

# Fairness Implications of Encoding Protected Categorical Attributes

CARLOS MOUGAN, University of Southampton, United Kingdom

JOSE M. ALVAREZ, Scuola Normale Superiore, Italy and University of Pisa, Italy

SALVATORE RUGGIERI, University of Pisa, Italy

STEFFEN STAAB, University of Southampton, United Kingdom and University of Stuttgart, Germany

Past research has demonstrated that the explicit use of protected attributes in machine learning can improve both performance and fairness. Many machine learning algorithms, however, cannot directly process categorical attributes, such as country of birth or ethnicity. Because protected attributes frequently are categorical, they must be encoded as features that can be input to a chosen machine learning algorithm, e.g. support vector machines, gradient boosting decision trees or linear models. Thereby, encoding methods influence how and what the machine learning algorithm will learn, affecting model performance and fairness. This work compares the accuracy and fairness implications of the two most well-known encoding methods: *one-hot encoding* and *target encoding*. We distinguish between two types of induced bias that may arise from these encoding methods and may lead to unfair models. The first type, *irreducible bias*, is due to direct group category discrimination and the second type, *reducible bias*, is due to the large variance in statistically underrepresented groups. We investigate the interaction between categorical encodings and target encoding regularization methods that reduce unfairness. Furthermore, we consider the problem of intersectional unfairness that may arise when machine learning best practices improve performance measures by encoding several categorical attributes into a high-cardinality feature.

CCS Concepts: • **Computing methodologies** → **Supervised learning by classification**; **Classification and regression trees**; **Supervised learning by regression**; • **Social and professional topics** → **Socio-technical systems**.

Additional Key Words and Phrases: Fairness, Algorithmic Accountability, Categorical Features, Bias

## ACM Reference Format:

Carlos Mougán, Jose M. Alvarez, Salvatore Ruggieri, and Steffen Staab. 2023. Fairness Implications of Encoding Protected Categorical Attributes. In *AAAI/ACM Conference on AI, Ethics, and Society (AIES '23)*, August 8–10, 2023, Montréal, QC, Canada. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3600211.3604657>

## 1 INTRODUCTION

Anti-discrimination laws [3, 19, 20] prohibit the unfair treatment of individuals based on *sensitive attributes* (also referred to as *protected attributes*). The list of sensitive attributes varies per country, though these usually include gender, ethnicity, and religion [62]. Following such legal motivations along with societal expectations, many studies have looked into discrimination in machine learning and proposed various ways to promote fairness (e.g., [22, 44, 56, 65]).

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

*AIES '23*, August 8–10, 2023, Montréal, QC, Canada

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0231-0/23/08.

<https://doi.org/10.1145/3600211.3604657>

Handling sensitive attributes throughout the machine learning pipeline is central to establishing fairness. An early common practice was removing data on sensitive attributes altogether. This technique has been questioned because sensitive attributes may be required for avoiding discrimination in data-driven decision models [37, 83]. Therefore, later work [30, 45, 79, 80] has aimed at how to obtain fairer models given the presence of sensitive attributes, formalizing the problem as an optimization trade-off between model quality in terms of performance and some fairness objective.

Sensitive attributes often come as categorical data. For instance, roughly 75% of the famous COMPAS dataset [35] consists of categorical attributes, including most of the sensitive ones (see Section 5.1.1 for more details). Many machine learning algorithms require categorical attributes to be suitably *encoded* as numerical data. Different ways of encoding categorical attributes into numerical features [32, 46, 49] have been proposed and extensively studied in the literature along with statistical regularization methods since the mid-1950s [54]. This has resulted in various methods that encode categorical attributes as numerical data to make them usable by popular machine learning models, such as support vector machines, gradient-boosting decision trees, or linear models.

In this paper, we study the broader implications that encoding categorical sensitive attributes can have on model accuracy and fairness. Despite being a common machine learning practice, often within the data pre-processing step, the effects of categorical attribute encodings on fairness remain largely unexplored. For instance, prior works on fair machine learning [30, 79, 80] also use the encoding of protected categorical attributes without discussing its implications or the choices of encodings.

We focus on the two most widely used encodings: *one-hot encoding* and *target encoding* [49, 54, 58, 73]. *One-hot encoding*, an unsupervised technique, produces orthogonal and equidistant vectors for each category [49, 52], thereby considering the categories to be equally independent of each other and other attributes. However, when dealing with high cardinality categorical variables, one-hot encoding suffers from a lack of scalability and sparsity issues due to the creation of many orthogonal dimensions (later discussed in Section 2 and 3). *Target encoding* [46, 49, 54] is a supervised technique that replaces a categorical attribute with the mean target value of each corresponding category. Thus, it can handle all categories together in one dimension.<sup>1</sup>

<sup>1</sup>Target encoding methods have become an industry standard for high-cardinality categorical data [28, 49, 54, 60] with algorithmic procedures being implemented in many open source packages [14, 73]. One of the most common is the Python package called *category encoders* ([https://pepy.tech/project/category\\_encoders](https://pepy.tech/project/category_encoders)). It achieves up to 1 million downloads per month. Target encoding is the default encoding method in some high-performance open-source software implementations such as *catboost* [17, 29] that has reached a total download of 74 million.

The first problem of unfairness related to sensitive categorical attributes, which we call *irreducible bias*, is associated with the statistical differences between two highly populated groups: more data about the compared groups will not diminish this type of bias. The second problem arises because sampling from small groups may exhibit large variance leading to unfairness constituting *reducible bias*. Many datasets contain distributions of data that are imbalanced over different values of categorical attributes, often leading to performance degradation of the learned models (known as the *classes imbalance problem* [34]). Using encodings in such datasets may naturally introduce disparity per the observed class imbalance.

Moreover, when a dataset contains several sensitive categorical attributes, and these are merged to become one feature (a strategy often followed to improve model quality substantially [28]), encodings may create fine-grained, sparsely populated intersectional features [39, 40, 72] increasing the chance for both types of induced biases [23].

The effects of encodings on model quality and fairness under the interplay of different encoding and regularization techniques have not been studied in the literature. However, they affect very commonly used machine learning practices. For target encoding, we study two popular statistical regularization methods called *smoothing* and *Gaussian noise regularization*. These both regularization provide new avenues for analyzing the implications of categorical encodings on fairness. Through both a theoretical analysis as well as an empirical analysis using two real-world datasets, we find that suitable regularization can address unfairness arising from the target encodings with only marginal losses in accuracy.<sup>2</sup>

In summary, we make the following contributions:

- We compare the best-known categorical feature encoding methods, one-hot encoding, and target encoding against learning without protected attribute(s) in terms of model performance and fairness.
- We study the relationship between the regularization of target encodings and fairness by evaluating smoothing and Gaussian noise, two common techniques used for regularization by data preprocessing.
- We provide evidence that creating intersectional features can worsen discrimination. We show that a regularized target encoder can retain the benefits of intersectional features without increasing unfair discrimination.
- We provide a theoretical analysis studying two types of induced biases, irreducible and reducible, that arise while encoding categorical protected attributes.

## 2 BACKGROUND

### 2.1 Categorical attribute encoding

Handling categorical attributes is a common problem in machine learning, given that many algorithms use numerical data [49, 69]. There are many well-known methods for approaching this problem [4, 8, 9, 53, 70].

<sup>2</sup>In this work, we use the term “accuracy” in a non-literal sense to refer to the model performance, rather than the statistical evaluation metric of the same name. Section 5.1.3 discusses the appropriate evaluation metric for our scenarios.

*One-hot encoding* (also known as dummy variables in the social sciences [74]) constructs orthogonal and equidistant vectors for each category. Given high cardinality categorical attributes, one-hot encoding suffers from shortcomings: (i) the dimension of the input space increases with the cardinality of the encoded variable, (ii) the derived features are rarely non-zero, and new and unseen categories cannot be handled [49, 68].

*Label/ordinal encoding* [6] uses a range of integers to represent different categorical values. These are assumed to have no true order and integers are assigned in the order of appearance of the categories. Label encoding suffers less from higher cardinalities of attribute values, but imposes an artificial random order on the categories, which may harm learning. This, in turn, obstructs the model to extract meaningful information from the categorical data.

Table 1. An illustrative example of one-hot and target encoding methods over the same data sample.

Ethnic	Encoding	Label
African-American	1	1
Caucasian	1/3	1
Caucasian	1/3	0
Caucasian	1/3	0
Hispanic	0	0

(a) Unregularized Mean Target Encoding

Ethnic	African-American	Caucasian	Hispanic
African-American	1	0	0
Caucasian	0	1	0
Caucasian	0	1	0
Caucasian	0	1	0
Hispanic	0	0	1

(b) One Hot Encoding

*Target encoding* replaces attribute categories by the mean target value<sup>3</sup> of each corresponding category. Thus, the high cardinality problem is addressed, and categories are ordered in a meaningful manner [8, 46]. The main drawback of target encoding appears when the target values of a category with few samples are averaged. The model may overly rely on the resulting target value, potentially suffering from inherent variance in the small sample of data points from this category. To overcome this problem, several strategies introduce regularization terms in the target estimation [46, 49, 54].

In Table 1, we illustrate one-hot encoding and target encoding for the category ethnicity using a five person sample from the COMPAS dataset [35]. The problem of over-fitting is evident for the cases of *African-American* and *Hispanic* where their encoding is replaced directly with the target, creating a data leakage that can potentially cause reducible induced bias (cf. Section 5).

<sup>3</sup>Throughout the paper, we assume a binary target feature with values {0, 1}.

Even though early works that have studied preprocessing techniques for classification without discrimination [37] do not discuss the fairness effects of encoding categorical protected attributes. To the best of our knowledge, no previous work studies the different effects of regularization on target encodings nor the fairness implications of encoding categorical protected attributes.

## 2.2 Group Fairness

Various definitions of fairness in machine learning have been proposed (see, e.g., [2, 22, 44] for recent overviews). They can be categorized into notions of individual fairness and group fairness. While metrics of individual fairness judge whether *similar individuals are treated similarly* [18], metrics of group fairness measure the disparate treatment of groups, which are assembled according to shared categories of sensitive attributes of their individuals such as gender or race [64].

Different disparity metrics emphasize varying aspects of disparate treatment. For comprehensive understanding, we investigate the effects of encoding methods according to three common disparity metrics, i.e. equal opportunity, statistical disparity (demographic disparity) and average absolute odds (equalized odds). All three metrics indicate equal treatments of different groups by values close to zero and highly disparate treatments by values different from zero.

In this work, we distinguish and define two types of induced bias or discrimination that the encoding of categorical attributes introduces. Borrowing terminology used about different types of uncertainty [15, 26, 41], we use *irreducible bias* to refer to (direct) group discrimination arising from the categorization of groups into labels: more data about the compared groups do not reduce this type of bias. *Reducible bias* occurs due when the variance of categories with few instances cannot be well contained.

## 2.3 Addressing Intersectionality

It is a common trick for boosting model performance to concatenate multiple categorical variables and encode them into a single feature [28]. This feature engineering procedure, which includes target encoding, parallels a possible implementation of *intersectionality* when we concatenated two or more protected attributes. Intersectionality refers to when an individual that belongs to more than one protected group experiences discrimination at the intersection of these groups. It broadly refers to how different identities interact to produce a unique new form of discrimination [72]. Crenshaw [13], for example, studied how *black women* in the United States experience discrimination beyond being either black or women.

Although individuals often belong to multiple protected groups, intersectionality is largely understudied within algorithmic fairness. With some exceptions (e.g., [1, 24, 72, 76, 77]), most works assume the single binary protected attribute or disregard intersectionality entirely when handling multiple protected attributes [72], which is unrealistic and reductive. This is a pertinent issue as it is possible for individuals not to suffer from multiple discrimination but to suffer from intersectional discrimination [63, 75].

We address the intersectionality concerns linked to target encoding. On one hand, it can boost model performance; on the other hand,

it can introduce new forms of discrimination. We add to this small but growing fairness literature by analyzing how target encoding can enable an implementation of intersectionality. In particular, we study how target encoding regularization can mitigate the potential biases induced by this feature engineering practice and compare it to the standard alternative of one-hot encoding.

## 3 FORMALIZATION AND REGULARIZATION OF TARGET ENCODING

Consider a categorical attribute  $Z$  with domain  $dom(Z) = \{z_1, \dots, z_c\}$ , a binary target attribute  $Y$  with  $dom(Y) = \{0, 1\}$ , and the joint probability of  $P(Z, Y)$  over the population of interest. Target encoding replaces  $Z$  with a continuous attribute  $\tilde{Z}$  with  $dom(\tilde{Z}) \in [0, 1]$ . Values  $z_i \in dom(Z)$ , for  $i = 1, \dots, c$ , are encoded to values  $\tilde{z}_i$  in a supervised way, as the posterior probability of positives:

$$\tilde{z}_i = p_i \quad \text{where } p_i = P(Y = 1|Z = z_i) \quad (1)$$

However, since  $P(\cdot)$  is typically unknown, an estimate of the posterior probability  $p_i$  is derived from a dataset  $\mathcal{D}_{tr}$  (called the *training set*) of i.i.d. realizations of  $Z, Y$ . Let  $n$  be the total number of observations,  $n_i$  the number of observations where  $Z = z_i$ , and  $n_Y$  the number of observations where  $Y = 1$ , and  $n_{i,Y}$  the number of observations where  $Z = z_i$  and  $Y = 1$ . A candidate estimator consists of the observed fraction of positives among those with  $Z = z_i$ , hence encoding:

$$\tilde{z}_i = \hat{p}_i \quad \text{where } \hat{p}_i = \frac{n_{i,Y}}{n_i} \quad (2)$$

Such an estimator is unbiased, namely  $E[\hat{p}_i] = p_i = P(Y = 1|Z = z_i)$ . More precisely, by Hoeffding bounds [33], for any  $\epsilon > 0$ ,  $P(|\hat{p}_i - p_i| \geq \epsilon) \leq 2e^{-2n_i\epsilon^2}$ , which already points out the dependence of the estimate on the number of observations  $n_i$  of  $z_i$ . Formally, the variance of the estimator  $Var[\hat{p}_i] = p_i(1 - p_i)/n_i$  is relatively large when  $n_i$  is small. Unregularized target encoding does not perform well on categories with little statistical mass [58] as it tends to overfit the training data, failing to generalize to new data. In the extreme case of only one observation, namely  $n_i = 1$ , it will replace the categorical value with the target of such an observation. Such an encoding will be unrepresentative of the category and introduces a sampling (or data collection) bias at the pre-processing stage. This type of bias is what we define as reducible bias and can be left unnoticed because extremely small categories do not significantly impact the overall loss of the problem but can still impact fairness metrics. To avoid overfitting, practitioners regularize using either (i) smoothing towards the global mean or (ii) Gaussian noise, which adds normal (Gaussian) distribution noise to training data to decrease overfitting. Other smoothing techniques can be found in the literature but are either minimal variations of those two techniques or less popular [73].

### 3.1 Smoothing regularization

Smoothing towards the global mean leads to the following target encoding:

$$\tilde{z}_i = \tilde{p}_i \quad \text{where } \tilde{p}_i = \lambda(n_i) \frac{n_{i,Y}}{n_i} + (1 - \lambda(n_i)) \frac{n_Y}{n} \quad (3)$$

Here, the proportion of positives among the observations with  $Z = z_i$  is interpolated with the proportion of positives among all observations. Formally, called  $\hat{p} = n_Y/n$  an estimate of the prior probability  $p = P(Y = 1)$ , we have  $\hat{p}_i = \lambda(n_i)\hat{p}_i + (1 - \lambda(n_i))\hat{p}$ . The choice of the prior probability  $P(Y = 1)$  is natural because, lacking a sufficient number of observations for  $Z = z_i$ , one resorts to the proportion of positives over the whole dataset of observations. The convex combination of the two estimators depends on  $\lambda(n_i) \in [0, 1]$ . The function  $\lambda(\cdot)$  is assumed to increase with  $n_i$ . Intuitively, the larger the number of observations with  $Z = z_i$ , the more weight we give to the first estimator. Thus, the smoothed estimator is asymptotically unbiased. Conversely, the smaller the number of observations, the more weight we give to the prior probability estimator. Therefore, the smoothed estimator has a small variance for small values of  $n_i$ —yet, it is biased towards the prior probability.

### 3.2 Gaussian noise regularization

Gaussian noise regularization adds normal (Gaussian) distribution noise into training data after encoding the categorical attribute as in (2). The intuition is to perturb the data to prevent overfitting the target encoded attribute values. During the prediction stage, testing data are encoded as in (2) with no perturbation. Formally, called  $z_{i,j}$  the  $j^{\text{th}}$  occurrence of  $z_i$  in the training set,  $z_{i,j}$  is replaced by:

$$\bar{z}_{i,j} = \hat{p}_{i,j} \quad \text{where } \hat{p}_{i,j} = \frac{n_{i,Y}}{n_i} + \epsilon_{i,j} \quad \epsilon_{i,j} \sim N(0, \lambda^2) \quad (4)$$

where the  $\epsilon_{i,j}$ 's are i.i.d. with mean 0 and standard deviation  $\lambda$ . Typical values for  $\lambda$  are set between 0.05 and 0.6 [73].

## 4 THEORETICAL ANALYSIS

We present a theoretical analysis under a number of assumptions that make it reasonably simple. First, we assume that  $\bar{Z}$  is the only predictive feature. Second, we consider a probabilistic binary classifier, which for an input  $\bar{Z} = \bar{z}$  outputs a score  $\hat{S}(\bar{z}) \in [0, 1]$ , and a prediction  $\hat{Y}(\bar{z}) = \mathbb{1}(\hat{S}(\bar{z}) > 1/2)$ . Third, the score is expected to approximate a Bayes optimal classifier, i.e.,  $\hat{S}(\bar{z}) \approx P(Y = 1 | \bar{Z} = \bar{z})$ . For notational convenience, we write  $a \succ b$  as a shorthand for  $a > 1/2 \Leftrightarrow b > 1/2$ , namely  $a$  and  $b$  are on the same side of the decision threshold  $1/2$ . We write  $a \not\succeq b$  when  $a \succ b$  does not hold.

**The case of perfect target encoding.** Under the (theoretical) assumption of knowing the true values  $p_i$ 's, the perfect target encoding would set  $\bar{z}_i = p_i$  as in (1). The score  $\hat{S}(\bar{z}_i) = p_i$  leads to the Bayes optimal classifier, hence maximizing AUC over the population and minimizing the classification error to the following:

$$\sum_{i=1}^c P(Z = z_i) \cdot \min\{p_i, 1 - p_i\} \quad (5)$$

Consider now the equal opportunity fairness metric, namely:

$$P(\hat{Y} = 1 | Y = 1, \bar{Z} = \bar{z}_i) - P(\hat{Y} = 1 | Y = 1, \bar{Z} = \bar{z}_r) \quad (6)$$

where  $\bar{z}_r$  is the encoding of the reference group in the protected attribute  $Z$ . By definition of  $\hat{Y}$ ,  $\hat{Y}(\bar{z}_i) = 1$  iff  $\bar{z}_i = p_i > 1/2$ , and analogously for  $r$ . Therefore, when both  $p_i > 1/2$  and  $p_r > 1/2$ :

$$P(\hat{Y} = 1 | Y = 1, \bar{Z} = \bar{z}_i) = P(\hat{Y} = 1 | Y = 1, \bar{Z} = \bar{z}_r) = 1$$

and then the difference is 0. A similar conclusion is obtained when both  $p_i \leq 1/2$  and  $p_r \leq 1/2$ . However, when the probabilities  $p_i$  and  $p_r$  lie on different sides of the threshold (i.e.,  $p_r \not\succeq p_i$ ), the equal opportunity metrics is non-zero (either  $-1$  or  $1$ ). In other words, the classifier is fair only if the prediction for the reference group is the same as for the protected group. But this will impact on accuracy. In fact, assuming a constant prediction over the groups, say  $\hat{Y}(\bar{z}_i) = 1$ , the classification error on the population becomes  $\sum_{i=1}^c P(Z = z_i) \cdot (1 - p_i)$ , which is clearly larger than (5).

In summary, even in the case of perfect target encoding and a Bayes optimal classifier, there is a tension between error and fairness metrics optimization: the amount of unfairness is *irreducible* as we assumed to know the posterior probabilities  $p_i$ 's, unless we admit increasing the error by not using the protected feature  $Z$  in the classification problem.

**The case of target encoding.** Let us consider now the encoding using the (un-regularized) estimator  $\hat{p}_i = n_{i,Y}/n_i$ , i.e., (2). The score  $\hat{S}(\bar{z}_i) = \hat{p}_i$  maximizes empirical AUC and minimizes the empirical error rate on the training set. When  $n_i$  is large,  $\hat{p}_i \approx p_i$  (since variance of the estimator is low), and then the contribution to the classification error (5) and to the AUC are approximately the same as in the case of perfect target encoding. Regarding the fairness metric, we can reasonably assume that  $n_r$  is large for the reference group, and then  $\hat{p}_r \approx p_r$ . Therefore, the equal opportunity metric is unchanged w.r.t. the case of perfect target encoding.

When  $n_i$  is small, the estimate  $\hat{p}_i = n_{i,Y}/n_i$  can be arbitrarily distant from  $p_i$ . The increment in classification error (5) is zero if  $p_i \succ \hat{p}_i$ , and it is  $P(Z = z_i) \cdot |1 - 2p_i|$  otherwise. Also, the AUC will possibly be smaller due to wrong ranking of instances with  $Z = z_i$ . The equal opportunity metric is, instead, independent of  $P(Z = z_i)$ . Compared to the perfect target encoding case, its value is unchanged if  $p_i \succ \hat{p}_i$ . Otherwise, it can either decrease (if  $p_r \succ \hat{p}_i$ ) or increase (if  $p_r \not\succeq \hat{p}_i$ ).

In summary, the variability of the estimator  $\hat{p}_i$  for  $n_i$  small, negatively impacts on the performance metrics, and it propagates to the the fairness metrics, unpredictably increasing or decreasing it compared to the perfect target encoding case. The increase in the fairness metrics is *reducible* bias, which can be corrected by increasing the number of observations of  $Z = z_i$ .

**The case of smoothing regularization.** Let us consider now the target encoding with smoothing regularization (3). Let  $\hat{S}(\cdot)$  be the score function that minimizes the empirical error rate over the training set. When  $n_i$  is large, then  $\hat{p}_i \approx \hat{p}_i \approx p_i$ , and then we fall back to the same situation as for (perfect) target encoding.

When  $n_i$  is small, we have  $\hat{p}_i \approx n_Y/n \approx p$ , and then instances of the training set for which  $Z = z_i$  are mapped close to  $\bar{Z} = p$ . This does not necessarily mean that the classification algorithm scores such instances as  $p$ —rather, it should score close to the mean target value of instances with  $\bar{Z} = p$ . Let us then be  $q$  such that  $\hat{S}(p) = q$ . We fall back then to the reasoning for the target encoding case. The increment in classification error (5) is zero if  $p_i \succ q$ , and  $P(Z = z_i) \cdot |1 - 2p_i|$  otherwise. Compared to the perfect target encoding case, the fairness metric value is unchanged if  $p_i \succ q$ . Otherwise, it can either decrease (if  $p_r \succ q$ ) or increase (if  $p_r \not\succeq q$ ).

In summary, the estimator  $\tilde{p}_i \approx p$  for  $n_i$  small is stable, but nevertheless, it can affect the performance metrics (negatively) and the fairness metrics (increase or decrease). The increase in the fairness metrics is *reducible* bias. Notice that the magnitude of the impact depends on the choice of  $q$  by the machine learning algorithm under consideration, which, in principle, could be controlled for.

**The case of Gaussian noise regularization.** Let us now consider the Gaussian noise regularization (4). Its expectation is  $E[\hat{p}_{i,j}] = E[\hat{p}_i] + E[\epsilon_{i,j}] = p_i$ , hence the estimator is unbiased. Its variance is  $Var[\hat{p}_{i,j}] = Var[\hat{p}_i] + \lambda^2$ . From this, we have that: (1) the variance is larger than in the case of target encoding, and, a fortiori, of the smoothing regularization; (2) the larger the regularization parameter  $\lambda$ , the larger the variance. Let us consider a partition of the instances with  $Z = z_i$  based on whether  $\hat{p}_{i,j} \succ p_i$  holds or not.

For the subset  $\hat{p}_{i,j} \succ p_i$ , there is no change in classification error, nor in the equal opportunity fairness metrics, when compared to the perfect target encoding case.

Consider instead the subset  $\hat{p}_{i,j} \not\succeq p_i$ . The increment in classification error (5) is  $\sum_{\bar{z}} P(Z = z_i, \bar{Z} = \bar{z}, \bar{z} \not\succeq p_i) \cdot |1 - 2p_i|$ . For  $n_i$  small, this is lower than in the cases of target encoding and smoothing regularization. For  $n_i$  large, this is greater than in those two cases, where it is  $\approx 0$ . However, since  $Var[\hat{p}_i] \approx 0$ , this case only occurs for a large  $\lambda^2$  that causes crossing the decision boundary, i.e., for which  $\hat{p}_{i,j} \not\succeq p_i$ . Compared to the perfect target encoding case, the fairness metric can either decrease (if  $p_r \succ \hat{p}_{i,j}$ ) or increase (if  $p_r \not\succeq \hat{p}_{i,j}$ ). Again, for small  $n_i$ 's the impact is smaller than for target encoding and smoothing regularization, and for large  $n_i$ 's, this can only occur if  $\lambda^2$  is large enough for crossing the decision boundary.

In summary, Gaussian noise regularization adds some controllable variability that impacts mainly on small  $n_i$ 's and for a subset of the data distribution for which a random perturbation may cross the decision boundary. If this happen, there is an increase in classification error, and some chance to increase/decrease the equal opportunity fairness metric. The increase in the fairness metrics is *reducible* bias.

**The case of one-hot encoding.** Consider a variant of one-hot encoding setting  $\bar{z}_i = 2^i$ , i.e., mapping  $z_i$  into a binary number with the  $i$ -th digit set to 1 and all others set to 0. Such a variant keeps our assumption of one predictive feature only. The previous subsections on perfect target encoding and on target encoding could be repeated, almost unchanged, as they only require  $\hat{S}(\bar{z}_i) = p_i$  and  $\hat{S}(\bar{z}_i) = \hat{p}_i$  respectively, ignoring the form of the coding of  $\bar{z}_i$ . We would therefore expect that the behavior of one-hot encoding and (unregularized) target encoding be very similar. What can make a difference is that most machine learning algorithms treat one-hot encoding as a collection of i.i.d. features, ignoring their dependencies (i.e., that one and only one digit must be 1). This may lead to a greater classification error when compared to target encoding.

## 5 EXPERIMENTS

In this section, we study the implications of model accuracy and fairness when encoding categorical protected attributes. (H1) The first main hypothesis is that encoding the protected attribute helps to improve accuracy. (H2) The second main hypothesis is that fairness is

worsened by encoding. To evaluate both (H1) and (H2) we compare two encoding methods, one-hot encoding and target encoding, versus not encoding the protected attribute. Our third hypothesis (H3) is that target encoding regularization can improve fairness without significantly impacting predictive performance, and we evaluate this by comparing two regularization techniques across various hyperparameters as part of the machine learning pipeline's preprocessing step. Additionally, in the last section, we explore the effects of intersectional protected categorical attributes, which augment the previous three hypotheses.

### 5.1 Experimental Setup

**5.1.1 Datasets: COMPAS and FolkTables.** We choose two datasets that happen to exhibit high-cardinality sensitive categorical attributes in a binary classification problem: COMPAS [35] and FolkTables [16]. We report our method and findings on the COMPAS dataset in the main body of this paper and apply the same methodology on FolkTables, but report findings from the latter in the appendix. Overall, the findings are very similar in both datasets.

COMPAS is an acronym for Correctional Offender Management Profiling for Alternative Sanctions, which is an assistive software and support tool used to predict the risk that a criminal defendant will re-offend. The dataset provides a category-based evaluation labelled as high risk of recidivism, medium risk of recidivism, or low risk of recidivism. We convert this multi-class classification problem into binary classification by combining the medium risk and high risk of recidivism and comparing them to low risk of recidivism. The input used for the prediction of recidivism consists of 11 categorical attributes, including gender, custody status, legal status, assessment reason, agency, language, ethnicity, and marital status. The sensitive attribute that we consider is *Ethnic* for the single discrimination case, whose protected group we define as the most represented group: African-American (cf. Figure 4).

To study fairness related to intersectional attributes, we created the variable *EthnicMarital*, engineered by concatenating *Ethnic* and *Marital* status. This new attribute has a high cardinality of 46 distinct values (cf. Figure 4). The most predominant category is *African-American Single*, and it will be the protected group (cf. Figure 4) for the intersectional fairness case. To compare disparate treatment between groups we will make use of *Caucasian Married* as the reference group. It is worth noting that the contribution of the attributes to the model performance, based on attribute importance explanation mechanism [43, 47, 51, 59], is highly relevant. The available data is split into a 50/50 stratified train/test split, maintaining the ratio of each category between train and test set. In the Figure 4 of the appendix, we can see how the group distributions are unbalanced with two groups, *African-American* and *Caucasian*, that account for the +80% of the data. For the intersectional fairness case, the number of groups increases, making room for more distinct, disparate, and imbalanced groups [23].

**5.1.2 Machine learning algorithms.** Our experiments involve a logistic regression model, a neural network (Multi-layer Perceptron classifier), and a gradient-boosting decision tree. All models are trained on the training set. These three models provide examples of a model with large bias (the linear regression model), a highly

complex model (the MLP classifier), and the extensively used, state-of-the-art gradient-boosting decision tree [7, 25, 48, 61, 82].

### 5.1.3 Choice of metrics and models.

*Model performance metrics.* Previous work on fair machine learning has evaluated their experiments on COMPAS using accuracy as a performance metric [78–80], but given that we want to study effects of group imbalance, we consider accuracy to be a less informative measure of model performance. Area Under the Curve (AUC) measures the diagnostic ability of a binary classifier as its discrimination threshold is varied. AUC is less susceptible to class imbalance than accuracy or precision and also accepts soft probabilistic predictions [32]. An AUC of 0.5 is equal to random predictions.

*Fairness metrics.* We use three different metrics  $\ell_{i,j}(f, X, y)$  to judge fairness of classifier  $f$  on data  $X$  between groups indexed by  $i, r$  and we denote  $\hat{Y} = f(X)$  for simplicity:

- *Statistical Parity (Strong Demographic Parity):* The difference between favourable outcomes received by the protected group and reference group [12, 21, 38, 81]. DP ensures that a fair decision does not depend on the protected attribute regardless of the classification threshold used [11, 36]

$$DP_{i,r} = d(P(\hat{Y}|Z = i), P(\hat{Y}|Z = r)) \quad (7)$$

where  $d(\cdot, \cdot)$  is a distance function. In this work, we use the Wasserstein distance as a measure between the two probabilistic distributions. The intuition behind Demographic Parity is that it states that the proportion of each segment of a protected attribute should receive a positive outcome at equal rates, a positive outcome is a preferred decision.

- *Equal opportunity fairness.* Following Hardt et al. [31]’s emphasis on ensuring fair opportunity instead of raw outcomes, we choose *equal opportunity* (EO) as a fairness notion and use the metric *disparate treatment* (difference between the true positive rates) to measure unfairness, which is estimated using the disparate treatment metric [78]. For simplicity, we refer to the interplay of these concepts as the *equal opportunity fairness (EOF)* metric. The value is the difference in the True Positive Rate (TPR) between the protected group and the reference group [50, 57]).

$$TPR_i = P(\hat{Y} = 1|Y = 1, Z = i) \quad EOF_{i,r} = TPR_i - TPR_r \quad (8)$$

A negative value in (8) is due to the worse ability of a Machine Learning model to find actual recidivists for the protected group (i) in comparison with the reference group (j).

- *Average Absolute Odds (Equalized Odds):* The sum of the absolute differences between the True Positive Rates and the False Positive Rates of the protected group plus the same ratio for the reference group.

$$FPR_i = P(\hat{Y} = 1|Y = 0, Z = i) \quad (9)$$

$$AAO_{i,r} = \frac{1}{2}(|FPR_i - FPR_r| + |TPR_i - TPR_r|) \quad (10)$$

The intuition is that an  $AAO = 0$  means the algorithm is fair because it results in the same False Positive Rate and True

Positive Rate for the reference group as a protected group. If the algorithm causes a difference in either, then  $AAO \neq 0$ . A deviation in each term contributes equally to AAO, then False Positives Rates might have different social implications than True Positives Rates [44, 66, 71].

All three metrics indicate better fairness between groups  $i, j$  by values closer to 0. We calculate the overall fairness  $\mathcal{L}$  of the model  $f$  on data of interest  $X$  given a fairness metrics  $\ell$ , reference group  $i$  and other groups  $\{i|i \neq r\}$  as:

$$\mathcal{L}(f, X, y, i, r) = \sum_{i \neq r} |\ell_{i,r}(f, X, y)| \quad (11)$$

where each group  $i$  contributes equally to the overall metric, meaning these are not weighted by the number of individuals in each group.

## 5.2 Experimental results: encoding categorical protected attributes

In this section, we evaluate hypotheses (H1), (H2), and (H3). The trade-offs between fairness metrics and predictive performance metrics (AUC) are analyzed using two different encoding techniques (Section 2), with two different regularization techniques (Section 3) and two different estimators (Section 5.1.3). The ranges of the regularization hyperparameters are:  $\lambda \in [0, 5]$  for the width of the Gaussian noise regularization;  $m \in [0, 1000000]$  for the additive smoothing using the  $m$ -probability estimate function  $\lambda(n_i) = n_i/(n_i + m)$  (see [46]). These hyperparameters will also be kept for the rest of the experiments for the COMPAS dataset.

Under **Gaussian noise regularization** (cf Figure 1 left images), evaluation supports our three hypotheses: (H1) predictive performance improves when encoding the categorical protected attributes. In all six experiments, the improvements reported are in the range of  $\sim 0.1$  AUC. (H2) All the experiments exhibit fairness degradation up to one order of magnitude. (H3) We observe that within low regularization ranges of hyperparameters (lighter dots), fairness improves without compromising the predictive performance of the model. However, for higher levels of regularization (darker dots), fairness metrics have a plateau while predictive performance (AUC) keeps degrading. At the highest regularization penalty, target encoding often matches performance and fairness with “no encoding” while with no regularization matches “one hot encoding”. Later in this section, we discuss this in depth.

We find similar results in the case of **smoothing regularization** (cf Figure 1 right images). But not for our regularization hypothesis. While it should be for the linear regression and the neural networks, it does not work for the gradient-boosting decision trees, whose target encoding regularization effects are negligible in both fairness and model performance. These can be due to smoothing producing a shrinking effect where decision tree-based models are generally not affected by monotonic attribute transformations [10].

In Figure 2, we analyze the target encoding hyperparameter fairness-accuracy trade-off deeper. We can see that there is an optimal trade-off value around 0.3, where the equal opportunity fairness and demographic parity have dropped down toward the fairness plateau, and the model performance has only slightly decreased. The

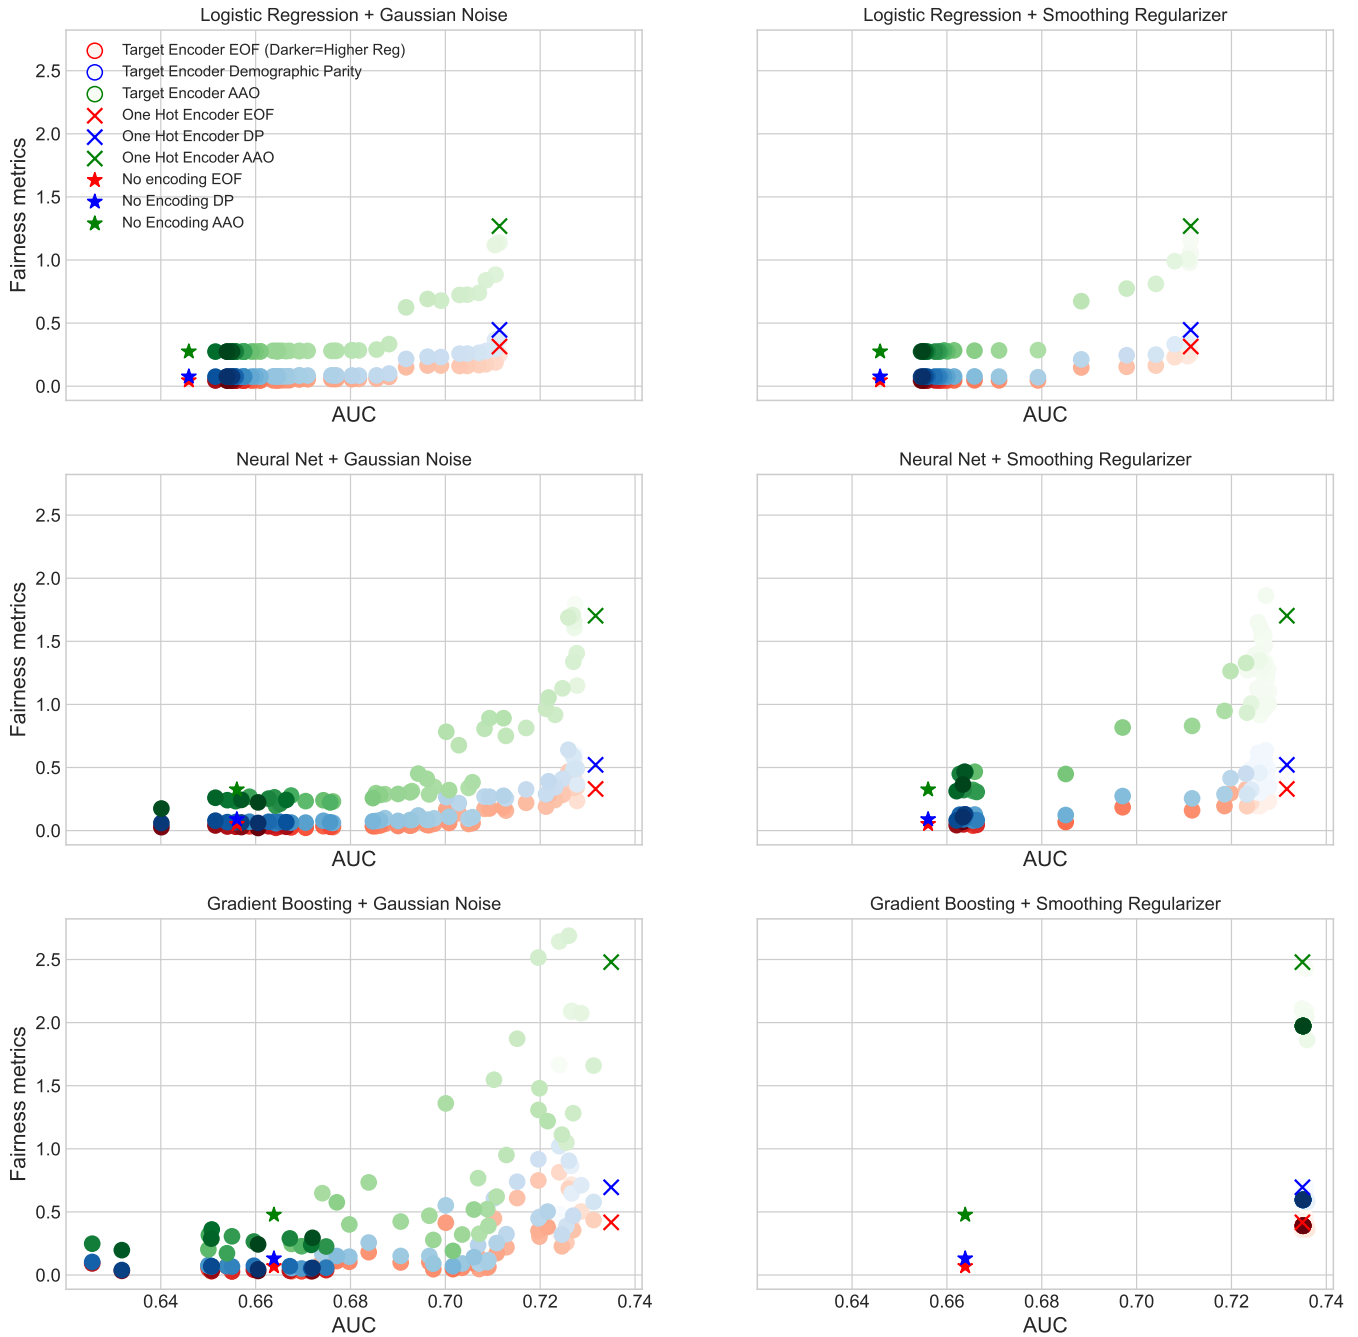


Fig. 1. Comparing one-hot encoding and target encoding regularization (Gaussian noise and smoothing) for the Logistic Regression, Neural Network, and Gradient Boosting classifiers over the test set of the COMPAS dataset. Coloured dots regard different regularization parameters: the darker the red, the higher the regularization. Different colours imply different fairness metrics. Crossed dots regard one-hot encoding, and starred dots are the results of models that exclude the use of the protected attribute.

predictive performances (AUC) of different groups have different negative slopes, ethnic groups as *Asian* or *Native-American* have a drastic drop in performance while groups as *African-American*

have only a small performance decay. *African-American* represents the 44.4% of the data while *Asian* or *Native-American* do not even achieve a statistical representation of 1%.

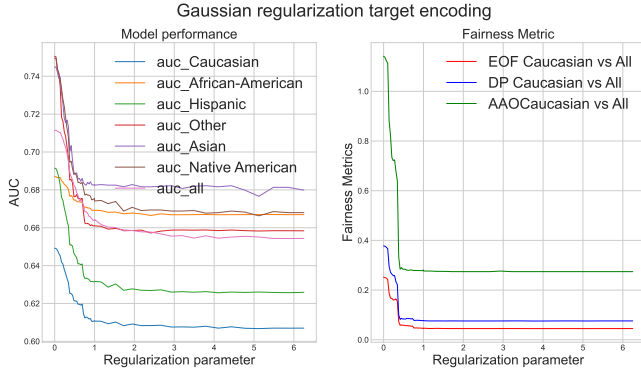


Fig. 2. Impact of the Gaussian noise regularization parameter  $\lambda$  on performance and fairness metrics over the test set of the COMPAS dataset using a Logistic Regression with L1 penalty. In the left image, the AUC of all the protected groups over the regularization hyperparameter. On the right, the equal opportunity fairness, demographic parity and average absolute odds variation throughout the regularization hyperparameter.

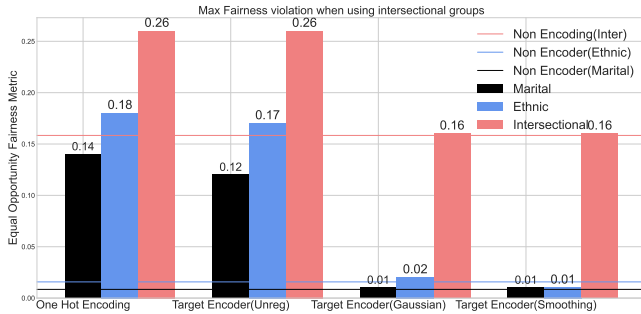


Fig. 3. Equal opportunity fairness implications of encoding categorical protected attributes and their regularization effects on the Compas Dataset. Horizontal lines are the baselines where the protected attribute is not included in the training data. Regularized target encoding does not harm fairness metrics, but it can improve predictive performance on this dataset.

Concerning the across-model comparison, the predictive performance of the gradient boosting decision tree model is best, followed by the neural network and then the linear model [5, 27]. From the fairness perspective, more complex models have a stronger fairness violation.

### 5.3 Experimental results: engineering intersectional fairness

Our intersectional fairness hypothesis are that (i) the engineering of intersectional features degrades fairness, (ii) that encoding the categorical protected attribute increases discrimination and (iii) that by regularizing target encoding, we can reduce intersectional discrimination to no-encoding levels.

To provide evidence of the potential effects of encodings on intersectional fairness, we concatenate *Ethnic* and *Marital* status of the COMPAS dataset. We select *Caucasian Married* as the reference group and compare the maximum fairness violation w.r.t. all other

groups. For visualization purposes, we choose the generalized linear model or the previous section and focus on the notion of Equal Opportunity Fairness since we have seen in the previous experimental section that the three fairness metrics exhibit the same behavior.

In Figure 3, we see how attribute concatenation creates intersectional attributes and boosts fairness violations. Validating our first hypothesis that fairness metrics increase just by the engineering of intersectional discrimination. Even when there is no-encoding the protected attributes (horizontal lines), the maximum fairness violation between groups is increased by an order of magnitude from 0.015 for *Ethnic* or 0.08 for *Marital Status* to 0.16 for the intersectional attribute of both. The increase of discrimination when engineering intersectional protected attributes align with the social findings presented originally back in 1958 when Kimberle Crenshaw [13] wrote her critique of the anti-discrimination doctrine, feminist theory, and anti-racist politics, to describe how different forms of oppression intersect and compound one another, increased discrimination for marginalized groups.

Our second hypothesis is validated as both encoding techniques achieve a higher equal opportunity violation than no-encoding of the protected attribute. Finally, we can see that fairness can be improved by regularizing the target encoding of protected attributes. This is not surprising, and, in general, attribute concatenation can worsen fairness both on the side of irreducible bias (because  $p_i$  and  $p_r$  become more distant) and on the side of reducible bias (because  $n_i$  becomes smaller) as we have seen in the theoretical section.

## 6 CONCLUSION

In this work, we have focused on how the encoding of categorical attributes can reconcile model quality and fairness. We have provided theoretical and empirical evidence that encoding categorical attributes could induce two different types of bias: an *irreducible bias*, due to the learning of discriminant information between the protected and reference groups, and a *reducible bias* due to the large variance of samples found in small protected groups.

Through theory and experiments, we showed that the most used categorical encoding method in the fair machine learning literature, one-hot encoding, consistently discriminates more than target encoding. However, we found some promising results using target encoding. Target encoding regularization showed fairness improvements with the risk of a noticeable loss of model performance in the case of over-parametrization. We also found that the type of regularization chosen is relevant depending on the algorithm used. These results support our view that (regularized) target encoding can be useful for fair machine learning. Furthermore, we discussed how attribute engineering could boost the performance of machine learning algorithms but can lead to fairness violations increase, potentially due to both reducible and irreducible biases.

These experiments aim to motivate industry practitioners, where in many situations, the usage of the protected attribute is not strictly prohibited, and with slight changes in the encoding of the protected attribute, improvements in fairness can be achieved without any noticeable detriment to predictive performance.

**Limitations and disclaimer:** In this work, we have used two models, two encodings, two regularization techniques, and two



datasets. To make a large-scale comparison, we must choose a single scalar metric that accounts for the trade-off between model accuracy and model fairness. Also, encodings are more impactful when the protected attribute is related to the target variable. This work aims to show what are some of the implications of encoding categorical protected attributes. At all times, it is important to understand that simply encoding categorical protected attributes may not necessarily lead to improved fairness metrics. We strongly advocate considering the effects of encoding regularization not only on predictive performance but also along the fairness axis. Using fair AI methods does not necessarily guarantee the fairness of AI-based complex socio-technical systems [42, 64, 67].

## Reproducibility Statement

We make our results open-source and reproducible: original data, data preparation routines, code repositories, and methods are all publicly available at <https://github.com/nobias-project/FairEncoding>. Note that throughout our work, we do not perform any hyperparameter tuning (except on the regularization); instead, we use default `scikit-learn` hyperparameters [55]. Our experiments were run on a four vCPU server with 15GB of RAM.

## ACKNOWLEDGMENTS

This work has received funding from the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie Actions (grant agreement number 860630) for the project: “NoBIAS - Artificial Intelligence without Bias”. Furthermore, this work reflects only the author’s view, and the European Research Executive Agency (REA) is not responsible for any use that may be made of the information it contains. We acknowledge the support of the Stuttgart Research Focus *Interchange Forum for Reflection on Intelligent Systems (IRIS)*. The authors thank Gourab K. Patro for his early-stage contributions.

## REFERENCES

- José M. Álvarez and Salvatore Ruggieri. 2023. Counterfactual Situation Testing: Uncovering Discrimination under Fairness given the Difference. *CoRR* abs/2302.11944 (2023).
- Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2019. *Fairness and Machine Learning*. fairmlbook.org. <http://www.fairmlbook.org>.
- Solon Barocas and Andrew D Selbst. 2016. Big data’s disparate impact. *California Law Review* 104, 3 (2016), 671–732.
- Rodrigo Kraus Barragán. 2022. Tratamiento de variables categóricas en modelos de Machine Learning. . (2022).
- Vadim Borisov, Tobias Leemann, Kathrin Seßler, Johannes Haug, Martin Pawelczyk, and Gjergji Kasneci. 2021. Deep Neural Networks and Tabular Data: A Survey. <https://doi.org/10.48550/ARXIV.2110.01889>
- Gregory Carey. 2003. Coding Categorical Variables. <http://psych.colorado.edu/~carey/Courses/PSYC5741/handouts/Coding%20Categorical%20Variables%202006-03-03.pdf>
- Rich Caruana and Alexandru Niculescu-Mizil. 2006. An empirical comparison of supervised learning algorithms. In *Machine Learning, Proceedings of the Twenty-Third International Conference (ICML 2006), Pittsburgh, Pennsylvania, USA, June 25-29, 2006 (ACM International Conference Proceeding Series, Vol. 148)*, William W. Cohen and Andrew W. Moore (Eds.). ACM, 161–168. <https://doi.org/10.1145/1143844.1143865>
- Patricio Cerda and Gaël Varoquaux. 2022. Encoding High-Cardinality String Categorical Variables. *IEEE Trans. Knowl. Data Eng.* 34, 3 (2022), 1164–1176. <https://doi.org/10.1109/TKDE.2020.2992529>
- Patricio Cerda, Gaël Varoquaux, and Balázs Kégl. 2018. Similarity encoding for learning with dirty categorical variables. *Mach. Learn.* 107, 8-10 (2018), 1477–1494. <https://doi.org/10.1007/s10994-018-5724-2>
- Tianqi Chen and Carlos Guestrin. 2016. XGBoost: A Scalable Tree Boosting System. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016*, Balaji Krishnapuram, Mohak Shah, Alexander J. Smola, Charu C. Aggarwal, Dou Shen, and Rajeev Rastogi (Eds.). ACM, 785–794. <https://doi.org/10.1145/2939672.2939785>
- Silvia Chiappa, Ray Jiang, Tom Stepleton, Aldo Pacchiano, Heinrich Jiang, and John Aslanides. 2020. A General Approach to Fairness with Optimal Transport. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*. AAAI Press, 3633–3640. <https://ojs.aaai.org/index.php/AAAI/article/view/5771>
- Sam Corbett-Davies, Emma Pierson, Avi Feller, Sharad Goel, and Aziz Huq. 2017. Algorithmic Decision Making and the Cost of Fairness. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Halifax, NS, Canada, August 13 - 17, 2017*. ACM, 797–806. <https://doi.org/10.1145/3097983.3098095>
- Kimberlé Crenshaw. 2013. Demarginalizing the intersection of race and sex: A black feminist critique of antidiscrimination doctrine, feminist theory and antiracist politics. In *Feminist legal theories*. Routledge, 23–51.
- David Masip, Carlos Mougán. 2020. Sktools:tools to extend sklearn, feature engineering based transformers. <https://sktools.readthedocs.io/> [Online; accessed 20-August-2022].
- Armen Der Kiureghian and Ove Ditlevsen. 2009. Aleatory or Epistemic? Does It Matter? *Structural Safety* 31 (2009), 105–112. <https://doi.org/10.1016/j.strusafe.2008.06.020>
- Frances Ding, Moritz Hardt, John Miller, and Ludwig Schmidt. 2021. Retiring Adult: New Datasets for Fair Machine Learning. *CoRR* abs/2108.04884 (2021). [arXiv:2108.04884](https://arxiv.org/abs/2108.04884) <https://arxiv.org/abs/2108.04884>
- Anna Veronika Dorogush, Vasily Ershov, and Andrey Gulin. 2018. CatBoost: gradient boosting with categorical features support. *ArXiv preprint* abs/1810.11363 (2018). <https://arxiv.org/abs/1810.11363>
- Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard S. Zemel. 2012. Fairness through awareness. In *Innovations in Theoretical Computer Science 2012, Cambridge, MA, USA, January 8-10, 2012*, Shafi Goldwasser (Ed.). ACM, 214–226. <https://doi.org/10.1145/2090236.2090255>
- European Commission. 2012. CHARTER OF FUNDAMENTAL RIGHTS OF THE EUROPEAN UNION. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:12012P/TXT>
- European Commission. 2019. 2018 Reform of EU data protection rules. [https://ec.europa.eu/commission/sites/beta-political/files/data-protection-factsheet-changes\\_en.pdf](https://ec.europa.eu/commission/sites/beta-political/files/data-protection-factsheet-changes_en.pdf)
- Michael Feldman, Sorelle A. Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. 2015. Certifying and Removing Disparate Impact. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Sydney, NSW, Australia, August 10-13, 2015*, Longbing Cao, Chengqi Zhang, Thorsten Joachims, Geoffrey I. Webb, Dragos D. Margineantu, and Graham Williams (Eds.). ACM, 259–268. <https://doi.org/10.1145/2783258.2783311>
- Jessie Finocchiaro, Roland Maio, Faidra Monachou, Gourab K Patro, Manish Raghavan, Ana-Andreea Stoica, and Stratis Tsirtsis. 2021. Bridging Machine Learning and Mechanism Design towards Algorithmic Fairness. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 489–503.
- James R. Foulds, Rashidul Islam, Kamrun Naher Keya, and Shimei Pan. 2020. Bayesian Modeling of Intersectional Fairness: The Variance of Bias. In *Proceedings of the 2020 SIAM International Conference on Data Mining, SDM 2020, Cincinnati, Ohio, USA, May 7-9, 2020*, Carlotta Demeniconi and Nitesh V. Chawla (Eds.). SIAM, 424–432. <https://doi.org/10.1137/1.9781611976236.48>
- James R. Foulds, Rashidul Islam, Kamrun Naher Keya, and Shimei Pan. 2020. An Intersectional Definition of Fairness. In *36th IEEE International Conference on Data Engineering, ICDE 2020, Dallas, TX, USA, April 20-24, 2020*. IEEE, 1918–1921. <https://doi.org/10.1109/ICDE48307.2020.00203>
- Jerome H. Friedman. 2001. Greedy function approximation: A gradient boosting machine. *The Annals of Statistics* 29, 5 (2001), 1189 – 1232.
- Yarin Gal. 2016. *Uncertainty in Deep Learning*. Ph.D. Dissertation. University of Cambridge.
- Léo Grinsztajn, Edouard Oyallon, and Gaël Varoquaux. 2022. Why do tree-based models still outperform deep learning on tabular data? (July 2022). <https://hal.archives-ouvertes.fr/hal-03723551> working paper or preprint.
- Alexander Guschin, Dmitry Ulyanov, Mikhail Trofimov, Dmitry Altukhov, and Mario Michaidilidis. 2018. How to Win a Data Science Competition: Learn from Top Kagglers - National Research University Higher School of Economics. <https://www.coursera.org/lecture/competitive-data-science/categorical-and-ordinal-features-qu1TF>. Accessed 02/11/20.
- John T Hancock and Taghi M Khoshgoffaar. 2020. CatBoost for big data: an interdisciplinary review. *Journal of big data* 7, 1 (2020), 1–45.

- [30] Moritz Hardt, Eric Price, and Nati Srebro. 2016. Equality of Opportunity in Supervised Learning. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, Daniel D. Lee, Masashi Sugiyama, Ulrike von Luxburg, Isabelle Guyon, and Roman Garnett (Eds.). 3315–3323. <https://proceedings.neurips.cc/paper/2016/hash/9d2682367c3935defcb1f9e247a97c0d-Abstract.html>
- [31] Moritz Hardt, Eric Price, and Nati Srebro. 2016. Equality of Opportunity in Supervised Learning. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, Daniel D. Lee, Masashi Sugiyama, Ulrike von Luxburg, Isabelle Guyon, and Roman Garnett (Eds.). 3315–3323. <https://proceedings.neurips.cc/paper/2016/hash/9d2682367c3935defcb1f9e247a97c0d-Abstract.html>
- [32] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. 2001. *The Elements of Statistical Learning*. Springer New York Inc., New York, NY, USA.
- [33] Wassily Hoeffding. 1963. Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc.* 58 (1963), 13–30. Issue 301. <https://doi.org/10.1080/01621459.1963.10500830>
- [34] Nathalie Japkowicz and Shaju Stephen. 2002. The class imbalance problem: A systematic study. *Intelligent data analysis* 6, 5 (2002), 429–449.
- [35] Lauren Kirchner Jeff Larson, Surya Mattu and Julia Angwin. 2016. How We Analyzed the COMPAS Recidivism Algorithm. <https://github.com/propublica/compas-analysis> [Online; accessed 21-December-2021].
- [36] Ray Jiang, Aldo Pacchiano, Tom Stepleton, Heinrich Jiang, and Silvia Chiappa. 2019. Wasserstein Fair Classification. In *Proceedings of the Thirty-Fifth Conference on Uncertainty in Artificial Intelligence, UAI 2019, Tel Aviv, Israel, July 22-25, 2019 (Proceedings of Machine Learning Research, Vol. 115)*, Amir Globerson and Ricardo Silva (Eds.). AUAI Press, 862–872. <http://proceedings.mlr.press/v115/jiang20a.html>
- [37] Faisal Kamiran and Toon Calders. 2011. Data preprocessing techniques for classification without discrimination. *Knowl. Inf. Syst.* 33, 1 (2011), 1–33. <https://doi.org/10.1007/s10115-011-0463-8>
- [38] Toshihiro Kamishima, Shotaro Akaho, Hideki Asoh, and Jun Sakuma. 2012. Fairness-aware classifier with prejudice remover regularizer. In *Joint European conference on machine learning and knowledge discovery in databases*. Springer, 35–50.
- [39] Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. 2018. An Empirical Study of Rich Subgroup Fairness for Machine Learning. arXiv:1808.08166 [cs.LG]
- [40] Michael J. Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. 2018. Preventing Fairness Gerrymandering: Auditing and Learning for Subgroup Fairness. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018 (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer G. Dy and Andreas Krause (Eds.). PMLR, 2569–2577. <http://proceedings.mlr.press/v80/kearns18a.html>
- [41] Armen Der Kiureghian and Ove Ditlevsen. 2009. Aleatory or epistemic? Does it matter? *Structural Safety* 31, 2 (2009), 105–112. <https://doi.org/10.1016/j.strusafe.2008.06.020> Risk Acceptance and Risk Communication.
- [42] Bogdan Kulynych, Rebekah Overdorf, Carmela Troncoso, and Seda F. Gürses. 2020. POTs: protective optimization technologies. In *FAT\* '20: Conference on Fairness, Accountability, and Transparency, Barcelona, Spain, January 27-30, 2020*, Mireille Hildebrandt, Carlos Castillo, L. Elisa Celis, Salvatore Ruggieri, Linnet Taylor, and Gabriela Zanfir-Fortuna (Eds.). ACM, 177–188. <https://doi.org/10.1145/3351095.3372853>
- [43] Scott M. Lundberg, Gabriel Erion, Hugh Chen, Alex DeGrave, Jordan M. Prutkin, Bala Nair, Ronit Katz, Jonathan Himmelfarb, Nisha Bansal, and Su-In Lee. 2019. Explainable AI for Trees: From Local Explanations to Global Understanding. arXiv:1905.04610 [cs.LG]
- [44] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2022. A Survey on Bias and Fairness in Machine Learning. *ACM Comput. Surv.* 54, 6 (2022), 115:1–115:35. <https://doi.org/10.1145/3457607>
- [45] Aditya Krishna Menon and Robert C. Williamson. 2018. The cost of fairness in binary classification. In *Conference on Fairness, Accountability and Transparency, FAT 2018, 23-24 February 2018, New York, NY, USA (Proceedings of Machine Learning Research, Vol. 81)*, Sorelle A. Friedler and Christo Wilson (Eds.). PMLR, 107–118. <http://proceedings.mlr.press/v81/menon18a.html>
- [46] Daniele Micci-Barreca. 2001. A Preprocessing Scheme for High-Cardinality Categorical Attributes in Classification and Prediction Problems. *SIGKDD Explor. Newsl.* 3, 1 (2001), 27–32.
- [47] Christoph Molnar. 2019. *Interpretable Machine Learning*. <https://christophm.github.io/interpretable-ml-book/>.
- [48] Carlos Mougan, Qian Fu, Joiceena Kolath, Huan Tong, Siddharth Dixit, Laurens Geffert, Hyesop Shin, Rabuh, Ahmad Abd, Caitlin Robinson, and Ella Gale. 2020. *Data Study Group Network Final Report: Bristol City Council (Get Bristol moving: tackling air pollution in Bristol city centre)*. Zenodo. <https://doi.org/10.5281/zenodo.3775497> Turing Network Data Study Group Bristol ; Conference date: 05-08-2019 Through 09-08-2019.
- [49] Carlos Mougan, David Masip, Jordi Nin, and Oriol Pujol. 2021. Quantile Encoder: Tackling High Cardinality Categorical Features in Regression Problems. In *Modeling Decisions for Artificial Intelligence - 18th International Conference, MDAI 2021, Umeå, Sweden, September 27-30, 2021, Proceedings (Lecture Notes in Computer Science, Vol. 12898)*, Vicenç Torra and Yasuo Narukawa (Eds.). Springer, 168–180. [https://doi.org/10.1007/978-3-030-85529-1\\_14](https://doi.org/10.1007/978-3-030-85529-1_14)
- [50] Cecilia Munoz, Megan Smith, and DJ Patil. 2016. Big data: A report on algorithmic systems, opportunity, and civil right. *United States. Executive Office of the President* (2016). <https://www.hsdl.org/?view&did=792977>
- [51] Carlos Mougan Navarro, Georgios Kanellos, and Thomas Gottron. 2021. Desiderata for Explainable AI in Statistical Production Systems of the European Central Bank. In *Machine Learning and Principles and Practice of Knowledge Discovery in Databases*. Springer International Publishing, Cham, 575–590.
- [52] Shuntaro Okada, Masayuki Ohzeki, and Shinichiro Taguchi. 2019. Efficient partition of integer optimization problems with one-hot encoding. *Scientific Reports* 9 (2019). <https://doi.org/10.1038/s41598-019-49539-6>
- [53] F. Pargent, B. Bischl, and J. Thomas. 2019. *A benchmark experiment on how to encode categorical features in predictive modeling*. Master's thesis. School of Statistics.
- [54] Florian Pargent, Florian Pfisterer, Janek Thomas, and Bernd Bischl. 2022. Regularized target encoding outperforms traditional methods in supervised machine learning with high cardinality features. *Comput. Stat.* 37, 5 (2022), 2671–2692.
- [55] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [56] Dino Pedreschi, Salvatore Ruggieri, and Franco Turini. 2008. Discrimination-aware data mining. In *KDD*. ACM, 560–568.
- [57] John Podesta, Penny Pritzker, Ernest J. Moniz, John Holden, and Jeffrey Zients. 2014. Big data: Seizing opportunities and preserving values. *United States. Executive Office of the President* (2014). <https://www.hsdl.org/?view&did=752636>
- [58] Liudmila Ostroumova Prokhorenkova, Gleb Gusev, Aleksandr Vorobev, Anna Veronika Dorogush, and Andrey Gulin. 2018. CatBoost: unbiased boosting with categorical features. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett (Eds.). 6639–6649. <https://proceedings.neurips.cc/paper/2018/hash/14491b756b3a51daac41c24863285549-Abstract.html>
- [59] Marco Túlio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016*, Balaji Krishnapuram, Mohak Shah, Alexander J. Smola, Charu C. Aggarwal, Dou Shen, and Rajeev Rastogi (Eds.). ACM, 1135–1144. <https://doi.org/10.1145/2939672.2939778>
- [60] Pau Rodriguez, Miguel A. Bautista, Jordi González, and Sergio Escalera. 2018. Beyond one-hot encoding: Lower dimensional target embedding. *Image and Vision Computing* 75 (2018), 21–31. <https://doi.org/10.1016/j.imavis.2018.04.004>
- [61] Byron P. Roe, Hai-Jun Yang, Ji Zhu, Yong Liu, Ion Stancu, and Gordon McGregor. 2005. Boosted decision trees as an alternative to artificial neural networks for particle identification. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 543, 2-3 (2005), 577–584.
- [62] Andrea Romei and Salvatore Ruggieri. 2014. A multidisciplinary survey on discrimination analysis. *Knowl. Eng. Rev.* 29, 5 (2014), 582–638. <https://doi.org/10.1017/S0269888913000039>
- [63] Arjun Roy, Jan Horstmann, and Eirini Ntoutsi. 2023. Multi-dimensional discrimination in Law and Machine Learning - A comparative overview. *CoRR* abs/2302.05995 (2023). <https://doi.org/10.48550/arXiv.2302.05995> arXiv:2302.05995
- [64] S Ruggieri, J. M. Alvarez, A. Pugnana, L. State, and F. Turini. 2023. Can We Trust Fair-AI?. In *The Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023*. AAAI Press.
- [65] Salvatore Ruggieri, Dino Pedreschi, and Franco Turini. 2010. Data mining for discrimination discovery. *ACM Trans. Knowl. Discov. Data* 4, 2 (2010), 9:1–9:40.
- [66] Pedro Saleiro, Benedict Kuester, Abby Stevens, Ari Anisfeld, Loren Hinkson, Jesse London, and Rayid Ghani. 2018. Aequitas: A Bias and Fairness Audit Toolkit. *ArXiv preprint abs/1811.05577* (2018). <https://arxiv.org/abs/1811.05577>
- [67] Kristen M. Scott, Sonja Mei Wang, Milagros Miceli, Pieter Delobel, Karolina Sztandar-Sztanderska, and Bettina Berendt. 2022. Algorithmic Tools in Public Employment Services: Towards a Jobseeker-Centric Perspective. In *FAccT '22: 2022 ACM Conference on Fairness, Accountability, and Transparency, Seoul, Republic of Korea, June 21 - 24, 2022*. ACM, 2138–2148. <https://doi.org/10.1145/3531146.3534631>
- [68] Austin Slakey, Daniel Salas, and Yoni Schamroth. 2019. Encoding Categorical Variables with Conjugate Bayesian Models for WeWork Lead Scoring Engine. *arXiv e-prints*, Article arXiv:1904.13001 (2019), arXiv:1904.13001 pages.

arXiv:1904.13001 [cs.LG]

[69] Gerhard Tutz. 2011. *Regression for Categorical Data*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511842061>

[70] Eric Valdez-Valenzuela, Angel Kuri-Morales, and Helena Gomez-Adorno. 2022. CESAMMO: Categorical Encoding by Statistical Applied Multivariable Modeling. In *Advances in Computational Intelligence*, Obdulia Pichardo Lagunas, Juan Martínez-Miranda, and Bella Martínez Seis (Eds.). Springer Nature Switzerland, Cham, 173–182.

[71] Sandra Wachter, Brent Mittelstadt, and Chris Russell. 2020. Bias preservation in machine learning: the legality of fairness metrics under EU non-discrimination law. *W. Va. L. Rev.* 123 (2020), 735.

[72] Angelina Wang, Vikram V. Ramaswamy, and Olga Russakovsky. 2022. Towards Intersectionality in Machine Learning: Including More Identities, Handling Underrepresentation, and Performing Evaluation. In *FAccT '22: 2022 ACM Conference on Fairness, Accountability, and Transparency*, Seoul, Republic of Korea, June 21–24, 2022. ACM, 336–349. <https://doi.org/10.1145/3531146.3533101>

[73] Will McGinnis. 2020. Category Encoders :A library of sklearn compatible categorical variable encoders. <https://contrib.scikit-learn.org/>

[74] Jeffrey M Wooldridge. 2015. *Introductory Econometrics: A Modern Approach*. Cengage Learning.

[75] Raphaële Xenidis. 2020. Tuning EU equality law to algorithmic discrimination: Three pathways to resilience. *Maastricht Journal of European and Comparative Law* 27, 6 (2020), 736–758.

[76] Ke Yang, Joshua R. Loftus, and Julia Stoyanovich. 2021. Causal Intersectionality and Fair Ranking. In *2nd Symposium on Foundations of Responsible Computing, FORC 2021, June 9–11, 2021, Virtual Conference (LIPIcs, Vol. 192)*, Katrina Ligett and Swati Gupta (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 7:1–7:20. <https://doi.org/10.4230/LIPIcs.FORC.2021.7>

[77] Seyma Yucer, Samet Akçay, Noura Al Moubayed, and Toby P. Breckon. 2020. Exploring Racial Bias within Face Recognition via per-subject Adversarially-Enabled Data Augmentation. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR Workshops 2020, Seattle, WA, USA, June 14–19, 2020*. Computer Vision Foundation / IEEE, 83–92. <https://doi.org/10.1109/CVPRW50498.2020.00017>

[78] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, and Krishna P. Gummadi. 2017. Fairness Beyond Disparate Treatment & Disparate Impact. *Proceedings of the 26th International Conference on World Wide Web (2017)*. <https://doi.org/10.1145/3038912.3052660>

[79] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez-Rodriguez, and Krishna P. Gummadi. 2017. Fairness Constraints: Mechanisms for Fair Classification. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, AISTATS 2017, 20–22 April 2017, Fort Lauderdale, FL, USA (Proceedings of Machine Learning Research, Vol. 54)*, Aarti Singh and Xiaojin (Jerry) Zhu (Eds.). PMLR, 962–970. <http://proceedings.mlr.press/v54/zafar17a.html>

[80] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez-Rodriguez, Krishna P. Gummadi, and Adrian Weller. 2017. From Parity to Preference-based Notions of Fairness in Classification. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4–9, 2017, Long Beach, CA, USA*, Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (Eds.), 229–239. <https://proceedings.neurips.cc/paper/2017/hash/82161242827b703e6acf9c726942a1e4-Abstract.html>

[81] Richard S. Zemel, Yu Wu, Kevin Swersky, Toniann Pitassi, and Cynthia Dwork. 2013. Learning Fair Representations. In *Proceedings of the 30th International Conference on Machine Learning, ICML 2013, Atlanta, GA, USA, 16–21 June 2013 (JMLR Workshop and Conference Proceedings, Vol. 28)*. JMLR.org, 325–333. <http://proceedings.mlr.press/v28/zemel13.html>

[82] Y. Zhang and A. Haghani. 2015. A gradient boosting method to improve travel time prediction. *Transportation Research Part C-emerging Technologies* 58 (2015), 308–324.

[83] Indre Zliobaite and Bart Custers. 2016. Using sensitive personal data may be necessary for avoiding discrimination in data-driven decision models. *Artif. Intell. Law* 24, 2 (2016), 183–201. <https://doi.org/10.1007/s10506-016-9182-5>

APPENDIX: EXPERIMENT RESULTS

Data: Compas data overview

In Figure 4 we complement the experimental section on the main body of the paper by showing the distributions of the ethnic groups. There are two groups (*African-American* and *Caucasian*) that account for the 80% of the data, while there are less represented groups such as Asian or Arabic that have a less significant statistical weight. For the intersectional fairness case, the number of groups is increased

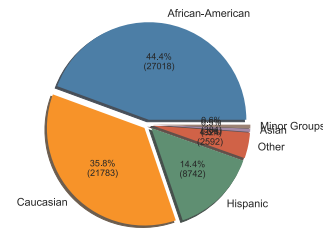


Fig. 4. Distribution of the protected attribute categories to be encoded and regularized for the COMPAS data [35]. Predominant *Ethnic* categories are *African-American* and *Caucasian*

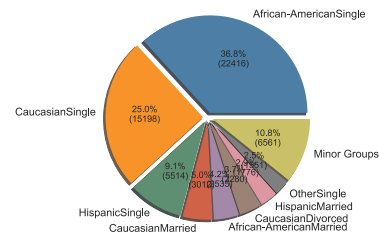


Fig. 5. Distribution of the intersectional protected attribute *Ethnic-Marital* to be encoded and regularized for the COMPAS data [35]. Predominant categories is categories distribution are *African-American Single* and *Caucasian Single*

to 46 distinct groups, making room for more distinct, disparate, and imbalanced groups[23].

Data: US Census Income

In this section, we provide experiments on the Adult Income data set<sup>4</sup> derived from the US census data [16]. Folktables package provides access to data-sets derived from the US Census, facilitating the bench-marking of fair machine learning algorithms. We select the data from California in 2014 that covers 60,729 individuals including their race, that has 8 unique groups. Aiming to predict whether an individual’s income is above 50,000. The data is split into a 50/50 train/test split, maintaining the ratio of each category between the train and test set.

Under **Gaussian noise regularization** (left images of Figure 6), for the logistic regression, we can validate our three hypotheses: (H1) first that predictive performance improves when encoding the categorical protected attributes, in this case, respect to the results on Compass dataset, the AUC improvements are smaller, this can be due to the lack of predictive power of the categorical protected attribute. (H2) that fairness metrics are worsened by the encoding

<sup>4</sup>Please see the ACS PUMS data dictionary for the full list of variables available <https://www.census.gov/programs-surveys/acs/microdata/documentation.html>

	Distribution	Ratio
White	117209	0.66
Asian	28817	0.16
Other	20706	0.11
Black	8435	0.05
Native	1121	0.005
Hawaiian	612	0.003
American Indian	379	0.002

Table 2. Statistical distribution of the protected attribute *Race* on the US census dataset.

versus one-hot encoding or non-regularized target encoding are substantial.

Our last hypothesis (H3) is that through regularisation predictive performance can be improved without compromising the fairness of the model. We can observe that during the low regularization range of hyperparameters (lighter dots), there are high fairness violations with only a small improvement in predictive performance. On the other side, for high regularization (darker dots), fairness metrics have a smaller value. At the highest regularization penalty, target encoding often matches performance and fairness with “no encoding” while with no regularization matches “one hot encoding”.



Fig. 6. Comparing one-hot encoding and target encoding regularization (Gaussian noise and smoothing) for the Logistic Regression classifier over the test set of the US Income dataset. The Reference group is *White*. Coloured dots regard different regularization parameters: the darker the red, the higher the regularization. Different colours imply different fairness metrics. Crossed dots regards one-hot encoding, and starred dots not including the protected attribute in the data.

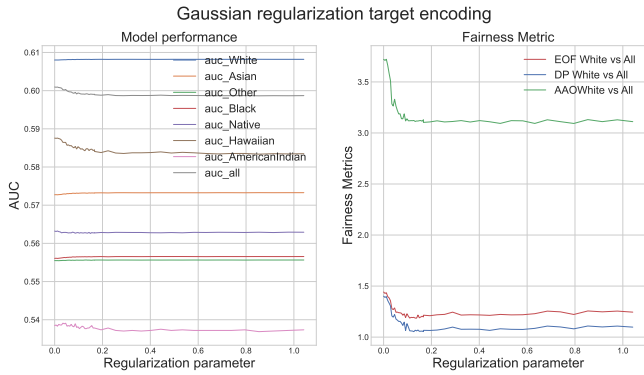


Fig. 7. Impact of the Gaussian noise regularization parameter  $\lambda$  on performance and fairness metrics over the test set of the US income dataset using a Logistic Regression with L1 penalty. In the left image, the AUC of all the protected groups over the regularization hyperparameter. On the right, the equal opportunity fairness, demographic parity and average absolute odds variation throughout the regularization hyperparameter.

of the protected attribute, the differences between no-encoding