LUCID–GAN: Conditional Generative Models to Locate Unfairness[*]

Andres Algaba[1], Carmen Mazijn[1], Carina Prunkl[2], Jan Danckaert[1], and Vincent Ginis[1,3]

[1] Vrije Universiteit Brussel, Pleinlaan 2, 1000 Brussels, Belgium
{andres.algaba,carmen.mazijn,jan.danckaert,vincent.ginis}@vub.be
[2] Oxford University, Wellington Square, OX1 2JD, Oxford, United Kingdom
carina.prunkl@philosophy.ox.ac.uk
[3] Harvard University, 150 Western Ave, Boston, USA

Abstract. Most group fairness notions detect unethical biases by computing statistical parity metrics on a model's output. However, this approach suffers from several shortcomings, such as philosophical disagreement, mutual incompatibility, and lack of interpretability. These shortcomings have spurred the research on complementary bias detection methods that offer additional transparency into the sources of discrimination and are agnostic towards an a priori decision on the definition of fairness and choice of protected features. A recent proposal in this direction is LUCID (Locating Unfairness through Canonical Inverse Design), where canonical sets are generated by performing gradient descent on the input space, revealing a model's desired input given a preferred output. This information about the model's mechanisms, i.e., which feature values are essential to obtain specific outputs, allows exposing potential unethical biases in its internal logic. Here, we present LUCID–GAN, which generates canonical inputs via a conditional generative model instead of gradient–based inverse design. LUCID–GAN has several benefits, including that it applies to non–differentiable models, ensures that canonical sets consist of realistic inputs, and allows to assess proxy and intersectional discrimination. We empirically evaluate LUCID–GAN on the UCI Adult and COMPAS data sets and show that it allows for detecting unethical biases in black–box models without requiring access to the training data.[1]

Keywords: algorithmic fairness · bias detection · discrimination · generative models.

---

[1] Code is available at https://github.com/Integrated-Intelligence-Lab/canonical_sets.

1   Introduction

The increasing use of Artificial Intelligence (AI) algorithms in (semi–)automated decision–making processes has raised concerns about harmful and discriminatory decision patterns observed in contexts such as healthcare [30,43], education [34], and hiring [12,46]. In these cases, the algorithmic decisions discriminate against people based on (legally) protected features, including gender, age, and ethnicity [32]. Often, the very detection of discrimination is difficult because protected characteristics are encoded in so–called proxies. Proxy discrimination is especially prevalent in the era of big data, where algorithms can reconstruct many protected features from non–protected data [7]. In this paper, we address this challenge by developing a fairness evaluation method that reveals an algorithm's desired feature values for a given outcome. Our method exposes proxies that embed potential unethical biases and enhances transparency in the algorithm's decision–making process.

Algorithmic discrimination can be direct or indirect [6]. Indirect discrimination focuses on the impact of a given decision on a protected group. Within US law, this is often labeled "disparate impact." While algorithmic decision–making tools have been shown to frequently put members of particular social groups at a disadvantage [6], indirect discrimination can often be justified as being a proportionate means of achieving a legitimate goal (e.g., in hiring decisions) [1]. Direct discrimination, on the other hand, focuses not so much on the impacts but on the *reasons* behind a given decision. In other words, a person not being hired because of their belonging to a particular social group would constitute a case of direct discrimination. While direct discrimination is illegal both under EU and US law[2], cases of direct discrimination often go unchallenged due to the difficulty of establishing a causal link between protected characteristics and decision outcomes. Within the context of algorithmic decision making, the direct discrimination doctrine is often translated into the requirement to abstain from using protected characteristics as input variables. In practice, however, such attempts of "fairness–through–unawareness" rarely work since protected features are often encoded through other features, giving rise to potential proxy discrimination [18].

Such proxy discrimination further exacerbates the challenge of establishing potential causal links between protected characteristics and less favorable decision outcomes. The identification of proxies and their relationship to protected characteristics is therefore crucial to identify cases of directly discriminating algorithms [1].

While proxy discrimination plays an important role for direct discrimination, it is more often considered in the context of indirect discrimination by algorithms [6,50]. Indirect discrimination has a strong focus on the *outcome* of a given decision (as opposed to the *reason* behind it) and so it more readily connects to

---

[2] Anti-discrimination law in the US speaks of "disparate treatment" which, while similar to direct discrimination, additionally requires there to be discriminatory intent [45].

the fairness literature's large focus on algorithmic outputs. The traditional approach involves translating philosophical or political notions of group fairness into a statistical parity metric on the model's output [34].[3] However, output–based fairness evaluations of this kind have several shortcomings. First, many notions of group fairness are incompatible, except for highly constrained special cases [26]. Second, the problem of many fairness notions only worsens as there is often substantial philosophical disagreement on which ones are genuinely fair in each specific context [8]. Third, by reducing the evaluation to a single number, output–based metrics make it hard to verify the validity of the results and to understand exactly why the model is unfair [37]. Fourth, most output–based fairness evaluations make it difficult to detect intersectional discrimination as they are limited to a specific selection of protected features. This selection entails the risk of missing discrimination against people at the intersection of different protected features [11,27] or against groups that do not share a protected feature [9]. Finally, the computation of these parity metrics often depends on a benchmark data set that may be biased or unbalanced to some extent [42,33]. These shortcomings motivate the research on complementary fairness evaluation methods, which offer additional transparency into the sources of discrimination and are agnostic towards an a priori decision on the choice of protected features, as this is often case-dependent and policy–related [20].

In this paper, we build on the LUCID (Locating Unfairness through Canonical Inverse Design) method proposed by Mazijn et al. [36]. LUCID generates canonical inputs by performing gradient descent on the input space, revealing a trained model's desired input given a preferred output. The resulting canonical set contains valuable information about the model's mechanisms, i.e., which feature values are essential to obtain specific outputs. This allows us to expose potential unethical biases in its internal logic by inspecting the distribution of the protected features. Despite LUCID's appealing properties as a fairness evaluation method, the canonical sets generated by gradient–based inverse design have some critical shortcomings. First, while the canonical sets in their current form are specifically suitable for tabular data, they require differentiable models. LUCID thereby omits the class of tree–based models, which are very effective for tabular data [47]. Additionally, the current gradient–based approach may lead to non–realistic canonical inputs, and it is not straightforward to assess proxy or intersectional discrimination [14].

We present **LUCID–GAN**, which generates canonical inputs via a conditional generative model instead of gradient–based inverse design. LUCID–GAN generates canonical inputs conditional on the predictions of the model under fairness evaluation (see Fig. 1a). Using a conditional generative model has two clear benefits. First, we only require a set of (test) samples and corresponding predictions from the model under fairness evaluation, making it a model–agnostic approach. Second, LUCID–GAN generates realistic samples as defined by its ob-

---

[3] Note that we focus on notions of group fairness which are most commonly used in practice, while many other definitions of fairness exist, including individual and counterfactual fairness [16,28].

jective function. Furthermore, the flexibility of LUCID–GAN is twofold. First, we can extend the canonical inputs with protected features that are not part of the input space of the model under fairness evaluation. Second, the categorical (one–hot encoded) features are part of the conditional vector in the generator, and we can thus condition the canonical inputs on specific feature values (e.g., setting "Male" for "sex"). The first flexibility allows us to assess proxy discrimination, and the second to explicitly check for sources of intersectional discrimination.

LUCID–GAN is an input–based fairness evaluation method which takes a somewhat reverse approach to the statistical output–based metrics. Instead of comparing the predictions, we compare the (protected) feature distributions corresponding to large positive and negative predictions. LUCID–GAN is agnostic towards an a priori decision on the definition of fairness and choice of protected features and instead provides results that suggest potential sources of discrimination [36]. One can then examine the resulting canonical sets from multiple points of view, which are often case–dependent and policy–related [20]. By learning the joint conditional distribution of the features on the model's predictions, LUCID–GAN generates a diverse set of realistic synthetic samples that get positive and negative outputs. Instead of a single number and focusing on a select number of protected features, we get an overview of the overall preferences of the model. We can detect combinations of feature values that often appear together (e.g., "White," "Male," "Married," and "Husband" in the UCI Adult data set), which gives us insights into the different channels of potential proxy discrimination [1].

While LUCID–GAN and output–based metrics may sometimes convey the same results, LUCID–GAN can shed light on the potential sources for the statistical disparities. This is crucial as enforcing statistical parity for the wrong reasons can actually harm the protected groups [13]. LUCID–GAN reveals a model's preferences for a specific output. In contrast, output–based metrics show the statistical disparities on a benchmark data set that may result from many unknown causes. Overall, we argue that both techniques can be part of the same toolbox as they yield different insights.

We provide a brief overview of the literature on algorithmic fairness, activation maximization, and generative models for tabular data in Section 2. In Section 3, we present LUCID–GAN and discuss how it solves LUCID's shortcomings. In Section 4, we show how to generate canonical sets via LUCID–GAN for various fairness evaluations, including direct, proxy, and intersectional discrimination, on the UCI Adult [15] and COMPAS [2] data sets. We find that LUCID–GAN is a valuable addition to the toolbox of algorithmic fairness evaluation, as it offers additional transparency into the sources of discrimination and is agnostic towards an a priori decision on the definition of fairness and choice of protected features.
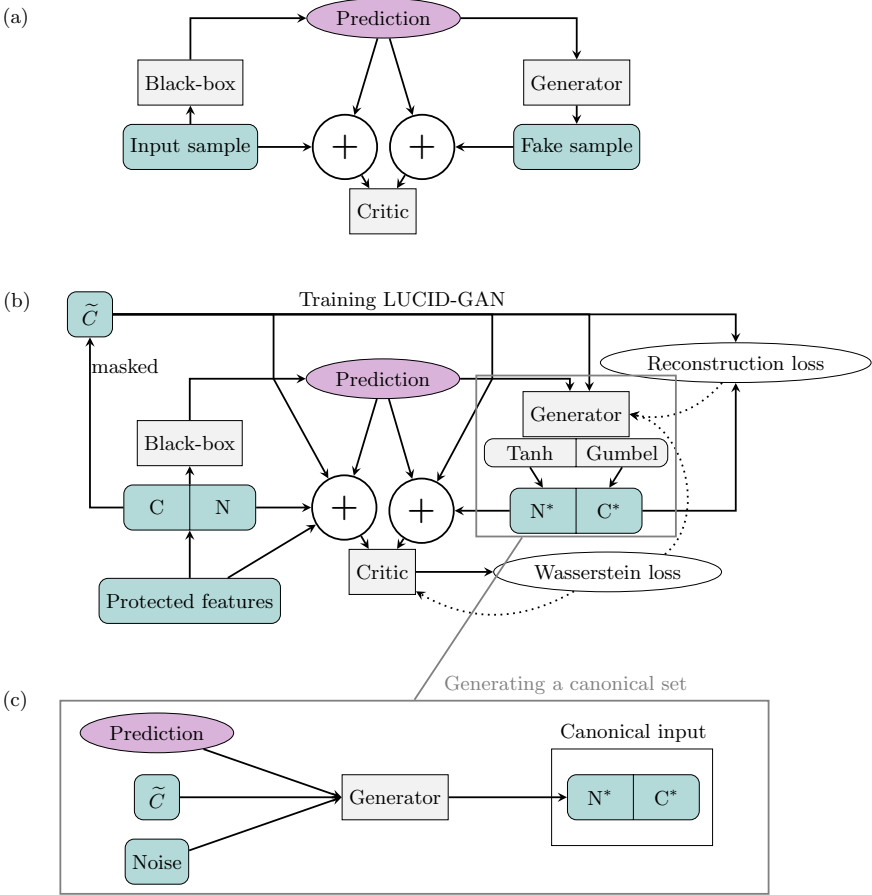
Fig. 1: The LUCID-GAN architecture. (a) High–level overview, (b) LUCID–GAN training mechanism, and (c) the details of the Generator. The input sample may consist of both numerical (N) and one–hot encoded (C) features, which may contain protected features that are not necessarily in the input space of the black–box model. The generator receives random noise, the black–box's predictions, and a masked version of C ($\widetilde{C}$) where only one category of a single one–hot encoded feature is still equal to one. The critic gets real samples from the data and synthetic samples from the generator, where N^* is generated via Tanh functions and C^* via Gumbel–Softmax functions. In addition, it also receives the black–box's predictions and $\widetilde{C}$ to check the validity of the sample. The critic is trained via the Wasserstein loss, while the generator also adds a reconstruction loss on C^* and $\widetilde{C}$. Note that if the protected features are not part of the black–box model's input space, we can still assess proxy discrimination by generating these features (as part of N^* or C^*) and concatenating them to the real samples (as part of N or C) given to the critic. We can use $\widetilde{C}$ to assess cases of intersectional discrimination by generating canonical inputs conditional on fixed values for protected features, such as setting "Male" for "Sex."

2   Background and Related Work

The development of LUCID–GAN to expose unethical biases in a model's internal logic connects ideas from the literature on algorithmic fairness, activation maximization using generative models, and recent advancements in GANs (Generative Adversarial Networks) for tabular data. We discuss these different fields below.

2.1   Algorithmic Fairness

The inherent ambiguity of viewing the concept of fairness spurred the development of many fairness notions, with over 19 widely accepted definitions [34]. Most group fairness notions focus on the equality of outcome by computing statistical parity metrics on a model's output. The two most prominent examples of these statistical output–based metrics are Demographic Parity (DP) and Equality Of Opportunity (EOP) [6]. In DP, we compare the Positivity Rate (PR) of the subpopulations under fairness evaluation, and in EOP, we compare the True Positive Rate (TPR). The choice between DP and EOP entirely depends on the underlying assumptions and worldview of the evaluator [19]. Indeed, even among those two most widely used metrics, substantial philosophical disagreement exists on which one is genuinely fair in each specific context [8]. Moreover, they are incompatible, except for highly constrained special cases [26], and it has been empirically shown that inherent trade–offs exist in many practical situations [35].

In addition, these statistical output–based metrics may suffer from sampling bias and variance as they depend on a benchmark data set [33]. By reducing the fairness evaluation to a single number, it is hard to verify the validity of the results and to understand exactly why the model is unfair. To enhance the transparency of the statistical parity metrics, there is a strong interaction between output–based fairness evaluations and interpretability methods [37]. However, there are many trade–offs and the interpretability methods themselves may suffer from unethical biases [5,25,49]. Finally, there is also the selection of protected attributes, which may lead to missing cases of intersectional discrimination, also known as fairness gerrymandering [24]. We argue that LUCID–GAN counters many of these shortcomings by offering additional transparency into the sources of discrimination and being agnostic towards an a priori decision on the definition of fairness and choice of protected features.

2.2   Activation Maximization

Performing gradient descent on the input space generates canonical inputs that maximize a specific output activation [48]. However, this gradient–based inverse design approach often leads to unrealistic inputs that obtain a high–confidence score on a specific class. These so–called fool inputs result from discriminative models allocating large areas of high confidence, often much larger than the area occupied by training samples for that class [41]. To avoid the generation of fool inputs, Nguyen et al. [39,40] perform gradient descent in the latent space of a

generator network to maximize the output activation in a separate discriminative model. Finally, Odena et al. [44] and Zhou et al. [53] train an Auxiliary Classifier GAN (AC–GAN) and Activation Maximization GAN (AM–GAN) to maximize the activation of its label in an auxiliary classifier.

Two important distinctions exist with our application of LUCID–GAN on fairness and tabular data. First, our goal is to generate realistic canonical inputs for any black–box model's predictions, not to stabilize the training of the GAN. Moreover, note that the AC– and AM–GAN require access to a differentiable model, which we avoid using a Conditional GAN (CGAN) [38]. Second, most previous methods focus on images, where the resulting canonical inputs are individually interpretable and difficult to aggregate [29]. In contrast, we infer the overall fairness of a model from the feature distributions by generating a set of canonical inputs.

2.3   Tabular GAN

The use of GANs (Generative Adversarial Networks) [21] for tabular data concerns specific challenges, such as mixed data types, non–Gaussian and multimodal distributions, and sparse and highly imbalanced categorical features [10]. To deal with mixed data types, i.e., numerical and one–hot encoded features, different types of output activation functions are used in the generator. For the numerical features, the hyperbolic tangent function (Tanh) is combined with a min–max normalization as a pre–processing step. However, this does not address the non–Gaussian and multimodal distributions. Therefore, Xu et al. [51] propose a Conditional Tabular GAN (CTGAN) with mode–specific (min–max) normalization by using a variational Gaussian mixture model. For the categorical features, the Gumbel–Softmax function is used to make the softmax operation in the generator differentiable [23]. To handle the sparse and highly imbalanced categorical features, Xu et al. [51] provide a conditional vector to the generator in the form of a masked version of the one–hot encoded features where only one category of a single one–hot encoded feature is still equal to one, and apply a reconstruction loss to the synthetic sample and the masked vector of the selected one–hot encoded feature. They further propose a training–by–sampling technique where the masked one–hot encoded feature is randomly sampled, and its category is sampled from the empirical log–frequency distribution.

The masked vector can be used to control the generation process similarly to the CGAN framework [38]. For example, [17] use the masked vector to oversample specific categories in an imbalanced learning problem. Zhao et al. [52] also add the components of the variational Gaussian mixture model to the masked vector to accommodate features with strict upper and lower bounds, e.g., an income of exactly zero. We continue the work on CTGAN by extending the conditional vector of the generator in LUCID–GAN with the black–box model's predictions, which are continuous values that are not part of the synthetic sample, and examining its interaction with the masked vector.

3   LUCID–GAN

We present LUCID–GAN to expose unethical biases in a model's internal logic by generating canonical inputs. First, we discuss LUCID and the canonical sets as proposed by Mazijn et al. [36]. Then, we introduce LUCID–GAN and show how to generate canonical sets.

3.1   LUCID

Shortcomings in the existing output–based group fairness metrics have spurred the research on complementary fairness evaluation methods, which offer additional transparency into the sources of discrimination and are agnostic towards an a priori decision on the definition of fairness and choice of protected features, as this is often case–dependent and policy–related [20]. In this spirit, Mazijn et al. [36] propose LUCID, and introduce the notion of a canonical set that allows to evaluate the fairness of a model's decision-making processes. Through gradient-based inverse design, LUCID generates canonical inputs, which can be considered the desired input given a preferred output for a trained model. By repeatedly generating canonical inputs, the resulting canonical set reveals which feature values are essential to obtain specific outputs. This allows for exposing potential unethical biases in the model's internal logic by inspecting the distribution of the protected features. In contrast to output metrics, there is no need for a specific fairness metric, a ground truth, or a benchmark data set.

Following Mazijn et al. [36], we use LUCID to generate a canonical set for a trained binary classifier.[4] First, we draw an extensive set of randomly initialized input samples from a uniform distribution. Then, we transform these random input samples into canonical inputs through gradient–based inverse design. Each subsequent transformation results from minimizing the (cross–entropy) loss between the model prediction and the preferred output (e.g., a loan is granted) until the model's prediction is close to its maximum (a predicted probability of 1 in this case). Note that we keep the model parameters fixed throughout the entire procedure. Finally, we inspect the distribution of each protected feature within the canonical set and compare it to the initial random distribution. Several design considerations impact the resulting canonical set, such as the initialization, choice of hyperparameters, and pre– and post–processing of categorical features. Besides these practical considerations, LUCID has some critical shortcomings, such as the requirement of differentiable models, unrealistic canonical inputs, and difficulties in assessing proxy and intersectional discrimination, which we solve by introducing LUCID–GAN.

---

[4] We use the same implementation and default values for the hyperparameters as LUCID. For more details see Mazijn et al. [36]. The code for LUCID is available at: `https://github.com/Integrated-Intelligence-Lab/canonical_sets`.

3.2   Specifications of LUCID–GAN

LUCID–GAN builds upon the CTGAN [51] framework. In the class of GANs, a generative model is trained through an adversarial process with a critic model (see Fig. 1a). The former aims to create synthetic samples from random noise that fool the latter into judging them as real. Besides the random noise, we provide a conditional vector to the generator, which allows us to control the generation process to some extent [38]. The CTGAN uses a Wasserstein GAN with gradient penalty (WGAN–GP) to prevent common problems in training GAN models, such as mode–dropping, vanishing gradients, and non–convergence [3,4,22]. Xu et al. [51] further propose stabilizing training by following PacGAN [31] and packaging multiple samples together in the critic. Both the generator and the critic consist of fully–connected hidden layers. In the generator, we use batch normalization, ReLU functions, and residual connections on each hidden layer. In the critic, we use leaky ReLU functions and dropout.[5]

3.3   Training LUCID–GAN

The generator in LUCID–GAN generates canonical inputs conditional on the predictions of the black–box under fairness evaluation (see Fig. 1b). The black–box's input samples may consist of both numerical (N) and one–hot encoded (C) features. Both N and C may contain protected features that are not necessarily in the black–box's input space for models which do not directly discriminate against specific features. The numerical features N are pre–processed via mode–specific min–max normalization using a variational Gaussian mixture model, and the categorical features C are one–hot encoded. Note that the processing of the samples for the training of the black–box model and LUCID–GAN does not need to be similar. For example, many tree–based models do not require the one–hot encoding of categorical features.

   The generator receives noise which is drawn from a standard normal distribution, and a conditional vector containing the black–box's predictions and a masked version of C ($\widetilde{C}$) where only one category of a single one–hot encoded feature is equal to one. The critic gets real samples from the data and synthetic samples from the generator, where N^* is generated via Tanh functions and C^* via Gumbel–Softmax functions.[6] In addition, it also receives the black–box's predictions and $\widetilde{C}$ to check the validity of the sample.

   We use $\widetilde{C}$ to handle the sparse and highly imbalanced categorical features via a training–by–sampling technique. It is constructed by randomly sampling a one–hot encoded feature, and subsequently sampling a category from the empirical

---

[5] We use the same implementation and default values for the hyperparameters as CTGAN. For more details see Xu et al. [51]. The code for CTGAN is available at: `https://github.com/sdv-dev/CTGAN`.

[6] The generator also outputs the components of the variational Gaussian mixture model for each numerical feature as a one–hot encoded vector. The components are part of the input for the critic to address the non–Gaussian and multimodal distributions. They are further used to reverse the mode–specific min–max normalization.

log–frequency distribution.[7] All the others values are set to zero. To enforce the generator to generate the sampled category in $\widetilde{C}$, we apply a reconstruction loss to C^* and $\widetilde{C}$. To ensure the validity of the conditional vector, we pick the prediction corresponding to the sample from the empirical distribution. After training, we can use $\widetilde{C}$ to generate canonical inputs conditional on fixed values for a single category, such as setting "Male" for "Sex."

The critic is trained via the Wasserstein loss, while the generator also adds the reconstruction loss. The reconstruction loss is the cross–entropy between the sampled one–hot encoded feature and its synthetic counterpart in C^*. The cross–entropy is only computed on a single one–hot encoded feature, which forces the generator to replicate this condition in the synthetic sample.

3.4   Generating Canonical Sets with LUCID–GAN

After training, we use the generator of LUCID–GAN to generate canonical inputs (see Fig. 1c). Similar to the CGAN framework [38], we can use the conditional vector to control the generation process. Indeed, setting the prediction in the conditional vector to a specific value corresponds to maximizing the output activation via gradient descent. For example, in the case of a binary classifier, we can generate a canonical set that reveals the preferred output for receiving a "positive" and "negative" decision by setting the prediction in the conditional vector equal to a predicted probability of 1 and 0, respectively. Instead of comparing the predictions, we compare the (protected) feature distributions corresponding to large positive and negative predictions.

Using a generative model has the benefits of working for any black–box model and generating realistic synthetic samples. Note that if the protected features are not part of the black–box model's input space, we can still assess proxy discrimination by generating these features (as part of N^* or C^*) and concatenating them to the real samples (as part of N or C) given to the critic. We use $\widetilde{C}$ to assess intersectional discrimination by generating canonical inputs conditional on fixed values for protected features, such as setting "Male" for "Sex." This allows us to compare feature distributions for many possible scenarios where otherwise, data would be scarce, and the estimates of output–based metrics would be unreliable [33].

---

[7] For example, let $C_1 = [a, b]$ and $C_2 = [c, d, e]$ be two one–hot encoded features with a possible empirical sample $C = [C_1, C_2] = [0, 1, 0, 0, 1]$, wherein the first one–hot encoded feature we observe the second category b, and in the second one–hot encoded feature the third category e. We construct a masked version of C by first randomly selecting either C_1 or C_2, and then by sampling a category from their respective empirical log–frequency distribution. A potentially masked version of C could be $\tilde{C} = [1, 0, 0, 0, 0]$ where the first category a of the first one–hot encoded feature C_1 is sampled.

4   Experiments and Discussion

We show how to generate canonical sets via LUCID–GAN for various fairness evaluations, including direct, proxy, and intersectional discrimination, on the UCI Adult [15] and COMPAS [2] data sets. In the UCI Adult data set, we predict if a person earns more or less than $50,000$ per year, with more being the preferred output. For COMPAS, the task is to predict if a person will commit recidivism in the next two years, with no recidivism being the preferred output as the person can be released on bail. For both data sets, we consider "Race" and "Sex" to be the (legally) protected features, and for UCI Adult we also consider "Marital Status" and "Relationship." The UCI Adult data set has a fixed test set, while for the COMPAS data set, we sample 20% from the training data as the test set. The samples and predictions of the test set are also used to train LUCID–GAN and to compute the statistical output–based metrics DP and EOP.

   We compare LUCID–GAN with LUCID and the statistical output–based metrics DP and EOP (by comparing the TP and TPR, respectively) for the case of direct discrimination (see Fig. 2 and Table 1). We further demonstrate how to apply LUCID–GAN to the cases of proxy (See Fig. 3) and intersectional (see Fig. 4) discrimination. As LUCID only works for differentiable models and DP and EOP are often used for binary classification tasks, we use binary fully-connected neural network classifiers. For LUCID–GAN, the extension to other black–box classifiers is trivial, as it only requires samples and their corresponding predictions, making it a model–agnostic approach.

   The classifiers consist of hidden layers with ReLU activation functions and a softmax output layer with two output nodes. The number of hidden layers and nodes is decided by the accuracy on a validation set (20% from the training set) which is 83.9% for UCI Adult and 64.0% for COMPAS. This performance is in line with the standard benchmarks. We choose this standard architecture as the point of this experiment is not to achieve state–of–the–art performance but to demonstrate the capabilities of LUCID–GAN, which does not depend on the quality of the underlying model. Note that the computational complexity of LUCID–GAN lies in the training of the generator and the critic, which for our experiments was only a matter of minutes on a consumer CPU. After training, the generation of canonical inputs for various fairness evaluations is done via single forward passes through the generator.

4.1   Direct Discrimination

We compare LUCID–GAN with LUCID and the statistical output–based metrics DP and EOP (by comparing the disparities between TP and TPR, respectively) for the case of direct discrimination (see Fig. 2 and Table 1). After the classifiers are trained, we train LUCID–GAN on the test set and then generate 1000 synthetic samples for a positive and negative output (i.e., a predicted probability of 1 and 0, respectively). We refer to them as the positive and negative canonical sets. Note that we require no access to the training data and treat the underlying

classifier as a black–box. For LUCID, we also generate 1000 synthetic samples for a positive output starting from an initial random uniform distribution (see also Appendix A). Finally, we compute the TP and TPR of all the subpopulations for the protected features. A disparity between these metrics indicates potential discrimination towards the group with lower values.

We locate unfairness in the model's decision–making process with LUCID–GAN by comparing the feature distributions of the positive and negative canonical sets. Only inspecting the feature distributions of either the positive or negative canonical set can lead to misleading results, as the generator is trained to mimic the underlying distribution of its training samples. For example, in both the UCI Adult and COMPAS data sets, more samples correspond to the "Male" category than the "Female" category for the feature "Sex" which results in more males in both the positive and negative canonical sets.

While the results in our experiments are quite clear on visual inspection and the data sets are not high–dimensional, this may not always be the case, and presenting the results as such can be unfeasible and ambiguous. In that case, we suggest using distance metrics, such as the Wasserstein or Jensen–Shannon distance. However, computing these metrics yields no additional insights into our experiments.

In Fig. 2, we see that LUCID–GAN generates more realistic canonical inputs than LUCID. The most prominent examples are the continuous features, where LUCID often generates unrealistic values while LUCID–GAN generates feature distributions that closely match the expected empirical distributions. For example, in the feature "Education Level" of the UCI Adult data set, we see a spike in high–school graduates (9 years), bachelors (13 years), and masters (14 years). The realism is a result of the adversarial process where the critic ensures that the generator outputs a diverse set of synthetic samples which look similar to the training samples. The realism also seems to lead to more outspokenness in LUCID–GAN's canonical inputs for at least two reasons. First, the synthetic samples must respect all the dependencies between the features. For example, the combination of "Male," "Husband," and "Married" is a strongly preferred sample in LUCID–GAN's positive canonical set. The strong co–occurrence of these feature values may indicate a potential for proxy discrimination. Second, it may be easier for LUCID to change the continuous features than the categories, as a small change in categorical features does not make any difference after post–processing. Any change in the continuous features remains, while the categories need to shift from one category to another due to the one–hot encoding.

Table 1: Positivity Rate (PR) and True Positive Rate (TPR) of the subpopulations for the protected features in UCI Adult and COMPAS.

UCI ADULT						
Sex	MALE	FEMALE				
PR	31.0	11.3				
TPR	73.0	72.3				
		ASIAN PAC.	AMER. INDIAN			
Race	WHITE	ISLANDER	ESKIMO	BLACK	OTHER	
PR	26.0	29.7	12.8	11.9	19.7	
TPR	73.2	72.8	66.7	64.9	77.8	
Relationship	WIFE	OWN CHILD	HUSBAND	NOT IN FAMILY	OTHER RELATIVES	UNMARRIED
PR	47.0	1.9	45.6	10.2	3.3	5.6
TPR	71.0	100.0	72.8	78.7	100.0	81.8

Marital Status	MARRIED	DIVORCED	NEVER MARRIED	SEPARATED	WIDOWED	SPOUSE ABSENT	MILITARY SPOUSE
PR	45.3	9.7	4.7	7.0	9.1	12.6	36.4
TPR	72.6	75.0	86.2	55.6	85.7	100.0	80.0

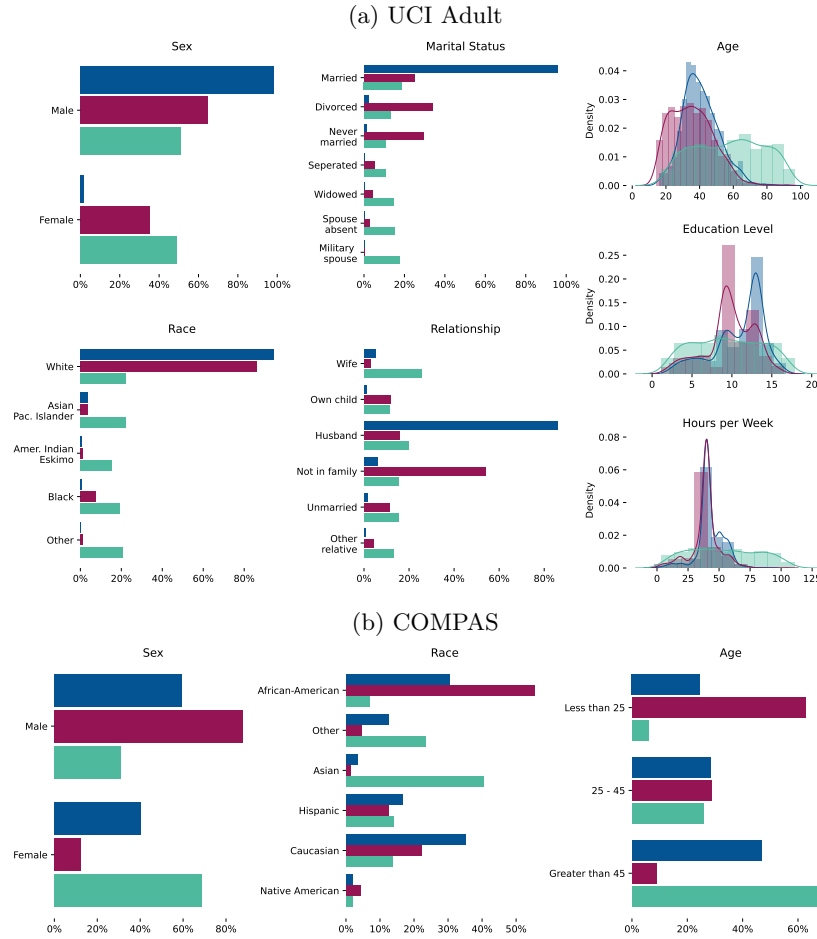| COMPAS						
Sex	MALE	FEMALE				
PR	53.6	62.7				
TPR	39.4	38.9				
	AFRICAN				NATIVE	
Race	AMERICAN	ASIAN	CAUCASIAN	HISPANIC	AMERICAN	OTHER
PR	49.2	71.4	61.9	62.5	60.0	54.0
TPR	37.6	0.00	41.9	52.0	60.0	25.0

(a) UCI Adult

(b) COMPAS

Fig. 2: Locating direct discrimination. The feature distributions for 1000 synthetic samples are shown for the positive and negative outputs in LUCID–GAN (dark blue and dark red, respectively) and positive outputs in LUCID (light green). For the UCI Adult data set, we show the protected features "Sex," "Race," "Marital Status" and "Relationship." We further show the features "Age," "Education Level," and "Hours per Week." For the COMPAS data set, we show the protected features "Sex" and "Race." We further show "Age," which is a categorical feature. Note that to interpret the results of LUCID–GAN, we need to compare the positive and negative canonical sets, while for LUCID we can compare the positive canonical set with an initial random uniform distribution (see also Appendix A). Only inspecting the feature distributions of either the positive or negative canonical set of LUCID–GAN can lead to misleading results as the generator is trained to mimic the underlying distribution of its training samples.

In Table 1, we see that the DP and EOP disparities do not always point to the same conclusions. For example, the disparity between the PR of "Male" and "Female" indicates a violation of DP, while their TPRs indicate that there is EOP for both the UCI Adult and COMPAS data sets. Additionally, many categories, including "Military Spouse" in the UCI Adult data set and "Asian" in the COMPAS data set, do not contain sufficient data points to make any conclusions [33]. Finally, the output–based metrics do not give any insights into the actual drivers of the unfairness. There are some notable differences between the output–based metrics and LUCID–GAN. For example, in the UCI Adult data set, the DP metric indicates that "Wife" is the preferred group in the "Relationship" feature. In contrast, the canonical sets indicate a strong preference for "Husband."

4.2   Proxy Discrimination

We retrain the models without the protected features "Race" and "Sex" and obtain a similar accuracy on the test set. We then generate a positive and negative canonical set including the left–out protected features, by generating these features (as part of C^*) and concatenating them to the real samples (as part of C) given to the critic. LUCID–GAN now receives predictions from the black–box model, which does not have "Race" and "Sex" in its input space. For the UCI Adult data set, we keep the protected features "Marital Status" and "Relationship" and assume that there are no protected features left in the COMPAS data set. For the PR and TPR, we obtain similar scores as for direct discrimination and therefore refer to Table 1 for a comparison. This confirms many previous findings that removing the protected attributes does not generally improve the disparity in the statistical output–based metrics [16].

By comparing the canonical sets from the case of proxy discrimination (see Fig. 3) with those of direct discrimination (see Fig. 2), we see that there are some notable differences (see also Appendix B). For example, in the COMPAS data set, the disparity between "Male" and "Female" in the positive and negative canonical sets has almost entirely disappeared. At the same time, the relative values in the "Race" feature remain mostly unchanged. This may indicate that the previous black–box model was directly discriminating based on "Sex," while the discrimination against "Race" is a combination of both direct and proxy discrimination. For the UCI Adult data set, we find that the discrimination against "Sex" and "Race" remains. This may result from the strong dependencies between many features, such as "White," "Male," "Married," and "Husband."
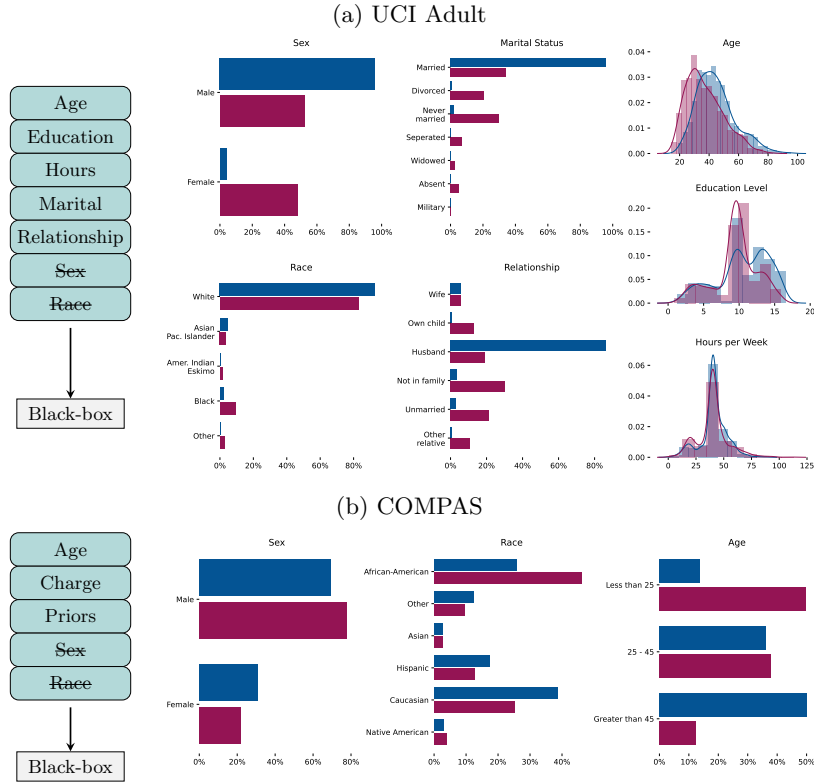
Fig. 3: Locating proxy discrimination. The feature distributions for 1000 synthetic samples are shown for the positive and negative outputs in LUCID–GAN (dark blue and dark red, respectively). For the UCI Adult data set, we show the protected features "Sex," "Race," "Marital Status" and "Relationship." We further show the features "Age," "Education Level," and "Hours per Week." For the COMPAS data set, we show the protected features "Sex" and "Race." We further show the feature "Age." Note that to interpret the results of LUCID–GAN, we need to compare the positive and negative canonical sets. Only inspecting the feature distributions of either the positive or negative canonical set of LUCID–GAN can lead to misleading results as the generator is trained to mimic the underlying distribution of its training samples.

(a) UCI Adult

All Marital Status

	Male	Female
White	0.98	0.93
Asian-Pac-Islander	0.01	0.02
Black	0.00	0.03
Amer-Indian-Eskimo	0.00	0.01
Other	0.00	0.00

Married

	Male	Female
White	0.99	1.00
Asian-Pac-Islander	0.01	0.00
Black	0.00	0.00
Other	0.00	0.00
Amer-Indian-Eskimo	0.00	0.00

(b) COMPAS

All Ages

	Male	Female
African-American	0.50	0.33
Caucasian	0.38	0.34
Hispanic	0.06	0.19
Other	0.03	0.10
Asian	0.01	0.03
Native American	0.01	0.01

Age > 45

	Male	Female
Caucasian	0.65	0.60
African-American	0.29	0.24
Other	0.03	0.06
Hispanic	0.02	0.08
Asian	0.01	0.01
Native American	0.00	0.01

Fig. 4: Locating intersectional discrimination. The percentage–wise amount of men and women per "Race" category which receive a positive output is shown for 1000 synthetic samples where "Sex" is fixed as "Male" and "Female," respectively. Additionally, for the UCI Adult data set, we show the percentage–wise amount of men and women per "Race" category which receive a positive output when the "Marital Status" is fixed as "Married," and for the COMPAS data set, when "Age" is fixed as "> 45."

4.3   Intersectional Discrimination

We use the original models from the direct discrimination evaluation with the protected features "Race" and "Sex" included. In this intersectional discrimination evaluation, we generate 1000 positive canonical inputs where we keep "Sex" fixed as "Male" and 1000 positive inputs where we keep "Sex" fixed as "Female" by using $\widetilde{C}$ in the conditional vector. Additionally, for the UCI Adult data set, we generate 1000 positive canonical inputs where we also keep "Marital Status" fixed as "Married," and for the COMPAS data set, we keep "Age" fixed as "> 45." We show in each case percentage–wise the amount of men and women per "Race" category which receive a positive output (see Fig. 4). On the heat maps, we show a two– and three–dimensional representation of the frequency of protected features at various intersections. For the UCI Adult data set, we find potential discrimination against women conditional on being "White," which entirely disappears if we add the condition of "Married." For the COMPAS data set, we see indications of bias toward women conditional on being "Caucasian" and "African-American." Additionally, when conditioning on "> 45" the number of positive outputs is considerably larger for "Caucasian" compared to "Hispanic" and "African-American." We believe that these heat maps are an ideal method for tracing potential sources of intersectional discrimination, especially when data samples are scarce.

5   Conclusion

The increasing use of Artificial Intelligence (AI) algorithms in (semi–)automated decision–making processes has raised concerns about discriminatory decision patterns. The literature on algorithmic fairness has primarily focused on defining unfairness and eliminating unethical biases, while recent efforts focus on rigorously detecting it. A recent proposal in this direction is LUCID (Locating Unfairness through Canonical Inverse Design), which generates canonical sets by performing gradient descent on the input space, revealing a model's desired input given a preferred output.

We present LUCID–GAN, which generates canonical inputs via a conditional generative model instead of gradient–based inverse design. Using a conditional generative model has several benefits, including that it applies to non–differentiable models, ensures that a canonical set consists of realistic inputs, and allows us to assess proxy and intersectional discrimination. LUCID–GAN is an input–based fairness evaluation method which takes a somewhat reverse approach to the statistical output–based metrics. Instead of comparing the predictions, we compare the (protected) feature distributions corresponding to large positive and negative predictions. The resulting canonical sets contain valuable information about the model's mechanisms, i.e., which feature values are essential to obtain specific outputs. This allows us to expose potential unethical biases in its internal logic by inspecting the distribution of the protected features.

We show how to generate canonical sets via LUCID–GAN for various fairness evaluations, including direct, proxy, and intersectional discrimination, on

the UCI Adult and COMPAS data sets. It allows for rigorously detecting unethical biases in black–box models without requiring access to the training data. Overall, we argue that LUCID–GAN is a valuable addition to the toolbox of algorithmic fairness evaluation, as it offers additional transparency into the sources of discrimination and is agnostic towards an a priori decision on the definition of fairness and choice of protected features.

References

1. Adams-Prassl, J., Binns, R., Kelly-Lyth, A.: Directly discriminatory algorithms. The Modern Law Review **86**(1), 144–175 (2023)
2. Angwin, J., Kirchner, L., Surya, M., Larson, J.: Machine Bias. There's software used across the country to predict future criminals. And it's biased against blacks. In: ProPublica (2016)
3. Arjovsky, M., Bottou, L.: Towards Principled Methods for Training Generative Adversarial Networks. In: International Conference on Learning Representations (2017)
4. Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein generative adversarial networks. In: Proceedings of the 34th International Conference on Machine Learning (2017)
5. Balagopalan, A., Zhang, H., Hamidieh, K., Hartvigsen, T., Rudzicz, F., Ghassemi, M.: The Road to Explainability is Paved with Bias: Measuring the Fairness of Explanations. In: Proceedings of the Conference on Fairness, Accountability, and Transparency (2022)
6. Barocas, S., Hardt, M., Narayanan, A.: Fairness and Machine Learning. fairmlbook.org (2019)
7. Barocas, S., Selbst, A.D.: Big Data's Disparate Impact. California Law Review **104**(3), 671–732 (2016)
8. Binns, R.: Fairness in machine learning: Lessons from political philosophy. In: Proceedings of the Conference on Fairness, Accountability, and Transparency (2018)
9. Binns, R.: On the apparent conflict between individual and group fairness. In: Proceedings of the Conference on Fairness, Accountability, and Transparency (2020)
10. Borisov, V., Leemann, T., Seßler, K., Haug, J., Pawelczyk, M., Kasneci, G.: Deep Neural Networks and Tabular Data: A Survey. arXiv preprint: 2110.01889 (2021)
11. Buolamwini, J., Gebru, T.: Gender shades: Intersectional accuracy disparities in commercial gender classification. In: Proceedings of the Conference on Fairness, Accountability, and Transparency (2018)
12. Chen, L., Ma, R., Hannák, A., Wilson, C.: Investigating the impact of gender on rank in resume search engines. In: Proceedings of the CHI Conference on Human Factors in Computing Systems (2018)
13. Corbett-Davies, S., Goel, S.: The measure and mismeasure of fairness: A critical review of fair machine learning. arXiv preprint: 1808.00023 (2018)
14. Crenshaw, K.: Mapping the Margins: Intersectionality, Identity Politics, and Violence against Women of Color. Stanford Law Review **43**(6), 1241–1299 (1991)
15. Dua, D., Graff, C.: UCI machine learning repository (2017)
16. Dwork, C., Hardt, M., Pitassi, T., Reingold, O., Zemel, R.: Fairness through awareness. In: Proceedings of the 3rd innovations in theoretical computer science conference (2012)
17. Engelmann, J., Lessmann, S.: Conditional Wasserstein GAN-based oversampling of tabular data for imbalanced learning. Expert Systems with Applications **174**, 114582 (2021)
18. Fazelpour, S., Danks, D.: Algorithmic bias: Senses, sources, solutions. Philosophy Compass **16**(8), e12760 (2021)
19. Friedler, S.A., Scheidegger, C., Venkatasubramanian, S.: On the (im) possibility of fairness. arXiv preprint: 1609.07236 (2016)
20. Goethals, S., Martens, D., Calders, T.: PreCoF: Counterfactual Explanations for Fairness. Working paper (2022)

21. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Advances in Neural Information Processing Systems (2014)
22. Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A.C.: Improved Training of Wasserstein GANs. In: Advances in Neural Information Processing Systems (2017)
23. Jang, E., Gu, S., Poole, B.: Categorical Reparameterization with Gumbel-Softmax. In: International Conference on Learning Representations (2017)
24. Kearns, M., Neel, S., Roth, A., Wu, Z.S.: Preventing Fairness Gerrymandering: Auditing and Learning for Subgroup Fairness. In: Proceedings of the 35th International Conference on Machine Learning (2018)
25. Kleinberg, J., Mullainathan, S.: Simplicity Creates Inequity: Implications for Fairness, Stereotypes, and Interpretability. In: Proceedings of the ACM Conference on Economics and Computation (2019)
26. Kleinberg, J., Mullainathan, S., Raghavan, M.: Inherent trade-offs in the fair determination of risk scores. arXiv preprint: 1609.05807 (2016)
27. Kong, Y.: Are "Intersectionally Fair" AI Algorithms Really Fair to Women of Color? A Philosophical Analysis. In: Proceedings of the Conference on Fairness, Accountability, and Transparency (2022)
28. Kusner, M.J., Loftus, J., Russell, C., Silva, R.: Counterfactual fairness. In: Advances in Neural Information Processing Systems (2017)
29. Lang, O., Gandelsman, Y., Yarom, M., Wald, Y., Elidan, G., Hassidim, A., Freeman, W.T., Isola, P., Globerson, A., Irani, M., et al.: Explaining in Style: Training a GAN to explain a classifier in StyleSpace. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (2021)
30. Ledford, H.: Millions of black people affected by racial bias in health-care algorithms. Nature **574**(7780), 608–610 (2019)
31. Lin, Z., Khetan, A., Fanti, G., Oh, S.: PacGAN: The power of two samples in generative adversarial networks. In: Advances in Neural Information Processing Systems (2018)
32. Lindholm, M., Richman, R., Tsanakas, A., Wüthrich, M.: Discrimination–free insurance pricing. ASTIN Bulletin **52**(1), 55–89 (2022)
33. Lum, K., Zhang, Y., Bower, A.: De-biasing "bias" measurement. In: Proceedings of the Conference on Fairness, Accountability, and Transparency (2022)
34. Makhlouf, K., Zhioua, S., Palamidessi, C.: On the applicability of machine learning fairness notions. ACM SIGKDD Explorations Newsletter **23**(1), 14–23 (2021)
35. Mazijn, C., Danckaert, J., Ginis, V.: How Do the Score Distributions of Subpopulations Influence Fairness Notions? In: Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (2021)
36. Mazijn, C., Prunkl, C., Algaba, A., Danckaert, J., Ginis, V.: LUCID: Exposing Algorithmic Bias through Inverse Design. In: Proceedings of the 37th AAAI Conference on Artificial Intelligence (2023)
37. Meng, C., Trinh, L., Xu, N., Enouen, J., Liu, Y.: Interpretability and fairness evaluation of deep learning models on MIMIC-IV dataset. Scientific Reports **12**(7166) (2022)
38. Mirza, M., Osindero, S.: Conditional Generative Adversarial Nets. arXiv preprint: 1411.1784 (2014)
39. Nguyen, A., Clune, J., Bengio, Y., Dosovitskiy, A., Yosinski, J.: Plug & Play Generative Networks: Conditional Iterative Generation of Images in Latent Space. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2017)

40. Nguyen, A., Dosovitskiy, A., Yosinski, J., Brox, T., Clune, J.: Synthesizing the preferred inputs for neurons in neural networks via deep generator networks. In: Advances in Neural Information Processing Systems (2016)
41. Nguyen, A., Yosinski, J., Clune, J.: Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2015)
42. Northcutt, C.G., Athalye, A., Mueller, J.: Pervasive label errors in test sets destabilize machine learning benchmarks. In: Advances in Neural Information Processing Systems Track on Datasets and Benchmark (2021)
43. Obermeyer, Z., Powers, B., Vogeli, C., Mullainathan, S.: Dissecting racial bias in an algorithm used to manage the health of populations. Science **366**(6464), 447–453 (2019)
44. Odena, A., Olah, C., Shlens, J.: Conditional Image Synthesis with Auxiliary Classifier GANs. In: Proceedings of the 34th International Conference on Machine Learning (2017)
45. Prince, A.E., Schwarcz, D.: Proxy Discrimination in the Age of Artificial Intelligence and Big Data. Iowa Law Review **105**, 1257–1318 (2019)
46. Raghavan, M., Barocas, S., Kleinberg, J., Levy, K.: Mitigating bias in algorithmic hiring: Evaluating claims and practices. In: Proceedings of the Conference on Fairness, Accountability, and Transparency (2020)
47. Shwartz-Ziv, R., Armon, A.: Tabular data: Deep learning is not all you need. Information Fusion **81**, 84–90 (2022)
48. Simonyan, K., Vedaldi, A., Zisserman, A.: Deep inside convolutional networks: Visualising image classification models and saliency maps. In: Workshop at International Conference on Learning Representations (2014)
49. Slack, D., Hilgard, S., Jia, E., Singh, S., Lakkaraju, H.: Fooling LIME and SHAP: Adversarial Attacks on Post hoc Explanation Methods. In: Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (2020)
50. Wachter, S., Mittelstadt, B., Russell, C.: Why fairness cannot be automated: Bridging the gap between EU non-discrimination law and AI. Computer Law & Security Review **41**, 105567 (2021)
51. Xu, L., Skoularidou, M., Cuesta-Infante, A., Veeramachaneni, K.: Modeling Tabular data using Conditional GAN. In: Advances in Neural Information Processing Systems (2019)
52. Zhao, Z., Kunar, A., Birke, R., Chen, L.Y.: CTAB-GAN: Effective Table Data Synthesizing. In: Proceedings of The 13th Asian Conference on Machine Learning (2021)
53. Zhou, Z., Cai, H., Rong, S., Song, Y., Ren, K., Zhang, W., Wang, J., Yu, Y.: Activation maximization generative adversarial nets. In: Proceedings of the 35th International Conference on Learning Representations (2018)
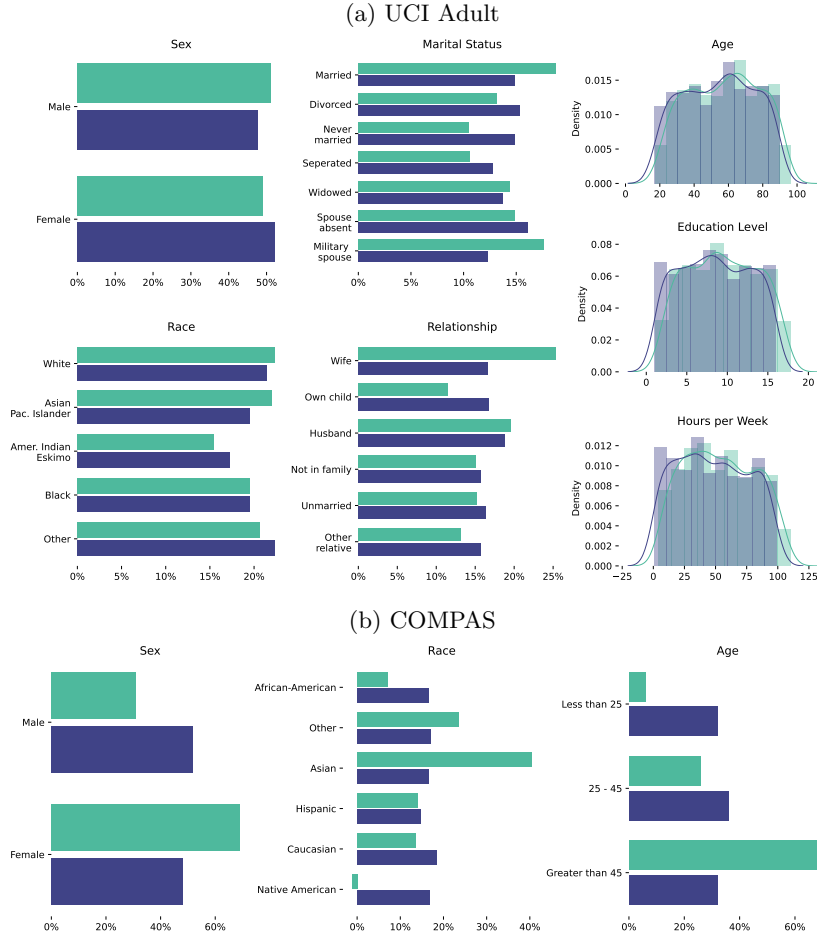
A   LUCID

Fig. 5: Locating direct discrimination. The feature distributions for 1000 synthetic samples are shown for the positive outputs in LUCID (light green) starting from an initial random uniform distribution (dark blue). For the UCI Adult data set, we show the protected features "Sex," "Race," "Marital Status," and "Relationship." We further show the features "Age," "Education Level," and "Hours per Week." For the COMPAS data set, we show the protected features "Sex" and "Race." We further show the feature "Age."

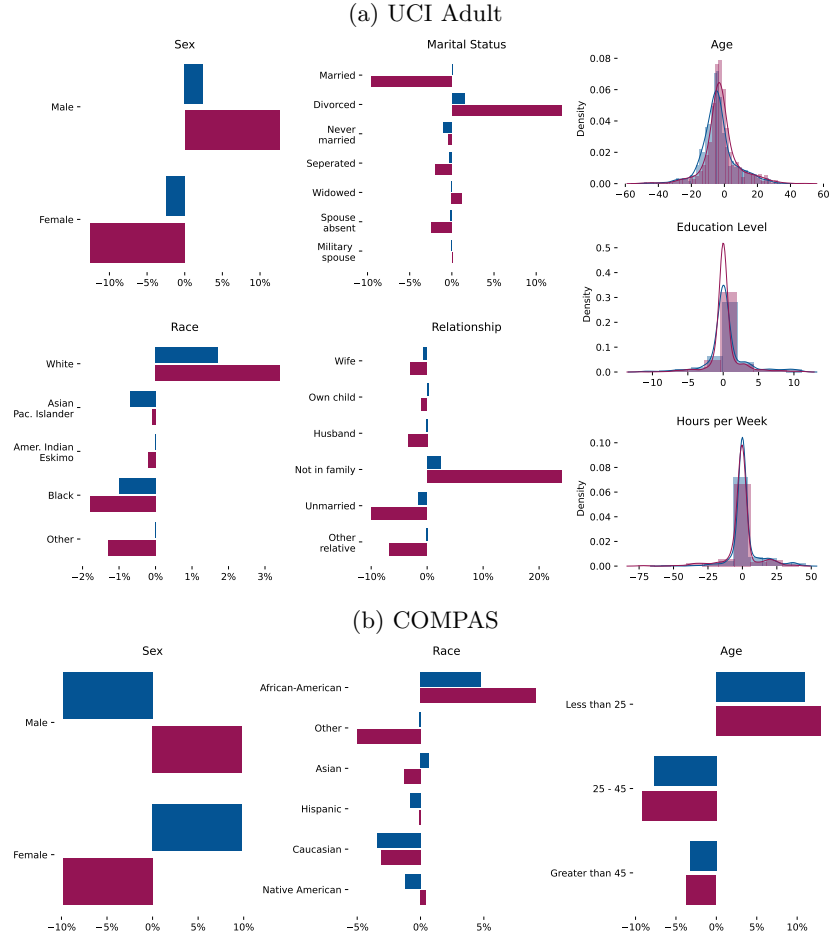# B   Comparing Direct and Proxy Discrimination

Fig. 6: The difference between the feature distributions of direct and proxy discrimination for the positive and negative outputs in LUCID–GAN (dark blue and dark red, respectively) are shown. For the UCI Adult data set, we show the protected features "Sex," "Race," "Marital Status," and "Relationship." We further show the features "Age," "Education Level," and "Hours per Week." For the COMPAS data set, we show the protected features "Sex" and "Race." We further show the feature "Age." A positive value indicates that the feature appears more frequently in the distribution of direct discrimination.