

Genetic programming approaches to learning fair classifiers

William La Cava*
University of Pennsylvania
Philadelphia, PA
lacava@upenn.edu

Jason H. Moore
University of Pennsylvania
Philadelphia, PA
jhmoore@upenn.edu

ABSTRACT

Society has come to rely on algorithms like classifiers for important decision making, giving rise to the need for ethical guarantees such as fairness. Fairness is typically defined by asking that some statistic of a classifier be approximately equal over protected groups within a population. In this paper, current approaches to fairness are discussed and used to motivate algorithmic proposals that incorporate fairness into genetic programming for classification. We propose two ideas. The first is to incorporate a fairness objective into multi-objective optimization. The second is to adapt lexicase selection to define cases dynamically over intersections of protected groups. We describe why lexicase selection is well suited to pressure models to perform well across the potentially infinitely many subgroups over which fairness is desired. We use a recent genetic programming approach to construct models on four datasets for which fairness constraints are necessary, and empirically compare performance to prior methods utilizing game-theoretic solutions. Methods are assessed based on their ability to generate trade-offs of subgroup fairness and accuracy that are Pareto optimal. The result show that genetic programming methods in general, and random search in particular, are well suited to this task.

CCS CONCEPTS

• **Mathematics of computing** → **Evolutionary algorithms**; • **Computing methodologies** → *Supervised learning by classification*; • **Applied computing** → Engineering;

KEYWORDS

genetic programming, pareto optimization, fairness, classification

ACM Reference Format:

William La Cava and Jason H. Moore. 2020. Genetic programming approaches to learning fair classifiers. In *Genetic and Evolutionary Computation Conference (GECCO '20), July 8–12, 2020, Cancún, Mexico*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3377930.3390157>

*Corresponding Author

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
GECCO '20, July 8–12, 2020, Cancún, Mexico
© 2020 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-7128-5/20/07.
<https://doi.org/10.1145/3377930.3390157>

1 INTRODUCTION

Machine learning (ML) models that are deployed in the real world can have serious effects on peoples' lives. In impactful domains such as lending [11], college admissions [24], criminal sentencing [3, 6], and healthcare [10, 29], there is increasing concern that models will behave in unethical ways [16]. This concern has led ML researchers to propose different measures of fairness for constraining and/or auditing classification models [8]. However, in many cases, desired notions of fairness require exponentially many constraints to be satisfied, making the problems of learning fair models, and also checking for fairness, computationally hard [14]. For this reason search heuristics like genetic programming (GP) may be useful for finding approximate solutions to these problems.

This paper is, to our knowledge, the first foray into incorporating fairness constraints into GP. We propose and study two methods for learning fair classifiers via GP-based symbolic classification. Our first proposal is a straightforward one: to add a fairness metric as an objective to multi-objective optimization [7]. This fairness metric works by defining *protected groups* within the data, which match individuals having a specific value of one protected attribute, e.g. "female" for a sex attribute. Unfortunately, simple metrics of fairness do not capture fairness over rich subgroups and/or intersections of groups - that is, over multiple protected attributes that intersect in myriad ways. With this in mind, we propose an adaptation of lexicase selection [18] designed to operate over randomized sequences of fairness constraints. This algorithm draws a connection between these numerous fairness constraints and the way in which lexicase samples fitness cases in random sequences for parent selection. We illustrate the ability of lexicase to sample the space of group intersections in order to pressure models to perform well on the intersections of groups that are most difficult in the current population. In our experiments, we compare several randomized search heuristics to a recent game-theoretic approach to capturing subgroup fairness. The results suggest that GP methods can produce Pareto-efficient trade-offs between fairness and accuracy, and that random search is a strong benchmark for doing so.

In the following section, we describe how fairness has been approached in the ML community and the challenges that motivate our study. Section 3 describes the algorithms we propose in detail, and Section 4 describes the experiment we conduct on four real-world datasets for which fairness concerns are pertinent. We present resulting measures of performance, statistical comparisons, and example fairness-accuracy trade-offs in Section 5, followed finally by a discussion of what these results entail for future studies.

2 BACKGROUND

Incorporating notions of fairness into ML is a fairly new idea [25], and early work in the field is reviewed in Chouldechova and Roth [5]. Algorithmic unfairness may arise from disparate causes, but often has to do with the properties of the data used to train a model. One major cause of bias is that data are often collected from unequal demographics of a population. In such a scenario, algorithms that minimize average error over all samples will skew towards fitting the majority population, since this leads to lower average error. One way to address this problem is to train separate models for separate demographic populations. In some scenarios, this method can reduce bias, but there are two main caveats, expounded upon in [28]. First, some application areas explicitly forbid demographic data to be used in prediction, meaning these models could not be deployed. The second, and more general, concern is that we may want to protect several sensitive features of a population (e.g., race, ethnicity, sex, income, medical history, etc.). In those cases, dividing data beforehand is non-trivial, and can severely limit the sample size used to train each model, leading to poor performance.

There is not a single agreed-upon definition of fairness for classification. The definitions put forth can be grouped into two kinds: *statistical* fairness, in which we ask a classifier to behave approximately equally on average across protected groups according to some metric; and *individual* fairness, in which we ask a classifier to perform similarly on similar pairs of individuals [8]. For this paper, we focus on statistical fairness, especially equality of false positive (FP), false negative (FN), and accuracy rates among groups. We essentially ask that the classifier’s errors be distributed among different protected groups as evenly as possible.

Fairness constraints have been proposed for classification algorithms, for example by regularization [2, 8], model calibration [11], cost-sensitive classification [1], and evolutionary multi-objective optimization [26]. For the most part, literature has focused on providing guarantees over a small number of protected groups that represent single attributes - for example, race and sex. However, a model that appears fair with respect to several individual groups may actually discriminate over specific *intersections* or conjunctions of those groups. Kearns et al. [14] refers to this issue as “fairness gerrymandering”. To paraphrase Example 1.1 of their work [14], imagine a classifier that exhibits equivalent error rates according to two protected groups: a race feature taking values in {“black”, “white”} and, separately, a sex feature taking values in {“male”, “female”}. This seemingly fair classifier could actually be producing 100% of its errors on black males and white females. In such a case the classifier would appear fair according to the individual race and sex groups, but unfair with respect to their conjunction.

If we instead wish to learn a classifier that is fair with respect to both individual groups defined over single attributes and boolean conjunctions of those groups, a combinatorial problem arises. For p protected attributes, we have to both learn and check for fairness over 2^p groups. It turns out that

the problems of auditing a classifier for fairness over boolean conjunctions of groups (as well as other group definitions) is computationally hard in the worst case, as is the classification problem [14].

Kearns et al. [14] proposed a heuristic solution to the problem of learning a classifier with rich subgroup fairness constraints by formulating it as a two-player game in which one player learns a classifier and the other learns to audit that classifier for fairness. They empirically illustrated the trade-off between fairness violations and model accuracy on four real-world problems [15]. In our study, we build upon their work by using their fairness auditor to compare performance of models on the same datasets. In their study, Kearns et al. focused on algorithmic characterization by reporting fairness and accuracy on the training samples. Conversely, we are interested in the generalization performance of the learned classification models; therefore we conduct our comparisons over cross-validated predictions, rather than reporting in-sample.

Our interest in applying GP to the problem of fair classification is motivated by three observations from this prior work. First, given that the learning and auditing problems for rich subgroup fairness are hard in the worst case means that a heuristic method such as GP may be able to provide approximate solutions with high utility, and therefore it is worth an empirical analysis. Second, many authors note the inherent trade-off that exists between fairness and accuracy [11, 15] and the need for Pareto-efficient solution sets. Multi-objective optimization methods that are typically used in GP (e.g., NSGA2 [7]) are well-suited to handle competing objectives during search. Finally, we note that demographic imbalance, one of the causes of model unfairness, is a problem due to the use of average error for guiding optimization. However, recent semantic selection methods [23] such as ϵ -lexicase selection [22] are designed specifically to move away from scalar fitness values that average error over the entire training set. The original motivation behind these GP methods is to prevent the loss of candidate models in the search space that perform well over difficult subsets of the data [22]. Furthermore, we hypothesize that ϵ -lexicase selection may be adapted to preserve models that perform well over structured subgroups of the protected attributes as well.

3 METHODS

We start with a dataset of triples, $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{x}'_i, y_i)\}_{i=1}^m$, containing m examples. Our labels $y \in \{0, 1\}$ are binary classification assignments and \mathbf{x} is a vector of d features. In addition to \mathbf{x} , we have a vector of p sensitive features, \mathbf{x}' , that we wish to protect via some fairness constraint. It is worth mentioning that for the purposes of this study, \mathbf{x} contains \mathbf{x}' , meaning that the learned classifier has access to the sensitive attribute observations in prediction; this is not always the case (e.g. [28]).

We also define protected groups \mathcal{G} , where each $g \in \mathcal{G}$ is an indicator function¹, mapping a set of sensitive features

¹We use $\mathbf{1}\{\}$ to denote indicator functions.

\mathbf{x}' to a group membership. It is useful to define a simple set of protected groups that correspond to the unique levels of each feature in \mathbf{x}' . We will call the set of these simple groups \mathcal{G}_0 . As an example, imagine we have two sensitive features corresponding to race and sex: $x'_1 \in \{\text{black, white}\}$ and $x'_2 \in \{\text{male, female}\}$. Then \mathcal{G}_0 would consist of four groups:

$$\begin{aligned} \mathcal{G}_0 = & \{g_1(\mathbf{x}') = \mathbf{1}\{x'_1 = \text{black}\}, \\ & g_2(\mathbf{x}') = \mathbf{1}\{x'_1 = \text{white}\}, \\ & g_3(\mathbf{x}') = \mathbf{1}\{x'_2 = \text{male}\}, \\ & g_4(\mathbf{x}') = \mathbf{1}\{x'_2 = \text{female}\}\} \end{aligned}$$

We make use of \mathcal{G}_0 in defining marginal fairness and in Algorithm 1.

We use a recent GP technique called FEAT [19, 21] that evolves feature sets for a linear model, in this case a logistic regression model. More details of this method are given in Section 4. As in other GP methods, FEAT trains a population of individuals, $n \in \mathcal{N}$, each of which produces binary classifications of the form $n(\mathbf{x}) \in \{0, 1\}$. The fitness of n is its average loss over the training samples, denoted $f(n)$. We refer to the fitness of n over a specific group of training samples as $f(n, g)$. With these definitions in mind, we can define the fairness of a classifier with respect to a particular group and fitness measure as:

$$f\text{-Fairness}(n, g) = |f(n) - f(n, g)| \quad (1)$$

FEAT uses logistic loss as its fitness during training, in keeping with its logistic regression pairing. However, we compare fairness on fitted models relative to the FP and FN rate, as in previous work [1, 15].

3.1 Multi-objective Approach

A straightforward way to incorporate fairness into FEAT is to add it as an objective to a multi-objective optimization algorithm like NSGA2. We use the term *marginal fairness* to refer to the first-level fairness of a model defined over simple groups \mathcal{G}_0 :

$$f\text{-Marginal Fairness}(n, \mathcal{G}_0) = \frac{1}{|\mathcal{G}_0|} \sum_{g \in \mathcal{G}_0} f\text{-Fairness}(n, g) \quad (2)$$

A challenge with using fairness as an objective is the presence of a trivial solution: a model that produces all 1 or all 0 classifications has perfect fairness, and will easily remain in the population unless explicitly removed.

A major shortcoming of optimizing Eqn. 2 is that it does not pressure classifiers to perform well over group intersections, and is therefore susceptible to fairness gerrymandering, as described in Section 2. Unfortunately, it is not feasible to explicitly audit each classifier in the population each generation over all possible combinations of structured subgroups. While an approximate, polynomial time solution has been proposed [14, 15], we consider it too expensive to compute in practice each iteration on the entire set of models. For these reasons, we propose an adaptation of lexicase selection [27] to handle this task in the following section.

3.2 Fair Lexicase Selection

Lexicase selection is a parent selection algorithm originally proposed for program synthesis tasks [13] and later regression [22]. Each parent selection event, lexicase selections filters the population through a newly randomized ordering of “cases”, which are typically training samples. An individual may only pass through one of these cases if it has the best fitness in the current pool of individuals, or alternately if it is within ϵ of the best for ϵ -lexicase selection. The filtering process stops when one individual is left (and is selected), or when it runs out of cases, resulting in random selection from the remaining pool.

Although different methods for defining ϵ have been proposed, we use the most common one, which defines ϵ as the median absolute deviation (λ) of the loss (ℓ) in the current selection pool:²

$$\lambda(\ell(n), n \in \mathcal{S}) = \text{median}(|\ell(n) - \text{median}(\ell(n))|), n \in \mathcal{S}$$

Lexicase selection has a few properties worth noting that are discussed in depth in [18]. First, it takes into account case “hardness”, meaning training samples that are very easy to solve apply very little selective pressure to the population, and vice versa. Second, lexicase selection selects individuals on the Pareto front spanned by the cases; this means that, in general, it is able to preserve individuals that only perform well on a small number of hard cases (i.e. specialists [12]). Third, and perhaps most relevant to rich subgroup fairness, lexicase selection does not require each individual to be run on each case/sample, since selection often chooses a parent before the cases have been exhausted [22]. The worst case complexity of parent selection is $O(|\mathcal{N}|^2 m)$, which only occurs in a semantically homogeneous population.

Because of the third point above, we can ask for lexicase selection to audit classifiers over conjunctions of groups without explicitly constructing those groups beforehand. Instead, in *fair lexicase* (FLEX, detailed in Alg. 1), we define “cases” to be drawn from the simple groups in \mathcal{G}_0 . A randomized ordering of these groups, i.e. cases, thereby assesses classifier performance over a conjunction of protected attributes. By defining cases in this way, selective pressure moves dynamically towards subgroups that are difficult to solve. For any given parent selection event, lexicase only needs to sample as many groups as are necessary to winnow the pool to one candidate, which is at most $|\mathcal{G}_0|$. Nonetheless, due to the conditional nature of case orderings, and the variability in case depth and orderings, lexicase effectively samples $|\mathcal{G}_0|!$ combinations of protected groups.

An illustration of three example selection events is shown in Figure 1. These events illustrate that FLEX can select on different sequences of groups and different sequence lengths, while also taking into account the easiness or hardness of the group among the current selection pool.

²Defining λ relative to the current selection pool is called “dynamic ϵ -lexicase selection” in [18].




Figure 1: Three example selection events with FLEX, with a population $\mathcal{N} = \{n_1, \dots, n_5\}$ and protected groups $\mathcal{G} = \{g_1, \dots, g_4\}$. Parent selection 1) selects on the conjunction of $g_1, g_3,$ and g_2 to select n_4 . Note that g_3 exerts no selection pressure because n_4 and n_5 both perform well on it. 2) Here a single group, $g_3,$ is enough to winnow the population to $n_3,$ which is selected. 3) Selection on g_4 and g_3 to select n_1 . Gray cases are paths that have already been visited for a given selection event.

A downside of FLEX versus the multi-objective approach is that it is not as clear how to pressure for both fairness and accuracy among cases. On one hand, selecting for accuracy uniformly over many group definitions could lead to fairness, but it may also preserve markedly unfair, and therefore undesirable, models. We address this issue by allowing both case definitions to appear with equal probability. This choice explains the random coin flip in Alg. 1.

4 EXPERIMENTS

We conduct our experiment on four datasets used in previous related work [15]. These datasets and their properties are detailed in Table 1. Each of these classification problems contain sensitive information for which one would reasonably want to assure fairness. Two of the datasets concern models for admissions decisions (Lawschool and Student); The other two are of concern for lending and credit assessment: one predicts rates of community crime (Communities), and the other attempts to predict income level (Adult). For each of these datasets we used the same cleaning procedure as this previous work, making use of their repository (available here: github.com/algowatchpenn/GerryFair).

We compared eight different modeling approaches in our study, the parameters of which are shown in Table 2. Here we briefly describe the two main algorithms that are used.

GerryFair. First, we used the “Fictitious Play” algorithm from [14, 15], trained for 100 iterations at 100 different levels of γ , which controls the trade-off between error and fairness. As mentioned earlier, GerryFair treats the problem of learning a fair classifier as a two player game in which one

Algorithm 1 : Fair ϵ -Lexicase Selection (FLEX) applied to individuals $n \in \mathcal{N}$ with loss $f(n, g)$ over protected groups $g \in \mathcal{G}_0$.

```

Selection( $\mathcal{N}, \mathcal{G}_0$ ) :
     $\mathcal{P} \leftarrow \emptyset$                                  $\diamond$  parents
    do N times:
         $\mathcal{P} \leftarrow \mathcal{P} \cup \text{GetParent}(\mathcal{N}, \mathcal{G}_0)$      $\diamond$  add selection
                                                                to  $\mathcal{P}$ 

GetParent( $\mathcal{N}, \mathcal{G}_0$ ) :
     $\mathcal{G}' \leftarrow \mathcal{G}_0$                              $\diamond$  protected
                                                                groups
     $S \leftarrow \mathcal{N}$                                  $\diamond$  selection pool
    while  $|\mathcal{G}'| > 0$  and  $|S| > 1$ :
         $g \leftarrow$  random choice from  $\mathcal{G}'$          $\diamond$  pick random
                                                                group
        if random number  $\in [0, 1] < 0.5$  then
             $\ell(n) \leftarrow f(n, g)$  for  $n \in S$      $\diamond$  loss over group
        else
             $\ell(n) \leftarrow f\text{-Fairness}(n, g)$  for  $n \in S$   $\diamond$  group fairness
             $\ell^* \leftarrow \min \ell(n)$  for  $n \in S$      $\diamond$  min fitness in
                                                                pool
             $\epsilon \leftarrow \lambda(\ell(n), n \in S)$      $\diamond$  deviation of
                                                                fitnesses
        for  $n \in S$ :
            if  $\ell(n) > \ell^* + \epsilon$  then
                 $S \leftarrow S \setminus \{n\}$          $\diamond$  filter selection
                                                                pool
             $\mathcal{G}' \leftarrow \mathcal{G}' \setminus \{g\}$      $\diamond$  remove  $g$ 
    return random choice from  $S$ 
    
```

player, the classifier, is attempting to minimize error over weighted training samples, and the other player, the auditor, is attempting to find the subgroup within the classifier’s predictions that produces largest fairness violation. The play continues for the maximum iterations or until the maximum fairness violation is less than γ . The final learned classifier is an ensemble of linear, cost-sensitive classification models. We make use of the auditor for validating the predictions of all compared models, so it is described in more detail in Section 4.1.

FEAT. Our GP experiments are carried out using the Feature Engineering Automation Tool (FEAT), a GP method in which each individual model n consists of a set of programs (i.e. engineered features) that are fed into a logistic regression model (see Figure 2). This allows FEAT to learn a feature space for a logistic regression model, where the number of features is learned via the search process. The features are comprised of continuous and boolean functions, including common neural network activation functions, as shown in Table 1 in [21]. We choose to use FEAT for this experiment because it performed well in comparison to other state-of-the-art GP methods on a battery of regression tests [19]. FEAT is also advantageous in this application to binary classification because it can be paired with logistic regression, which provides probabilistic outputs for classification. These probabilities are necessary for assessing model performance

Table 1: Properties of the datasets used for comparison.

Dataset	Source (link)	Outcome	Samples	Features	Sensitive features	Protection Types	Number of simple groups ($ \mathcal{G}_0 $)
Communities	UCI	Crime rates	1994	122	18	race, ethnicity, nationality	1563
Adult	Census	Income	2020	98	7	age, race, sex	78
Lawschool	ERIC	Bar passage	1823	17	4	race, income, age, gender	47
Student	Secondary Schools	Achievement	395	43	5	sex, age, relationship status, alcohol consumption	22

using certain measures such as the average precision score, as we will describe later in Eqn. 5.

FEAT trains models according to a common evolutionary strategy. This strategy begins with the construction of models, followed by selection for parents. The parents are used to produce offspring via mutation and crossover. Depending on the method used, parents and offspring may then compete in a survival step (as in NSGA2), or the offspring may replace the parents (LEX, FLEX). For further details of FEAT we refer the reader to [20] and to the github project (github.com/lacava/feat).

We test six different selection/survival methods for FEAT, shown in Table 2. FLEX-NSGA2 is a hybrid of FLEX and NSGA2 in which selection for parents is conducted using FLEX and survival is conducted using the survival step of NSGA2. Each GP method was trained for 100 generations with a population of 100, except for Random, which returned the initial population. These parameters were chosen to approximately match those of GerryFair, and to produce the same number of final models (100). However, since the GP methods are population-based, they train 100 models per generation (except Random). GerryFair only trains two models per iteration (the classifier and the auditor); thus, at a first approximation we should expect the GP models aside from Random to require roughly 50 times more computation.

In our experiments, we run 50 repeat trials of each method on each dataset, in which we split the data 50/50 into training and test sets. For each trial, we train models by each method, and then generate predictions on the test set over each returned model. Each trial is run on a single core in a heterogeneous cluster environment, consisting mostly of 2.6GHz processors with a maximum of 8 GB of RAM.

There are inherent trade-offs between notions of fairness and accuracy that make it difficult to pick a definitive metric by which to compare models [17]. We compute several metrics of comparison, defined below.

4.1 Auditing Subgroup Fairness

In order to get a richer measure of subgroup fairness for evaluating classifiers, Kearns et al. [15] developed an auditing method that we employ here for validating classifiers. The auditor uses cost-sensitive classification to estimate the group that most badly violates a fairness measure they propose, which we refer to as a subgroup FP- or FN- Violation. We

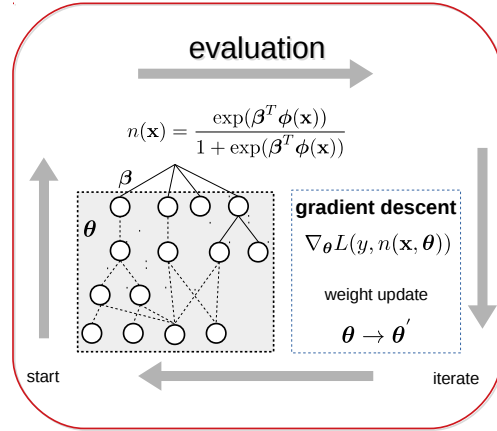


Figure 2: Diagram of the evaluation of a single FEAT individual, which produces a logistic regression model over program outputs ϕ . The internal weights θ are trained via gradient descent each generation for a set number of iterations.

can define this relative to FP rates as

$$\begin{aligned}
 \alpha_{FP}(g, \mathcal{P}) &= Pr_{\mathcal{P}}[g(\mathbf{x}') = 1, y = 0] \\
 \beta(n, g) &= |FP(n) - FP(g)| \\
 FP\text{-Violation}(n, g, \mathcal{P}) &= \alpha_{FP}(g) \beta(n, g, \mathcal{P})
 \end{aligned}
 \tag{3}$$

here, \mathcal{P} is the distribution from which the data \mathcal{D} is drawn. In Eqn. 3, $\alpha_{FP}(g, \mathcal{P})$ is estimated by the fraction of samples group g covers, so that larger groups are more highly weighted. β measures fairness equivalently to Eqn. 1. This metric can be defined equivalently for FN subgroup violations, and we report both measures in our experiments. The auditing algorithm’s objective is to return an estimate of the group g with the highest FP- or FN-Violation, and this violation is used as a measure of classifier unfairness.

4.2 Measures of Accuracy

In order to compare the accuracy of the classifiers, we used two measures. The first is accuracy, defined as

$$Accuracy(n) = \frac{1}{m} \sum_i^m \mathbf{1}[n(\mathbf{x}_i) = y_i]
 \tag{4}$$

Table 2: Settings for the methods in the experiments.

Algorithm	Settings
GerryFair [14]	iterations=100, $\gamma=100$ values $\in [0.001, \dots, 1]$, ml = logistic regression
- GerryFairGB	“”, ml = gradient boosting
FEAT [21]	generations=100, pop size=100, max depth=6, max dim=20
- Tourn	selection: size 2 tournament selection
- LEX [22]	selection: ϵ -lexicase selection
- FLEX (Alg. 1)	selection: Fair ϵ -lexicase selection
- NSGA2 [7]	NSGA2 selection and survival
- FLEX-NSGA2	selection: ϵ -lexicase selection, survival: NSGA2
- Random	return initial random population

The second is average precision score³, which is the mean precision of the model at different classification thresholds, t . APS is defined as

$$APS(n) = \sum_t (R_t(n) - R_{t-1}(n))P_t(n) \quad (5)$$

where $R(n) = Pr[n = 1, y = 1] / Pr[y = 1]$ is the recall and $P(n) = Pr[n = 1, y = 1] / Pr[n = 1]$ is the precision of $n(\mathbf{x})$.

4.3 Comparing Accuracy-Fairness Trade-offs

It is well known that there is a fundamental trade-off between the different notions of fairness described here and classifier accuracy [3, 11, 17]. For this reason, recent work has focused on comparing the Pareto front of solutions between methods [15]. For GerryFair, this trade-off is controlled via the parameter γ described in Table 2. For the GP methods, we treat the final population as the solution set to be evaluated.

In order to compare sets of solutions between methods, we compute the hypervolume of the Pareto front [9] between competing pairs of accuracy objectives (Accuracy, APS) and fairness objectives (FP Subgroup Violation, FN Subgroup Violation). This results in four hypervolume measures of comparison. For two objectives, the hypervolume provides an estimate of the area of the objective space that is covered/dominated by a set of solutions. Thus, the hypervolume allows us to compare how well each method is able to characterize the fairness-accuracy trade-off [4].

5 RESULTS

In Figure 3, we show the distributions of the hypervolume of the FP violation-APS Pareto front across trials and problems for each method. Each subplot shows the test results for each method on a single dataset, with larger values indicating better performance. In general, we observe that the GP-based approaches do quite well compared to GerryFair in

³This is a pessimistic version of estimating area under the precision-recall curve. See https://scikit-learn.org/stable/modules/generated/sklearn.metrics.average_precision_score.html.

terms of finding good trade-offs along the Pareto front. Every GP variant generates a higher median hypervolume measure than GerryFair and GerryFairGB on every problem.

Among GP variants, we observe that Random, LEX and FLEX tend to produce the highest hypervolume measures. Random search works best on the Communities and Student datasets; LEX performs best on Adult, and there is a virtual tie between Random, LEX and FLEX on LawSchool. NSGA2, FLEX-NSGA2 and Tourn all perform similarly and generally worse than Random, LEX and FLEX.

The hypervolume performance results are further summarized across problems in Figure 4. Here, each subplot shows the distribution of rankings according to a different hypervolume measurement, shown on the y axis. The significance of pairwise Wilcoxon tests between methods are shown as asterisks between bar plots. Since all pairwise comparisons are cumbersome to show, the complete pairwise Wilcoxon tests for FP Violation-APS hypervolume are shown in Table 3, corresponding to the bottom right subplot of Figure 4.

In general, the differences in performance between methods are significant. We observe that Random search, which has the best rankings across hypervolume measures, significantly outperforms all methods but LEX across problems. LEX and FLEX are significantly different only by one comparison, and the effect size is noticeably small. In addition, Tourn and NSGA2 are not significantly different, while NSGA2 and FLEX-NSGA2 are significantly different for two of the four measures.

Since the hypervolume measures only give a coarse grained view of what the Pareto fronts of solutions look like, we plot the Pareto fronts of specific trials of each method on two problems in Figures 5 and 6. The first figure shows results for the Adult problem, and presents a typical solution set for this problem. It’s noteworthy that, despite having 100 models produced by each method, only a fraction of these models produce Pareto-efficient sets on the test data. The small numbers of Pareto optimal models under test evaluation suggest that most classifiers are overfit to the training data to some degree, in terms of error rate, unfairness, or both. We also find it interesting that the combined front of solutions to this problem samples includes models from six different methods. In this way we see the potential for generating complimentary, Pareto-optimal models from distinct methods.

By contrast, models for the Student dataset shown in Figure 6 are dominated by one method: Random search. Random produces high hypervolume measures for this problem compared to other methods, and the Pareto fronts in this figure shows an example: in this case, Random is able to find three Pareto-optimal classifiers with very low error (high APS) and very low unfairness. These three models dominate all other solutions found by other methods.

Each method is evaluated on a single core, and the wall clock times of these runs are shown in Figure 7. Random is the quickest to train, followed by the two GerryFair variants. Compared to the generational GP methods, GerryFair exhibits runtimes that are between 2 and 5 times faster. Interestingly, the NSGA2 runs finish most quickly among the GP




Figure 3: Normalized hypervolume of the Pareto front for test values of FP violation and average precision score.

methods. This suggests that NSGA2 may be biased toward smaller models during optimization.

6 DISCUSSION AND CONCLUSION

The purpose of this work is to propose and evaluate methods for training fair classifiers using GP. We proposed two main ideas: first, to incorporate a fairness objective into NSGA2, and second, to modify lexicase selection to operate over subgroups of the protected attributes at each case, rather than on the raw samples. We evaluated these proposals relative to baseline GP approaches, including tournament selection, lexicase selection, and random search, and relative to a game-theoretic approach from the literature. In general we found that the GP-based methods perform quite well, in terms of the hypervolume dominated by the Pareto front of accuracy-fairness trade-offs they generate. An additional advantage of this family of methods is that they may generate intelligible models, due to their symbolic nature. However, the typical evolutionary strategies used by GP did not perform significantly better than randomly generated models, except for one tested problem.

Our first idea, to incorporate a marginal fairness objective into NSGA2, did not result in model sets that were better than tournament selection. This suggests that the marginal fairness objective (Eq. 2) does not, in and of itself, produce model sets with better subgroup fairness (Eq. 3). An obvious next step would be to incorporate the auditor (Section 4.1) into NSGA2 in order to explicitly minimize the subgroup fitness violation. The downside to this is its computational

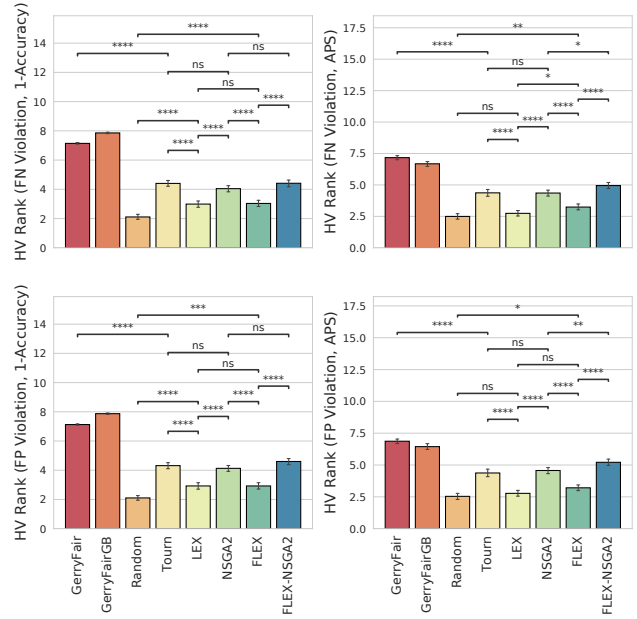


Figure 4: Rankings of methods by four different hypervolume (HV) measurements, across all problems. Asterisks denote statistical comparisons, conducted by a corrected pairwise Wilcoxon test. ns: $5e - 02 < p \leq 1.0$; *: $1e - 02 < p \leq 5e - 02$; **: $1e - 03 < p \leq 1e - 02$; ***: $1e - 04 < p \leq 1e - 03$; ****: $p \leq 1e - 04$.

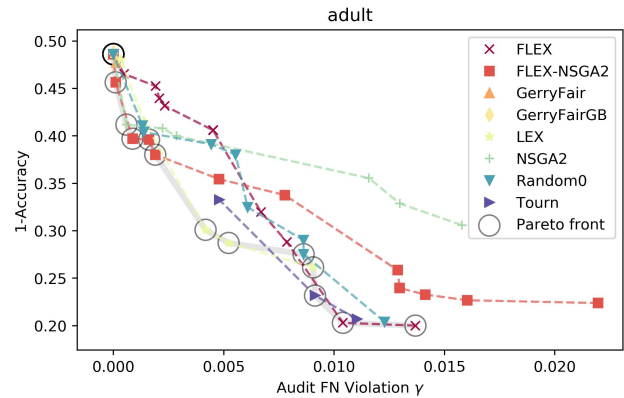


Figure 5: An example Pareto front of error (1-Accuracy) and unfairness (Audit FN Violation) based on test predictions on the adult dataset. The test set Pareto fronts for each method are plotted separately with dotted lines. The combined Pareto front is circled, and consists of models from six different methods in this case.

Table 3: Bonferroni-adjusted p -values using a Wilcoxon signed rank test of (FP-Violation, APS) hypervolume scores for the methods across all problems. Bold: $p < 0.05$.

	FLEX	FLEX-NSGA2	GerryFair	GerryFairGB	LEX	NSGA2	Random
FLEX-NSGA2	2.6e-16						
GerryFair	1.3e-30	7.4e-15					
GerryFairGB	4.2e-27	9.1e-10	3.1e-02				
LEX	1.5e-01	1.7e-20	7.9e-32	9.5e-27			
NSGA2	1.7e-09	5.7e-03	5.6e-21	1.9e-14	2.1e-13		
Random	1.7e-02	1.3e-22	5.9e-32	3.5e-27	1.0e+00	3.7e-18	
Tourn	3.5e-06	2.3e-03	5.2e-20	7.3e-16	3.7e-11	1.0e+00	2.4e-12

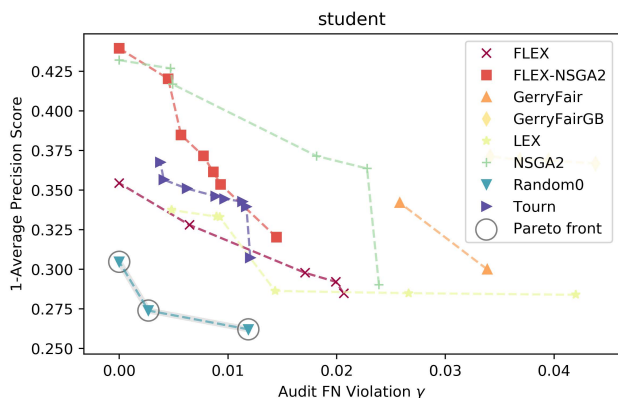


Figure 6: An example Pareto front of error (APS) and unfairness (Audit FN Violation) based on test predictions on the student dataset. The test set Pareto fronts are plotted for each method separately with dotted lines. The combined Pareto front is circled, and consists of three models generated by random search that dominate all other models.




Figure 7: Wall clock runtime comparisons for all methods across all datasets.

complexity, since it would require an additional iteration of model training per individual per generation.

Our proposal to modify lexibase selection by regrouping cases in order to promote fairness over subgroups (FLEX) did not significantly change the performance of lexibase selection. It appeared to improve performance on one dataset

(adult), worsen performance on another (communities), and overall did not perform significantly differently than LEX. This comparison is overshadowed by the performance of random search over these datasets, which gave comparable, and occasionally better, performance than LEX and FLEX for a fraction of the computational cost.

In light of these results, we want to understand why random search is so effective on these problems. There are several possible avenues of investigation. For one, the stability of the unfairness estimate provided by the auditor on training and test sets should be understood, since the group experiencing the largest fairness violation may differ between the two. This difference may make the fairness violation differ dramatically on test data. Unlike typical uses of Pareto optimization in GP literature that seek to control some static aspect of the solution (e.g., its complexity), in application to fairness, the risk of overfitting exists for both objectives. Therefore the robustness of Pareto optimal solutions may suffer. In addition, the study conducted here considered a small number of small datasets, and it is possible that a larger number of datasets would reveal new and/or different insights.

The field of fairness in ML is nascent but growing quickly, and addresses very important societal concerns. Recent results show that the problems of both learning and auditing classifiers for rich subgroup fairness are computationally hard in the worst case. This motivates the analysis of heuristic algorithms such as GP for building fair classifiers. Our experiments suggest that GP-based methods can perform competitively with methods designed specifically for handling fairness. We hope that this motivates further inquiry into incorporating fairness constraints into models using randomized search heuristics such as evolutionary computation.

7 SUPPLEMENTAL MATERIAL

The code to reproduce our experiments is available from https://github.com/lacava/fair_gp.

8 ACKNOWLEDGMENTS

The authors would like thank the Warren Center for Data Science and the Institute for Biomedical Informatics at Penn for their discussions. This work is supported by National Institutes of Health grants K99 LM012926-02, R01 LM010098 and R01 AI116794.

REFERENCES

- [1] Alekh Agarwal, Alina Beygelzimer, Miroslav Dudik, John Langford, and Hanna Wallach. 2018. A Reductions Approach to Fair Classification. In *International Conference on Machine Learning*. 60–69. <http://proceedings.mlr.press/v80/agarwal18a.html>
- [2] Richard Berk, Hoda Heidari, Shahin Jabbari, Matthew Joseph, Michael Kearns, Jamie Morgenstern, Seth Neel, and Aaron Roth. 2017. A convex framework for fair regression. *arXiv preprint arXiv:1706.02409* (2017).
- [3] Richard Berk, Hoda Heidari, Shahin Jabbari, Michael Kearns, and Aaron Roth. 2018. Fairness in Criminal Justice Risk Assessments: The State of the Art. *Sociological Methods & Research* (July 2018), 004912411878253. <https://doi.org/10.1177/0049124118782533>
- [4] Shelvin Chand and Markus Wagner. 2015. Evolutionary many-objective optimization: A quick-start guide. *Surveys in Operations Research and Management Science* 20, 2 (Dec. 2015), 35–42. <https://doi.org/10.1016/j.sorms.2015.08.001>
- [5] Alexandra Chouldechova and Aaron Roth. 2018. The Frontiers of Fairness in Machine Learning. *arXiv:1810.08810 [cs, stat]* (Oct. 2018). <http://arxiv.org/abs/1810.08810> arXiv: 1810.08810.
- [6] Sam Corbett-Davies, Emma Pierson, Avi Feller, Sharad Goel, and Aziz Huq. 2017. Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 797–806.
- [7] Kalyanmoy Deb, Samir Agrawal, Amrit Pratap, and T Meyarivan. 2000. A Fast Elitist Non-dominated Sorting Genetic Algorithm for Multi-objective Optimization: NSGA-II. In *Parallel Problem Solving from Nature PPSN VI*, Marc Schoenauer, Kalyanmoy Deb, Günther Rudolph, Xin Yao, Evelyne Lutton, Juan Julian Merelo, and Hans-Paul Schwefel (Eds.). Vol. 1917. Springer Berlin Heidelberg, Berlin, Heidelberg, 849–858. <http://repository.ias.ac.in/83498/>
- [8] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*. ACM, 214–226.
- [9] Carlos M. Fonseca, Luís Paquete, and Manuel López-Ibáñez. 2006. An improved dimension-sweep algorithm for the hypervolume indicator. In *2006 IEEE international conference on evolutionary computation*. IEEE, 1157–1163.
- [10] Milena A. Gianfrancesco, Suzanne Tamang, Jinoos Yazdany, and Gabriela Schmajuk. 2018. Potential biases in machine learning algorithms using electronic health record data. *JAMA internal medicine* 178, 11 (2018), 1544–1547.
- [11] Moritz Hardt, Eric Price, and Nathan Srebro. 2016. Equality of Opportunity in Supervised Learning. (Oct. 2016). <https://arxiv.org/abs/1610.02413v1>
- [12] Thomas Helmuth, Edward Pantridge, and Lee Spector. 2019. Lexicase selection of specialists. In *Proceedings of the Genetic and Evolutionary Computation Conference*. 1030–1038.
- [13] T. Helmuth, L. Spector, and J. Matheson. 2014. Solving Uncompromising Problems with Lexicase Selection. *IEEE Transactions on Evolutionary Computation* PP, 99 (2014), 1–1. <https://doi.org/10.1109/TEVC.2014.2362729>
- [14] Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. 2017. Preventing Fairness Gerrymandering: Auditing and Learning for Subgroup Fairness. *arXiv:1711.05144 [cs]* (Nov. 2017). <http://arxiv.org/abs/1711.05144> arXiv: 1711.05144.
- [15] Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. 2018. An Empirical Study of Rich Subgroup Fairness for Machine Learning. *arXiv:1808.08166 [cs, stat]* (Aug. 2018). <http://arxiv.org/abs/1808.08166> arXiv: 1808.08166.
- [16] Michael Kearns and Aaron Roth. 2019. *The Ethical Algorithm: The Science of Socially Aware Algorithm Design*. Oxford University Press.
- [17] Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. 2016. Inherent trade-offs in the fair determination of risk scores. *arXiv preprint arXiv:1609.05807* (2016).
- [18] William La Cava, Thomas Helmuth, Lee Spector, and Jason H. Moore. 2018. A probabilistic and multi-objective analysis of lexicase selection and ϵ -lexicase selection. *Evolutionary Computation* (May 2018), 1–28. <https://doi.org/10.1162/evco.a.00224>
- [19] William La Cava and Jason H. Moore. 2019. Semantic variation operators for multidimensional genetic programming. In *Proceedings of the 2019 Genetic and Evolutionary Computation Conference (GECCO '19)*. ACM, Prague, Czech Republic. <https://doi.org/10.1145/3321707.3321776> arXiv: 1904.08577.
- [20] William La Cava and Jason H. Moore. 2020. Learning feature spaces for regression with genetic programming. *Genetic Programming and Evolvable Machines* (March 2020). <https://doi.org/10.1007/s10710-020-09383-4>
- [21] William La Cava, Tilak Raj Singh, James Taggart, Srinivas Suri, and Jason H. Moore. 2019. Learning concise representations for regression by evolving networks of trees. In *International Conference on Learning Representations (ICLR)*. <https://arxiv.org/abs/1807.00981>
- [22] William La Cava, Lee Spector, and Kourosh Danai. 2016. Epsilon-Lexicase Selection for Regression. In *Proceedings of the Genetic and Evolutionary Computation Conference 2016 (GECCO '16)*. ACM, New York, NY, USA, 741–748. <https://doi.org/10.1145/2908812.2908898>
- [23] Pawel Liskowski, Krzysztof Krawiec, Thomas Helmuth, and Lee Spector. 2015. Comparison of Semantic-aware Selection Methods in Genetic Programming. In *Proceedings of the Companion Publication of the 2015 Annual Conference on Genetic and Evolutionary Computation (GECCO Companion '15)*. ACM, New York, NY, USA, 1301–1307. <https://doi.org/10.1145/2739482.2768505>
- [24] Frank Marcinkowski, Kimon Kieslich, Christopher Starke, and Marco Lünich. 2020. Implications of AI (un-)fairness in higher education admissions: the effects of perceived AI (un-)fairness on exit, voice and organizational reputation. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT* '20)*. Association for Computing Machinery, Barcelona, Spain, 122–130. <https://doi.org/10.1145/3351095.3372867>
- [25] Dino Pedreshi, Salvatore Ruggieri, and Franco Turini. 2008. Discrimination-aware data mining. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*. 560–568.
- [26] Novi Quadrianto and Viktoriia Sharmanska. 2017. Recycling privileged learning and distribution matching for fairness. In *Advances in Neural Information Processing Systems*. 677–688.
- [27] Lee Spector. 2012. Assessment of problem modality by differential performance of lexicase selection in genetic programming: a preliminary report. In *Proceedings of the fourteenth international conference on Genetic and evolutionary computation conference companion*. 401–408. <http://dl.acm.org/citation.cfm?id=2330846>
- [28] Philip S. Thomas, Bruno Castro da Silva, Andrew G. Barto, Stephen Giguere, Yuriy Brun, and Emma Brunskill. 2019. Preventing undesirable behavior of intelligent machines. *Science* 366, 6468 (Nov. 2019), 999–1004. <https://doi.org/10.1126/science.aag3311>
- [29] Anna Zink and Sherri Rose. 2019. Fair Regression for Health Care Spending. *arXiv:1901.10566 [cs, stat]* (Jan. 2019). <http://arxiv.org/abs/1901.10566> arXiv: 1901.10566.