-FB 88-28, Institut für Dynamik der Flugsysteme, Deutsche Forschungs- und Versuchsanstalt für Luft- und Raumfahrt (DFVLR), 1988.
- Kraft, D. Algorithm 733: TOMP—Fortran modules for optimal control calculations. *ACM Transactions on Mathematical Software*, 20(3), 1994. doi: 10.1145/192115.192124.
- Liu, L. T., Simchowitz, M., and Hardt, M. The implicit fairness criterion of unconstrained learning. In *Proceedings of the 36th International Conference on Machine Learning*, 2019.
- Madras, D., Creager, E., Pitassi, T., and Zemel, R. Learning adversarially fair and transferable representations. In *Proceedings of the 35th International Conference on Machine Learning*, 2018.
- Menon, A. K. and Williamson, R. C. The cost of fairness in binary classification. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, 2018.
- Narayanan, A. Translation tutorial: 21 fairness definitions and their politics. In *Proceedings of the Conference on Fairness, Accountability and Transparency, New York, USA*, 2018.
- Pleiss, G., Raghavan, M., Wu, F., Kleinberg, J., and Weinberger, K. Q. On fairness and calibration. In *Advances in Neural Information Processing Systems*, 2017.
- Samadi, S., Tantipongpipat, U., Morgenstern, J. H., Singh, M., and Vempala, S. The price of fair PCA: One extra dimension. In *Advances in Neural Information Processing Systems*, 2018.
- Song, J., Kalluri, P., Grover, A., Zhao, S., and Ermon, S. Learning controllable fair representations. In *Proceedings of Machine Learning Research*, 2019.
- Tan, Z., Yeom, S., Fredrikson, M., and Talwalkar, A. Learning fair representations for kernel models. In *Proceedings of The 23rd International Conference on Artificial Intelligence and Statistics*, 2020.
- Verma, S. and Rubin, J. Fairness definitions explained. In *Proceedings of the International Workshop on Software Fairness*, 2018. doi: 10.1145/3194770.3194776.

Žliobaitė, I. On the relation between accuracy and fairness in binary classification. In *Proceedings of the Second Workshop on Fairness, Accountability, and Transparency in Machine Learning*, 2015.

Zafar, M. B., Valera, I., Rodriguez, M. G., and Gummadi, K. P. Fairness constraints: Mechanisms for fair classification. In *Proceedings of the Second Workshop on Fairness, Accountability, and Transparency in Machine Learning*, 2015.

Zafar, M. B., Valera, I., Gomez Rodriguez, M., and Gummadi, K. P. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In *Proceedings of the 26th International Conference on World Wide Web*, 2017. doi: 10.1145/3038912.3052660.

Zemel, R., Wu, Y., Swersky, K., Pitassi, T., and Dwork, C. Learning fair representations. In *Proceedings of the 30th International Conference on Machine Learning*, 2013.

Zhao, H. and Gordon, G. J. Inherent tradeoffs in learning fair representations. In *Advances in Neural Information Processing Systems*, 2019.