
Deontological Ethics By Monotonicity Shape Constraints

Serena Wang
Google Research

Maya Gupta
Google Research

Abstract

We demonstrate how easy it is for modern machine-learned systems to violate common deontological ethical principles and social norms such as “favor the less fortunate,” and “do not penalize good attributes.” We propose that in some cases such ethical principles can be incorporated into a machine-learned model by adding shape constraints that constrain the model to respond only positively to relevant inputs. We analyze the relationship between these deontological constraints that act on individuals and the consequentialist group-based fairness goals of one-sided statistical parity and equal opportunity. This strategy works with sensitive attributes that are Boolean or real-valued such as income and age, and can help produce more responsible and trustworthy AI.

1 Introduction

As the use of machine-learned (ML) models broadens, strategies are sought to ensure machine-learned systems behave *responsibly*, and are *ethical* and *fair*. There may be many different reasonable and conflicting ethical stances for a given problem (Haidt, 2013; Singer, 2016; Gray and Graham, 2018; Corbett-Davies and Goel, 2018; Binns, 2018). Thus, no single strategy for making machine learning fairer is likely to be sufficient.

In this paper, we show that nonlinear machine learning can easily produce trained models that violate social norms or ethics that can be described as: *certain inputs should not have a negative effect on an outcome*. For example, the toy example in Fig. 1 shows a nonlinear 1-d model trained to predict how highly a resume will be scored based on the candidate’s years of experience.

The best fit model (purple dashed line) sometimes penalizes candidates for having *more* job experience. For candidates with many years of experience, the dashed model picked up the age discrimination in the biased training samples. In addition, the dashed model also penalizes job experience in the lower-end of years of experience simply due to overfitting. Such unfair responses are very much a danger in sparse regions of a feature space with modern over-parameterized nonlinear models.

Fig. 1 also illustrates our proposed solution (blue solid line): train the model with monotonicity shape constraints to guarantee that the model can only use job experience as positive evidence. This is a *deontological* solution, and thus differs from many existing mathematical expressions of fairness that are *consequentialist* and *statistical* (Friedler et al., 2019; Hardt et al., 2016; Binns, 2018; Sandvig et al., 2016). Both types of goals may be of interest to a practitioner (Sandvig et al., 2016), and we explore how they relate theoretically and experimentally.




Figure 1: Toy example showing how a monotonicity constraint can protect against unfair penalization.

The main contributions of this paper are: (i) identifying that violations of monotonicity in ML models may pose significant ethical issues and risk of societal harm, (ii) demonstrating experimentally that these violations can easily occur in practice with nonlinear machine learned models, (iii) showing that existing shape constrained ML can be used effectively to ameliorate such problems, and (iv) theoretically and empirically analyzing how the proposed monotonicity constraints relate to statistical fairness notions.

Model type	# shape constraints	Train Acc.	Test Acc.
GAM	0	94.93%	94.89%
GAM	2	94.90%	94.97%
DNN	0	94.97%	94.89%
GBT	0	95.04%	94.80%

Table 1: Law School Admissions Experiment Results. Two monotonicity constraints ensure that individuals aren’t penalized for higher GPA or LSAT score.

2 Illustrating Unfair Penalization

To further illustrate the potential for machine-learning to produce objectionable models, consider the *Law School Admissions* dataset (Wightman, 1998). Suppose this data is used to predict whether a person would pass the bar exam based on their LSAT score and undergraduate GPA, and that the classifier’s score was used to guide law school admissions or scholarships. We trained a standard two-layer deep neural network (DNN) (more experimental details in Sec. 8) and show the model’s output for each possible input in Fig. 2. The DNN sometimes penalizes people for having a higher GPA: for example, with an LSAT score of 15, the DNN rewards students with a lower 2.7 GPA over students who earned a higher 3.5 GPA. Similarly, if a student has a GPA of 2.5, the DNN gives that student a higher score for scoring 10 on the LSAT than if they had scored 15. Thus this model violates merit-based social norms and the “best-qualified” ethical principle (Hunter and Schmidt, 1976) (we also acknowledge that the very use of standardized test scores or GPA for allocating goods may raise other ethical issues (Hunter and Schmidt, 1976)). Training with gradient boosted trees has the same problem: the model penalizes some people for raising their GPA or LSAT score.

Each of the two-dimensional models shown in Fig. 2 were trained on 19,064 training examples, which may sound like plenty of data to learn a model on two inputs, but the non-uniform distribution of the training data means that some regions of the feature space were sparse and the model may have overfit (for a density plot, see Fig. 9 in the Appendix). These monotonicity violations can also occur from legitimate, clean data. To see this, note that how hard it is to get a good GPA can vary greatly between schools, and imagine as an extreme example that there was a large university that simply gave every student a 2.7 GPA, which could cause some of the problems seen in this model.

Our proposal is to train the model with monotonicity constraints. An example monotonic model is shown in (c) of Fig. 2. It is a generalized additive model (GAM) trained with the constraints that it never penalizes

higher GPA’s for any LSAT score, and that it never penalizes higher LSAT scores for any GPA. Training the same GAM *without* monotonicity constraints is shown in the lower right, and also produces an objectionable model. The test accuracies of these four models are given in Table 1 are similar (more experimental details in Sec. 8).

3 Monotonicity Fairness Constraints

The motivating examples in Figures 1 and 2 focused on the concern of *unfair penalization* – that there may be inputs that a responsible model may reward but should never penalize. A second ethical pattern we consider is to *favor the less fortunate* – there may be inputs that help us identify the less fortunate and favor them if there are no other relevant differences. Policies favoring the less fortunate have been well studied in economics (Coate and Loury, 1993). As shown in Figures 1 and 2, we propose addressing these two principles by constraining the ML model to only respond positively to relevant inputs if all other inputs are fixed.

Throughout, we focus on ML models that produce a score that is used to determine some benefit, such as a better credit rating or a scholarship. Specifically, consider a learned model $f(x, z)$ where $x \in \mathbb{R}^D$ is a D -dimensional feature vector, and $z \in \mathbb{R}$ is another input that is of ethical interest (such as age or income). The model $f(x, z)$ satisfies a *positive monotonicity shape constraint* (Groeneboom and Jongbloed, 2014) with respect to z if for any $\delta > 0$ and any choice of x and z , $f(x, z + \delta) \geq f(x, z)$. If f is differentiable with respect to z , this constraint is equivalent to non-negative slope $\frac{\partial(f(x, z))}{\partial z} \geq 0$ for any x, z . *Strict monotonicity* replaces the \geq sign above with $>$. Reverse definitions produce negative monotonicity constraints. To apply a monotonicity shape constraint to a categorical feature, one can express each category as a Boolean feature indicating membership in the category.

As defined above, and as is standard in the shape constraints literature (Groeneboom and Jongbloed, 2014), our proposed monotonicity shape restrictions are *ceterus paribus*: that is, $f(x)$ should behave monotonically with respect to increases in each protected feature z , but only *when all other features x are held fixed*. Here we take the input features x as given, but of course it is also important to have the best possible set of features x .

4 More Example Scenarios

We further illustrate the critical importance and breadth of situations in which we need ML models to behave consistently with a society’s norms or ethical




Figure 2: Real data example predicting if a student will pass the bar. Predictions by a neural network (a) and gradient boosted trees (b) penalize some students if they increase their GPA or LSAT score. We trained a generalized additive model (GAM) with constraints to be monotonically increasing in both GPA and LSAT score (c). The same GAM model trained *without* monotonicity constraints also violates monotonicity in GPA (d).

policies. Then in Section 8, we show experimentally on two more problems that nonlinear ML models do not naturally pick up such norms and policies, and that training with monotonicity constraints can be used to incorporate desired policies.

Crimes and Misdemeanors: An ML system trained to determine fines for misdemeanors would likely be considered fairer if the fines were monotonically increasing with respect to the magnitude of the illegality, e.g. a larger fine for illegally parking one’s car for longer, or for exceeding the speed limit by more (Allen et al., 2015). Also, in many societies it is expected that a juvenile will not be penalized more heavily than an adult for the same crime, all else equal (Allen et al., 2015). Many societies prefer not to give harsher penalties to first-time offenders than to repeat offenders (Allen et al., 2015).

Pay: People generally feel it is unfair to get paid less for doing more of the same work (Fehr and Schmidt, 1999). Suppose a model is trained to advise parents on how much to pay their babysitter. It may be desirable to workers if the recommended pay were constrained to be a monotonically increasing function of the number of hours worked, all else equal. Similarly, consider an app that uses an ML model to advise people on how much to tip their waiter in America. Such a model would be better aligned with American societal norms if it recommended larger tips for more expensive meals, and if it recommended higher tip percentages for more upscale establishments, all else equal (Azar, 2004).

Medical Triage: In some medical contexts, it is considered more ethical or more responsible to prioritize patients based on their risk or neediness (Iserson and Moksop, 2007). For example, the United States Transplant Board has a policy of giving a sicker person a higher score to receive a transplant for ethical reasons (L. Bernstein, 2017). Similarly an emergency room scoring patients for prioritization may require that patients that have waited longer are treated first, if all other relevant characteristics are equal (Iserson and Moksop,

2007). More generally, *first-come first-served* is a common principle underpinning civil society (Beauchamp and Childress, 2001; Sugden, 1989).

5 Related Work in Ethical ML

This work fits into the broader literature on *machine ethics* (Anderson and Anderson, 2007; Moor, 2006). Ethics itself is a broad field of philosophy that encompasses many questions, and there may be many different reasonable ethical stances and criteria for a given problem (Binns, 2018; Thiroux and Krasemann, 2017; Haidt, 2013; Gray and Graham, 2018; Corbett-Davies and Goel, 2018; Singer, 2016; Gruetzemacher, 2018; Allen et al., 2015; Rawls, 1971).

Deontological: Most recent work in machine learning fairness has been *consequentialist*: focusing on ensuring that machine-learned models deliver statistically-similar performance for different groups (Binns, 2018; Zafar et al., 2015; Goh et al., 2016; Hardt et al., 2016; Zafar et al., 2017; Donini et al., 2018; Agarwal et al., 2018; Goel et al., 2018; Cotter et al., 2019a,b). Those efforts solve a different important problem than the one we target here, and are complementary to this proposal. This proposal is instead *deontological*, in that it enables people to impose rules on how the model can respond to inputs, producing an *implicit ethical agent* in the terminology of Moor (Moor, 2006). Sandvig et al. recently called for more research into deontological algorithms, noting that, “Applied ethics in real-world settings typically incorporates both rule-based and consequences-based reasoning” (Sandvig et al., 2016).

Monotonicity and Fairness: Concurrently, Cole and Williamson (2019) also recognized the importance of monotonicity in a fairness context. That work differs in framing: they use monotonicity to reduce unfair *resentment*. Our work further differs in its theoretical results and comparisons.

Asymmetric: Another fairness principle is that *similar individuals should receive similar treatment* (Dwork

et al., 2012); while also important, that principle aims for equal outcomes for two examples, whereas this proposal aims to make sure that any unequal treatment is unequal in the appropriate direction, such as *favor the less fortunate*.

Counterfactual Fairness: Counterfactual fairness (Kusner et al., 2017; Pearl et al., 2016) says that changing a protected attribute A while holding things not causally dependent on A constant will not change the distribution of the model output. This is similar to the definition of monotonicity in section 3, but is focused on treating certain cases the *same* rather than preferring one case to another.

Continuous Sensitive Features: The proposed monotonicity constraints handles real-valued attributes natively. Other recent efforts have also been devised to handle continuous protected attributes, like age, for consequentialist fairness goals (Raff et al., 2018; Kearns et al., 2018; Komiyama et al., 2018).

Fair Ranking and Fair Regression: Monotonicity constraints can be applied to ranking and regression models. Other notions of ethics have also been considered for ranking models (Berk et al., 2017; Zehlike et al., 2017; Celis et al., 2018; Beutel et al., 2019; Singh and Joachims, 2018) and regression (Komiyama et al., 2018; Pérez-Suay et al., 2017; Berk et al., 2017; Agarwal et al., 2019).

6 Related Work in Training Models with Monotonicity Constraints

Monotonicity shape constraints have long been used to capture prior knowledge and regularize estimation problems to improve a model’s generalization to new test examples (Barlow et al., 1972; Chetverikov et al., 2018; Ben-David, 1992; Archer and Wang, 1993; Sill and Abu-Mostafa, 1997; Kotlowski and Slowinski, 2009; Groeneboom and Jongbloed, 2014; Gupta et al., 2016; Canini et al., 2016; You et al., 2017; Bonakdarpour et al., 2018). In this paper, we point out that monotonicity shape constraints can and should *also* be used to ensure that machine-learned models behave consistently with societal norms and prima facie duties (Ross, 2002) that can be expressed as monotonic relationships.

Unlike its use for regularization, applying monotonicity constraints to impose ethical principles may actually hurt test accuracy if the training and test data is biased (as in the toy example of Fig. 1). However, in all three of our real-data experiments the test accuracy was little changed by adding these constraints (see Section 8 and Fig. 2 and Appendix D).

Constraining multi-dimensional models to obey mono-

tonic shape constraints has been shown to work for a variety of function classes, including neural networks (Archer and Wang, 1993; Sill and Abu-Mostafa, 1997) and trees (Ben-David, 1992; Kotlowski and Slowinski, 2009; Bonakdarpour et al., 2018).

Our experiments use the open-source TensorFlow Lattice 2.0 package (Google AI Blog, 2020), which enables training GAMs and lattice models with monotonicity constraints (Gupta et al., 2016; Canini et al., 2016; You et al., 2017).

7 Relationship to Statistical Fairness

In this section we analyze how the proposed deontological monotonicity constraints interact with consequentialist statistical fairness goals that are based on aggregate outcomes. For example, suppose a national funding agency enforces that poorer schools are funded more often as richer schools on average (we will call this *one-sided statistical parity*). By only considering the national average funding rates, richer schools may still get funded more often than poorer schools in some states. In contrast, if the state were the only input X , then a deontological monotonicity constraint would guarantee that for each state, poorer schools would get a higher funding rate than richer schools.

Hardt et al. (2016) showed that any *oblivious* statistical fairness measure that doesn’t depend on X or the function form $f(X, Z)$ can fail to identify forms of discrimination. Popular statistical notions like *statistical parity* (Dwork et al., 2012) and *equal opportunity* (Hardt et al., 2016) are oblivious, whereas monotonicity is not. In the next section we show specifically that satisfying *one-sided statistical parity* does not imply monotonicity, but that satisfying monotonicity can impact and bound one-sided statistical parity and equal opportunity (proofs in Appendix).

7.1 Bounds on One-sided Statistical Parity

Statistical parity is a well known measure of fairness for ML models (Dwork et al., 2012; Zafar et al., 2015; Cotter et al., 2019b). Consider a model $f(X, Z)$ that takes as input a random feature vector $X \in \mathbb{R}^D$ and a random protected attribute Z , which may be categorical or real-valued. Suppose the model f outputs a real-valued score (we treat the special case of classifiers next). Then *statistical parity* requires $E[f(X, Z)|Z = j] = E[f(X, Z)|Z = k]$ for any j, k . Because we focus on asymmetric goals like favoring the less fortunate, we consider *one-sided* statistical parity: $E[f(X, Z)|Z = j] \leq E[f(X, Z)|Z = k]$ for any $j \leq k$.

We show in Appendix B.1 that due to Simpson’s paradox (Bickel et al., 1975), that a monotonicity constraint

on Z is not sufficient to guarantee one-sided statistical parity with respect to Z . However, monotonicity does imply a *bound* on the one-sided statistical parity violation: in Lemma 1, we show that if f is monotonic with respect to Z , then the one-sided statistical parity violation between $Z = j$ and $Z = k$ will be bounded by the maximum density ratio between the two groups.




Figure 3: Illustration of Lemma 1. The dotted line shows the X value achieving the max density ratio C .

Lemma 1. *Let (Ω, \mathcal{F}) be a measurable space with a regular conditional probability property, and let $X : \Omega \rightarrow \mathbb{R}^D$, $Z : \Omega \rightarrow \mathbb{R}$ be \mathcal{F} -measurable random variables. Suppose P_j and P_k are σ -finite probability measures on (Ω, \mathcal{F}) , where P_j denotes the conditional probability measure of X given that $Z = j$, and P_k denote the same for $Z = k$, and P_j is absolutely continuous with respect to P_k . Let $f : \mathbb{R}^D \times \mathbb{R} \rightarrow \mathbb{R}$ be defined as in Section 3, and $f(x, z) \geq 0$ for all $x \in \mathbb{R}^D$, $z \in \mathbb{R}$. If the function f satisfies monotonicity in the second argument such that for $j \leq k$, $f(x, j) \leq f(x, k)$ for all $x \in \mathbb{R}^D$, and if the Radon-Nikodym derivative $\frac{dP_j}{dP_k}$ is bounded almost everywhere with respect to P_k by a finite constant $C > 0$, then*

$$E[f(X, Z)|Z = j] \leq CE[f(X, Z)|Z = k]. \quad (1)$$

If the conditional probability distribution of X given $Z = j$ has density $p_{X|Z=j}(x)$, and X given $Z = k$ has density $p_{X|Z=k}(x)$, then Lemma 1 says that if the likelihood ratio $\frac{p_{X|Z=j}(x)}{p_{X|Z=k}(x)} \leq C$ almost everywhere with respect to P_k for finite $C > 0$, then the one-sided statistical parity bound (1) holds.

Fig. 3 illustrates Lemma 1 with an example: suppose Z is categorical and denotes either a poorer or richer high school, and suppose $D = 1$ with $X \in [0, 20]$ being the number of hours of extracurricular activities a student does each week. Let $f(X, Z)$ be a score used to determine a student’s admission to some college. Then if f is monotonic in Z such that f never gives a lower admissions score to a poorer student if their X value is the same as a richer student, then Lemma 1 says that the average score for richer students will be no more than C times the average score for poorer students, where C is given by the maximum ratio of the two distributions over X .

While Lemma 1 shows that monotonicity implies a bound on the one-sided statistical parity violations, the converse does not hold: a model satisfying statistical parity can have arbitrarily high monotonicity violations (proof in Appendix B.2). This can be ethically problematic if overlooked by practitioners.

While Lemma 1 provides a *worst case bound* on the statistical parity violations of monotonic functions, we next ask, can imposing monotonicity ever make statistical parity violations worse? Lemma 2 shows that for any model f , the monotonic projection of f cannot have worse statistical parity violations *on average*.

Lemma 2. *Let $f : \mathcal{X} \times \mathcal{Z} \rightarrow \mathbb{R}$, where $\mathcal{X} \subseteq \mathbb{R}^D$, $\mathcal{Z} \subseteq \mathbb{R}$. Assume that \mathcal{X}, \mathcal{Z} are both finite, with $X \in \mathcal{X}$, $Z \in \mathcal{Z}$. Let \tilde{f} be the projection of f onto the set of functions over $\mathcal{X} \times \mathcal{Z}$ that are monotonic with respect to Z such that for $j \leq k$, $\tilde{f}(x, j) \leq \tilde{f}(x, k)$. For $z_{(i)} \in \mathcal{Z}$, let $z_{(1)} \leq z_{(2)} \leq \dots \leq z_{(|\mathcal{Z}|)}$. Define the average statistical parity violation:*

$$R_f \triangleq \sum_{i=1}^{|\mathcal{Z}|} \frac{E[f(X, Z)|Z = z_{(i)}] - E[f(X, Z)|Z = z_{(i+1)}]}{|\mathcal{Z}|}$$

Then $R_{\tilde{f}} \leq R_f$.

When $|\mathcal{Z}| = 2$, Lemma 2 also bounds the worst case violation of \tilde{f} . However, for $|\mathcal{Z}| > 2$, there is no such worst case guarantee, and there may be pairs j, k where \tilde{f} has a worse one-sided statistical parity violation than f (proof in Appendix B.3).

7.2 Bounds for Binary Classifiers

In binary classification, an example has an associated true label $Y \in \{0, 1\}$, and the model outputs a binary decision $\hat{Y} \in \{0, 1\}$. By definition, a monotonicity shape constraint on Z implies a one-sided bound on the conditional probabilities: for any $j \leq k$ and for all x , $P(\hat{Y} = 1|X = x, Z = j) \leq P(\hat{Y} = 1|X = x, Z = k)$.

Because the label Y can be modeled as a Bernoulli random variable, the goal of statistical parity is equivalent to *marginal independence*, that is, for any j, k , $P(\hat{Y} = 1|Z = j) = P(\hat{Y} = 1|Z = k)$. Correspondingly, the bound (4) on the statistical parity becomes a bound on the marginal probabilities: $P(\hat{Y} = 1|Z = j) \leq CP(\hat{Y} = 1|Z = k)$.

For binary classifiers, we can give a more explicit bound:

Lemma 3. *Suppose X is a continuous (or with a straightforward extension, discrete) random variable, and let \mathcal{S} be a nonempty set such that for all $x \in \mathcal{S}$, the joint probability density values $p_{X, \hat{Y}|Z=z}(x, 1) > 0$ for $z = j, k$. Suppose we have monotonicity where $f(x, j) \leq f(x, k)$ for $j \leq k$ for all $x \in \mathcal{S}$. For a binary classifier this implies $P(\hat{Y} = 1|X = x, Z = j) \leq$*

$P(\hat{Y} = 1|X = x, Z = k)$. Then we can bound one-sided statistical parity as follows:

$$\frac{P(\hat{Y} = 1|Z = j)}{P(\hat{Y} = 1|Z = k)} \leq \inf_{x \in \mathcal{S}} \frac{p_{X|Z=j}(x)p_{X|\hat{Y}=1,Z=k}(x)}{p_{X|Z=k}(x)p_{X|\hat{Y}=1,Z=j}(x)}$$

The bound in Lemma (3) contains two likelihood ratios: $\frac{p_{X|Z=j}(x)}{p_{X|Z=k}(x)}$ and $\frac{p_{X|\hat{Y}=1,Z=k}(x)}{p_{X|\hat{Y}=1,Z=j}(x)}$. The first is the same as in Lemma 1. The second is the inverse of that likelihood ratio, conditioned on $\hat{Y} = 1$. When the first likelihood ratio is low, the second inverse likelihood ratio may be high, producing a trade-off between these two ratios. We describe an example in Appendix C.

Similarly, for *equal opportunity* (Hardt et al., 2016), we have Lemma 4:

Lemma 4. Let $Y \in \{0, 1\}$ be a random variable representing the target. Let \mathcal{S} be a nonempty set such that for all $x \in \mathcal{S}$, the following joint probability density values are non-zero for $z = j, k$: $p_{X,Y,\hat{Y}|Z=z}(x, 1, 1) > 0$ and $p_{X,Y|\hat{Y}=1,Z=z}(x, 1) > 0$. Then,

$$\frac{P(\hat{Y} = 1|Y = 1, Z = j)}{P(\hat{Y} = 1|Y = 1, Z = k)} \leq \inf_{x \in \mathcal{S}} \frac{c_j(x)}{c_k(x)}$$

$$\text{where } c_z(x) = \frac{p_{X|Z=z}(x)P(Y = 1|\hat{Y} = 1, Z = z)}{p_{X|\hat{Y}=1,Z=z}(x)P(Y = 1|Z = z)}$$

We supplement these bounds with empirical results.

8 Experiments

We demonstrate experimentally with three public datasets that (i) nonlinear ML can violate common ethical policies or norms, (ii) training with monotonicity constraints can be used to impose such policies to the extent that the choice of model inputs x enables, (iii) test accuracy may not be hurt, (iii) statistical fairness violations may be reduced.

Model training details: Code for these experiments and further tutorials for using the open-source TensorFlow Lattice 2.0 library (Google AI Blog, 2020) are available at <https://github.com/tensorflow/lattice/blob/master/docs/tutorials/>. All experiments used a nonlinear generalized additive model (GAM) (Hastie and Tibshirani, 1990), also called a *calibrated linear* model in the TensorFlow Lattice, of the form $f(x) = \sum_{d=1}^D c_d(x[d]; \beta_d)$, where each $c_d(\cdot)$ is an one-dimensional piecewise-linear function parameterized by K interpolated (key, value) pairs where the keys are set to match the K quantiles of the training data and their corresponding values $\beta_d \in R^K$ are trained (Gupta et al., 2016). We used a fixed $K = 20$ parameters for each of the D one-dimensional transforms, for all experiments. All DK parameters were

jointly trained using projected stochastic gradient descent. For each input d , one can choose to constrain $f(x)$ to be monotonically increasing with respect to d by constraining its one-dimensional curve $c_d(x[d])$ to be monotonic (Gupta et al., 2016). The TensorFlow Lattice package enables fitting more flexible monotonic models, but we felt it was most compelling to show that these monotonicity violations occur even with simple nonlinear models like GAMs, which are popular in the statistics literature. We randomly uniformly split each dataset into 70% training, 10% validation, and 20% test examples. We validated the learning rates using grid search by powers of 10 ensuring that the optimal learning rates did not fall on the extremes, and trained for 1000 epochs (more than sufficient for convergence).

8.1 Law School Experiments

In Section 1 we partially described an experiment using the *Law School Admissions* dataset (Wightman, 1998). The dataset has 27,234 total law students, and we use only two features in all models: GPA and LSAT scores. We plot the densities of these conditioned on whether the student passed the bar or not in Fig. 9 in Appendix D. Models were trained as described above, where for the constrained model we constrained both LSAT score and undergraduate GPA to only be positive evidence, *ceterus paribus*. The two-layer neural network and gradient boosted trees models were trained with a similar number of model parameters to the GAMs and no monotonicity constraints. Resulting models were shown in Fig. 2, and train/test accuracies in Table 1.

8.2 Credit Default Experiments

Next, we consider the Default of Credit Card Clients benchmark dataset from the UCI repository (Lichman, 2013; Yeh and Lien, 2009). The data was collected from 30,000 Taiwanese credit card users and contains a binary label of whether or not a user defaulted on a payment in a time window. Features include marital status, gender, education, and how long a user is behind on payment of their existing bills, for each of the months of April-September 2005.

Repayment status is the integer number of months it has been since the user has repaid, and negative values mean the user has already repaid. Here we illustrate using monotonicity constraints to avoid *unfair penalization*: if the model were to be used to determine a user’s credit score, it could feel unfair to many if they were penalized for paying their bills sooner, all else equal. Thus, we apply a monotonicity constraint that keeps the model from penalizing early payments.

Figure 4 (top) shows the average and standard deviation of the default rate of the training examples as a

function of the months since the bills were paid. The data is very noisy for 3+ months overdue payment (there are only 122 such training examples), and these noisy averages do not follow the reasonable principle of predicting a higher default rate if your bills are more over-due.

For ease of visualization, our first experiment uses only $D = 2$ features: marital status and April repayment status, with a monotonicity shape constraint on the repayment status to ensure that paying your bills on time doesn't hurt you. Our second experiment uses all $D = 24$ features, and we impose monotonicity shape constraints on all 6 repayment features to guarantee that the person is not penalized for paying early/on-time during any of the 6 months.

Fig. 5 (*top*) shows the predicted default rate for the unconstrained GAM. It mimics the average training labels, and thus unfairly rewards people who are 5-6 months overdue on their bills with a lower defaulting score than people who are only 2-3 months overdue. However, Figure 5 (*bottom*) shows the GAM trained with a monotonicity shape constraint: as requested, it does not penalize people for paying their bills early.

Table 2 shows that adding the monotonicity constraints for the Credit Default problem had only a tiny effect on the train and test accuracy for the $D = 2$ experiments. For the $D = 24$ experiment, the test accuracy is slightly worse, but since the train accuracy was not hurt, we believe the lower test accuracy is simply due to the randomness of the sample.

8.3 Funding Proposals Experiments

Next we demonstrate the use of monotonicity constraints to favor the less fortunate. We use the dataset from the KDD Cup 2014: Predicting Excitement at DonorsChoose.org (KDD Cup, 2014). DonorsChoose.org is a platform through which teachers in K-12 schools can request funding and materials for proposed projects. Donors can search for and donate to projects. The dataset contains 619,327 examples of projects proposed by teachers. The label is a binary label that represents the outcome of the project, where the positive class means the project was deemed "exciting", a definition determined by DonorsChoose.org that includes whether the project was fully funded. Only 5.91% of examples are labeled "exciting". Figure 4 (*bottom*) shows the mean and std. dev. of the training examples for 2 of the 28 features: poverty level and number of students impacted. A machine learning model trained on this dataset could be used to rank projects to display to potential donors. Standard machine learning trained on just these two features would not favor schools at higher poverty levels, or projects

with greater impact. We compare to training with monotonicity constraints for those two principles, both on just the shown two features, and on all 28 features (but only these two features constrained).

The results in Table 3 for the AUC metric (the same metric used in the KDD Cup 2014 competition) show that adding monotonicity constraints to prefer poorer schools and greater student impact did not hurt the test AUC. Figures comparing the standard and constrained $D = 2$ models are in Appendix E.

# features	# shape constraints	Train Acc.	Test Acc.
2	0	82.07%	81.55%
2	1	82.06%	81.60%
24	0	82.44%	82.02%
24	6	82.35%	80.86%

Table 2: Credit Default experiment results, where the monotonicity shape constraints ensure the model does not penalize people for paying their bills earlier.

# features	# shape constraints	Train AUC	Test AUC
2	0	0.520	0.517
2	2	0.514	0.518
28	0	0.752	0.746
28	2	0.751	0.746

Table 3: Funding Proposals experiment results, where the monotonicity shape constraints ensure that the model gives higher scores to projects for higher poverty schools or impacting more students, all else equal.

8.3.1 Comparison to Statistical Fairness

In Section 6 we gave *theoretical* results on how monotonicity constraints can bound and affect *one-sided statistical parity* and *one-sided equal opportunity* fairness violations. In Table 4 we show empirically what happens to the one-sided (1-s) statistical parity and equal opportunity violations (defined below) for the Funding Proposals experiment, where the protected groups $j, k \in \{0, 1, 2, 3\}$ are four different *poverty levels*, which we treat as four different ordinal protected groups for calculating these one-sided statistical metrics.

Max *One-sided Statistical Parity* Violation:

$$\max_{j < k} (0, P(\hat{Y} = 1 | Z = j) - P(\hat{Y} = 1 | Z = k)) \quad (2)$$

Max *One-sided Equal Opportunity* Violation:

$$\max_{j < k} (0, P(\hat{Y} = 1 | Z = j, Y = 1) - P(\hat{Y} = 1 | Z = k, Y = 1)) \quad (3)$$

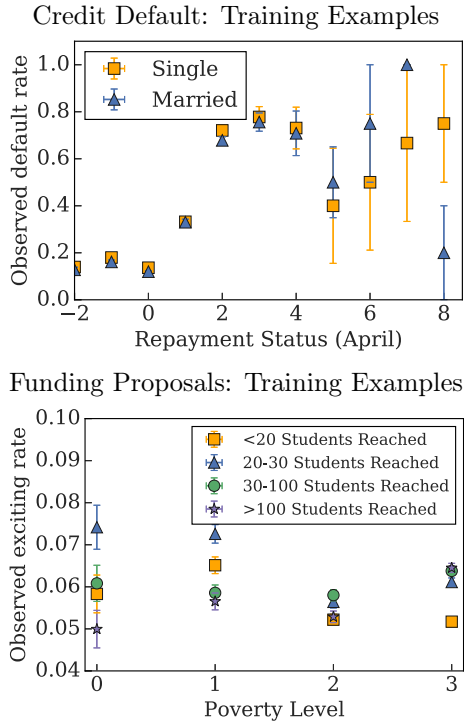



Figure 4: Mean and standard error of the label as a function of the inputs over the training dataset.

Credit Default: Unconstrained Model Predictions



Credit Default: Monotonic Model Predictions




Figure 5: Credit Default: unconstrained model predictions (*top*) and constrained model predictions (*bottom*).

Table 4 shows that the monotonically constrained models do lower violations of both fairness goals. The improvement is smaller when the model has $D = 28$ features than $D = 2$ features, which we believe is due to a weakening of the maximum likelihood ratios as described by Lemmas 3 and 4.

D	# shape constraints	Max 1-s Stat. Par. Viol. (2)	Max 1-s Eq. Opp. Viol. (3)
2	0	0.00704	0.00707
2	2	0.00017	0.00037
28	0	0.00751	0.00397
28	2	0.00261	0.00331

Table 4: Statistical Fairness Violations for Funding Proposals Experiment.

9 Discussion

We have demonstrated that nonlinear machine learned models can easily overfit noise or learn bias in a way that violates social norms or ethics about whether certain inputs should be allowed to negatively affect a score or decision. We have also shown that this problem can be ameliorated by training with monotonicity constraints to reflect the desired principle.

An advantage of monotonicity constraints is that their effect does not depend on the data distribution, so there are no questions of generalization to test data as there are with statistical fairness measures. While enforcing such constraints addresses deontological rather than consequentialist ethics, we have also shown theoretically and experimentally that monotonicity constraints can improve or bound (consequentialist) one-sided statistical fairness violations.

Monotonicity constraints are fairly easy to explain and reason about for laypeople, as illustrated in the examples in this paper. The examples also illustrate the broad applicability of constraints to situations where the model should avoid unfair penalization of good attributes, and favor the less fortunate.

We conclude that monotonicity constraints are a necessary and useful tool for creating responsible AI, but certainly not sufficient or applicable to all situations. For this method to be completely effective, all relevant features must be identified and constrained. Non-sensitive information can be highly correlated with sensitive information, causing indirect discrimination (Hajian and Domino-Ferrer, 2013), which is also a problem for statistical fairness measures (Hardt et al., 2016). This method is also not directly applicable to unordered inputs like addresses, photos, or voice signals. Thus, this will be one of many tools and strategies needed to achieve responsible AI.

Bibliography

- Agarwal, A., Beygelzimer, A., Dudík, M., Langford, J., and Wallach, H. M. (2018). A reductions approach to fair classification. In *ICML*, pages 60–69.
- Agarwal, A., Dudík, M., and Wu, Z. S. (2019). Fair regression: Quantitative definitions and reduction-based algorithms. In *ICML*.
- Allen, H. E., Latessa, E. J., and Ponder, B. S. (2015). *Corrections in America: An Introduction*. Pearson.
- Anderson, M. and Anderson, S. (2007). The status of machine ethics: A report from the AAAI Symposium. *Minds and Machines*, 17:1–10.
- Archer, N. P. and Wang, S. (1993). Application of the back propagation neural network algorithm with monotonicity constraints for two-group classification problems. *Decision Sciences*, 24(1):60–75.
- Ayer, M., Brunk, H. D., Ewing, G. M., William, W. T., and Silverman, E. (1955). An empirical distribution function for sampling with incomplete information. *The Annals of Mathematical Statistics*, 26:641–647.
- Azar, O. (2004). What sustains social norms and how they evolve? the case of tipping. *Journal of Economic Behavior and Organization*, 54:49–64.
- Barlow, R. E., Bartholomew, D. J., and Bremner, J. M. (1972). *Statistical inference under order restrictions; the theory and application of isotonic regression*. Wiley.
- Beauchamp, T. L. and Childress, J. F. (2001). Principles of biomedical ethics. In *Oxford University Press*.
- Ben-David, A. (1992). Automatic generation of symbolic multiattribute ordinal knowledge based DSS: methodology and applications. *Decision Sciences*, pages 1357–1372.
- Berk, R., Heidari, H., Jabbari, S., Joseph, M., Kearns, M., Morgenstern, J., Neel, S., and Roth, A. (2017). A convex framework for fair regression. *FATML*.
- Beutel, A., Chen, J., Doshi, T., Qian, H., Wei, L., Wu, Y., Heldt, L., Zhao, Z., Hong, L., Chi, E. H., and Goodrow, C. (2019). Fairness in recommendation through pairwise experiments. *KDD Applied Data Science Track*.
- Bickel, P. J., Hammel, E. A., and O’Connell, J. W. (1975). Sex bias in graduate admissions: Data from Berkeley. *Science*, pages 398–404.
- Binns, R. (2018). Fairness in machine learning: Lessons from political philosophy. *Conf. on Fairness, Accountability, and Transparency*.
- Bonakdarpour, M., Chatterjee, S., Barber, R. F., and Lafferty, J. (2018). Prediction rule reshaping. *Proc. ICML*.
- Canini, K., Cotter, A., Fard, M. M., Gupta, M. R., and Pfeifer, J. (2016). Fast and flexible monotonic functions with ensembles of lattices. *Advances in Neural Information Processing Systems (NIPS)*.
- Celis, L. E., Straszak, D., and Vishnoi, N. K. (2018). Ranking with fairness constraints. In *ICALP*.
- Chetverikov, D., Santos, A., and Shaikh, A. M. (2018). The econometrics of shape restrictions. *Annual Review of Economics*.
- Coate, S. and Loury, G. C. (1993). Will affirmative-action policies eliminate negative stereotypes? *The American Economic Review*, 83:1220–1240.
- Cole, G. W. and Williamson, S. A. (2019). Avoiding resentment via monotonic fairness. *arXiv preprint arXiv:1909.01251*.
- Corbett-Davies, S. and Goel, S. (2018). The measure and mismeasure of fairness: A critical review of fair machine learning. *arXiv preprint arXiv:1808.00023*.
- Cotter, A., Jiang, H., and Sridharan, K. (2019a). Two-player games for efficient non-convex constrained optimization. In *ALT*.
- Cotter, A., Jiang, H., Wang, S., Narayan, T., Gupta, M. R., You, S., and Sridharan, K. (2019b). Optimization with non-differentiable constraints with applications to fairness, recall, churn, and other goals. *Journal of Machine Learning Research (JMLR)*.
- Donini, M., Oneto, L., Ben-David, S., Shawe-Taylor, J., and Pontil, M. (2018). Empirical risk minimization under fairness constraints.
- Dwork, C., Hardt, M., Pitassi, T., Reingold, O., and Zemel, R. (2012). Fairness through awareness. In *Proc. 3rd Innovations in Theoretical Computer Science*, pages 214–226. ACM.
- Fehr, E. and Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. In *The Quarterly Journal of Economics*, volume 114, pages 817–868.
- Friedler, S. A., Scheidegger, C., Venkatasubramanian, S., Choudhary, S., Hamilton, E. P., and Roth, D. (2019). A comparative study of fairness-enhancing interventions in machine learning. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT*)*.
- Goel, N., Yaghini, M., and Faltings, B. (2018). Non-discriminatory machine learning through convex fairness criteria. *AAAI/ACM Conference on Artificial Intelligence, Ethics, Society*.
- Goh, G., Cotter, A., Gupta, M., and Friedlander, M. P. (2016). Satisfying real-world goals with dataset constraints. In *Advances in Neural Information Processing Systems*, pages 2415–2423.
- Google AI Blog (2020). Tensorflow lattice: Flexible, controlled and interpretable ml.

- Gray, K. and Graham, J. (2018). *Atlas of Moral Psychology*. Guilford Press.
- Groeneboom, P. and Jongbloed, G. (2014). *Nonparametric estimation under shape constraints*. Cambridge Press, New York, USA.
- Gruetzemacher, R. (2018). Rethinking AI strategy and policy as entangled super wicked problems. *AAAI/ACM Conference on Artificial Intelligence, Ethics, Society*.
- Gupta, M. R., Cotter, A., Pfeifer, J., Voevodski, K., Canini, K., Mangylov, A., Moczydlowski, W., and Esbroeck, A. V. (2016). Monotonic calibrated interpolated look-up tables. *Journal of Machine Learning Research*, 17(109):1–47.
- Haidt, J. (2013). *The Righteous Mind*. Random House, New York, USA.
- Hajian, S. and Domino-Ferrer, J. (2013). A methodology for direct and indirect discrimination prevention in data mining. *IEEE Trans. on Knowledge and Data Engineering*, 25(7):1445–1459.
- Hardt, M., Price, E., and Srebro, N. (2016). Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems*, pages 3315–3323.
- Hastie, T. and Tibshirani, R. (1990). *Generalized Additive Models*. Chapman Hall, New York.
- Hunter, J. E. and Schmidt, F. L. (1976). Critical analysis of the statistical and ethical implications of various definitions of test bias. *Psychological Bulletin*.
- Iserson, K. V. and Moksop, J. C. (2007). Triage in medicine: Concept, history, types. *Annals of Emergency Medicine*.
- JB, K. (1964). Nonmetric multidimensional scaling: A numerical method. *Psychometrika*, 29:115–129.
- KDD Cup (2014). Predict funding requests that deserve an A+.
- Kearns, M., Neel, S., Roth, A., and Wu, S. (2018). Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. *Proc. ICML*.
- Komiyama, J., Takeda, A., Honda, J., and Shima, H. (2018). Nonconvex optimization for regression with fairness constraints. *Proc. ICML*.
- Kotlowski, W. and Slowinski, R. (2009). Rule learning with monotonicity constraints. In *ICML*, pages 537–544.
- Kusner, M., Loftus, J., Russell, C., and Silva, R. (2017). Counterfactual fairness. *NIPS*.
- L. Bernstein (2017). *Who Deserves a Liver?* Washington Post.
- Lichman, M. (2013). UCI machine learning repository.
- Moor, J. H. (2006). The nature, importance, and difficulty of machine ethics. *IEEE Intelligent Systems*, pages 18–21.
- Pearl, J., Glymour, M., and Jewell, N. (2016). *Causal Inference in Statistics: a Primer*. Wiley.
- Pérez-Suay, A., Laparra, V., Mateo-García, G., Muñoz-Marí, J., Gómez-Chova, L., and Camps-Valls, G. (2017). Fair kernel learning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*.
- Raff, E., Sylvester, J., and Mills, S. (2018). Fair forests: Regularized tree induction to minimize model bias. *AAAI/ACM Conference on Artificial Intelligence, Ethics, Society*.
- Rawls, J. (1971). *A Theory of Justice*. Belknap Press.
- Ross, W. D. (2002). *The Right and the Good*. Oxford Press Reprint of 1930 Original, Oxford.
- Sandvig, C., Hamilton, K., Karahalios, K., and Langbort, C. (2016). When the algorithm itself is racist: diagnosing ethical harm in the basic components of software. *Intl. Journal of Communication*.
- Sill, J. and Abu-Mostafa, Y. S. (1997). Monotonicity hints. *Advances in Neural Information Processing Systems (NIPS)*, pages 634–640.
- Singer, P. (2016). *Ethics in the Real World*. Princeton University Press, New York.
- Singh, A. and Joachims, T. (2018). Fairness of exposure in rankings. In *KDD*.
- Sugden, R. (1989). Spontaneous order. In *Journal of Economic Perspectives*.
- Thiroux, J. P. and Krasemann, K. W. (2017). *Ethics Theory and Practice*. Pearson, New York.
- Wightman, L. (1998). LSAC national longitudinal bar passage study. *Law School Admission Council*.
- Yeh, I. and Lien, C. (2009). The comparisons of data mining techniques for the predictive accuracy of probability of default of credit card clients. *Expert Systems with Applications*, 36(2):2473–2480.
- You, S., Canini, K., Ding, D., Pfeifer, J., and Gupta, M. R. (2017). Deep lattice networks and partial monotonic functions. *Advances in Neural Information Processing Systems (NIPS)*.
- Zafar, M. B., Valera, I., Rodriguez, M. G., and Gummadi, K. P. (2015). Fairness constraints: A mechanism for fair classification. In *ICML Workshop on Fairness, Accountability, and Transparency in Machine Learning*.
- Zafar, M. B., Valera, I., Rodriguez, M. G., and Gummadi, K. P. (2017). Fairness constraints: Mechanisms for fair classification. In *AISTATS*.

Zehlike, M., Bonchi, F., Castillo, C., Hajian, S., Megahed, M., and Baeza-Yates, R. (2017). Fa*ir: A fair top-k ranking algorithm. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, CIKM '17*.