

Active Fairness in Algorithmic Decision Making

Anonymous Authors*

Abstract

Society increasingly relies on machine learning models for automated decision making. Yet, efficiency gains from automation have come paired with concern for algorithmic discrimination that can systematize inequality. Recent work has proposed optimal post-processing methods that randomize classification decisions for a fraction of individuals, in order to achieve fairness measures related to parity in errors and calibration. These methods, however, have raised concern due to the information inefficiency, intra-group unfairness, and Pareto sub-optimality they entail. The present work proposes an alternative *active* framework for fair classification, where, in deployment, a decision-maker adaptively acquires information according to the needs of different groups or individuals, towards balancing disparities in classification performance. We propose two such methods, where information collection is adapted to group- and individual-level needs respectively. We show on real-world datasets that these can achieve: 1) calibration and single error parity (e.g., *equal opportunity*); and 2) parity in both false positive and false negative rates (i.e., *equal odds*). Moreover, we show that by leveraging their additional degree of freedom, *active* approaches can substantially outperform randomization-based classifiers previously considered optimal, while avoiding limitations such as intra-group unfairness.

1 Introduction

As automated decision-making systems (ADMs) have become increasingly ubiquitous—e.g., in criminal justice (Kleinberg, Mullainathan, and Raghavan 2016), medical diagnosis and treatment (Kleinberg et al.), human resource management (Chalfin et al.), social work (Gillingham 2015), credit (Huang, Chen, and Wang 2007), and insurance (Siegel 2013)—there is widespread concern about how these can deepen social inequalities and systematize discrimination. Consequently, substantial work on defining and optimizing for algorithmic fairness has surged in the last few years.

Inspired by domains such as race biases in criminal risk predictions (Flores, Bechtel, and Lowenkamp 2016), a substantial body of literature has focused on the problem of balancing classification errors across protected population subgroups, towards achieving equal false positive rates, false negative rates, or both (*equal odds*). To that end, recent research has proposed “optimal” post-processing methods

that randomize decisions of a fraction of individuals to attain *group fairness* (Hardt et al. 2016; Pleiss et al. 2017). Yet, strong limitations of randomized approaches have been noted, such as information wastefulness, Pareto sub-optimality, and intra-group unfairness (Hardt et al. 2016; Pleiss et al. 2017; Corbett-Davies and Goel 2018).

Our work aims at overcoming such limitations. We propose a complementary approach, *active fairness*, where, in deployment, an ADM adaptively collects information (features) about decision subjects; gathering more information about groups or individuals harder to classify, towards achieving equity in predictive performance. Thereby, the approach leverages a natural affordance of many real-world decision systems—adaptive information collection—and allocates an ADM’s information budget according to group- or individual-level needs.

Summary of contributions. We propose two methods for achieving fairness, based on group-level and individual-level budgets. We show that, without resorting to randomization, these methods are able to achieve: a) calibration and a single-error parity constraint, and b) parity in both false positive and false negative rates (i.e., *equal odds*). We show in four real-world datasets that, with constrained information budgets, active approaches can substantially outperform randomized approaches previously considered optimal (lower false positive and false negative rates). Finally, we show that classifiers using individual-level budgets in combination with active inquiry tend to dominate classifiers that use group-level budget constraints.

Intuition and motivating contexts. Consider a patient entering a hospital seeking diagnosis, typically undergoing a progressive inquiry—measuring vitals, procuring lab tests, specialists’ opinions, etc. At each step, absent sufficient certainty, the inquiry continues. Intuitively, a fair health system allocates resources to provide all patients similar-quality diagnoses. Likewise, active inquiry under cost constraints underlies contexts like disaster response, poverty mapping, homeland security, recruitment, telemedicine, refugee status determination, credit and insurance pricing, etc.

Problem formulation. Let X be an $n \times d$ feature matrix. Let $X^{(q)} \subset X$ denote a query on a subset of features in X , with $q \subset \{0, \dots, d\}$, and $x_i^{(q)}$ the partial feature vector of individual i ; and let $f(X^{(q)})$ be a predictor of class probability $P\{Y = 1|X^{(q)}\}$. We study the classification context where a decision-maker can choose what information to collect about each decision subject, and seeks to maximize accuracy and fairness under an information budget constraint $\bar{b} = \frac{1}{n} \sum_{\forall i} b_i < b_{max}$, where $b_i = |q_i| \in [0, d]$ is the

amount of information collected for individual i .

Although this setting is natural to many real-world decision systems, its affordances and implications to algorithmic decision systems—at the intersection of accuracy, fairness, and cost-efficiency—have not been thoroughly studied. Here, we focus on contexts with constant costs across features. Yet we note that the active fairness framework allows generalizations to contexts with varied costs across features, as well as richer and context-specific utility functions with potential costs to decision-subjects, such as monetary, opportunity, or privacy costs.

2 Related Work

Active feature acquisition (AFA). Several probabilistic and non-probabilistic methods exist for sequential feature querying under budget constraints (Gao and Koller ; Liu et al.), dating at least back to (MacKay 1992), and applied in domains such as medical diagnosis (Gorry and Barnett 1968), customer targeting (Kanani and Melville 2008), and image classification (Gao and Koller). To the best of our knowledge, this work is the first to study the implications AFA has to the algorithmic fairness literature and policy debate. Here, we use an approach based on probabilistic random forests, but more sophisticated methods can be used, for example, for dealing with domains with very high-dimensional input like medical images (Trapeznikov and Saligrama 2013).

Active learning. Similar to the general AFA setting, this paper assumes that a fixed set of training data is used and that incremental features of a test sample can be queried. This differs from the active learning setting in which the system actively queries training examples that optimize learning e.g. by balancing exploration and exploitation or maximizing the expected model change (Settles 2012). We foresee future work studying synergies in systems that attain fairness by actively choosing training samples using active learning while also applying AFA at test time.

Notions of fairness. Several notions of fairness and their corresponding formalizations have been proposed, most of which require that statistical properties hold across two or more population subgroups. *Demographic* or *statistical parity* requires that decision rates are independent from group membership (Calders, Kamiran, and Pechenizkiy ; Zafar et al. 2015; Louizos et al. 2015), such that $P\{\hat{Y} = 1|A = 0\} = P\{\hat{Y} = 1|A = 1\}$, for the case of binary classification and a sensitive attribute $A \in \{0, 1\}$. Most recent work focuses on meritocratic notions of fairness, or *error rate matching* (Hardt et al. 2016; Bechavod and Ligett 2017), such as requiring population subgroups to have equal false positive rates (FPR), equal false negative rates (FNR), or both, i.e., $P\{\hat{Y} = 1|A = 0, Y = y\} = P\{\hat{Y} = 1|A = 1, Y = y\}$, $y \in 0, 1$. In this work we focus on the latter set of fairness notions, although generalizations to others, such as statistical parity, are possible. Refer to (Žliobaitė) for a survey on computational measures of fairness.

Achieving Equal Opportunity and Equal Odds

Hardt et al. 2016 proposes parity in FNRs and/or parity in

FPRs as a measure of unfair discrimination across population subgroups (Hardt et al. 2016). Parity in both types of error is referred to as *equal odds*, and its relaxation, equality in only FPRs, is conceptualized as *equal opportunity*, as in contexts of positive classification it means that subjects within the positive class have an equal probability of being correctly classified positive, regardless of group membership.

Equal opportunity can be achieved simply by shifting up or down the decision threshold t_A —where $\hat{Y} = \mathbb{1}_{[t_A, 1]}$ ($\hat{P}(Y = 1|X, A)$)—for group A or A^c . Yet, doing so also directly affects FNRs, impeding achievement of *equal odds*. In this context, Hardt et al. 2016 propose a classifier that balances both FPRs and FNRs, based on naive randomization of a fraction of individuals in the advantaged group; and prove conditions under which the classifier is optimal with respect to accuracy (Hardt et al. 2016).

Although effective in achieving *equal odds*, these randomization-based results have been considered discouraging for reasons outlined below, and, as shown in Section 6, are outperformed by active approaches.

Achieving Calibration and Error Parity

In many real-world uses of algorithms for risk estimation, it is common practice to require that predictions are *calibrated*—e.g., in recidivism (Flores, Bechtel, and Lowenkamp 2016; Corbett-Davies et al. 2017), child maltreatment hotlines (Gillingham 2015; Chouldechova et al. 2018), and credit risk assessments (Huang, Chen, and Wang 2007). A *calibrated* estimator is one where, if we look at the subset of people who receive any given probability estimate $p \in [0, 1]$, we find indeed a p fraction of them to be positive instances of the classification problem. In the context of credit assignment, for example, we would expect a p fraction of credit applicants with estimated default risk of p to default. Moreover, in the context of algorithmic fairness across population groups, it is desired that calibration holds for each group (Flores, Bechtel, and Lowenkamp 2016).

Calibration is not necessary nor sufficient to achieve parity in classification errors (Corbett-Davies and Goel 2018). However, it is particularly desirable in cases where the output of an algorithm is not directly a decision but used as input to the subsequent judgment of a human decision-maker. In such contexts, risk estimates of an uncalibrated algorithm would carry a different meaning for different groups (e.g., African-American and white defendants), and hence their use in informing human judges’ decisions would likely entail disparate impact.

Recently, Kleinberg et al. 2016 demonstrated that a tension exists between minimizing error disparity across different population groups and maintaining calibrated probability estimates (Kleinberg, Mullainathan, and Raghavan 2016). In particular, it showed that calibration is compatible only with a single error constraint (i.e. equal FNR or equal FPR). On the same vein, Pleiss et al. 2017 showed that the results hold for even a strong relaxation of *equal odds*, named *equal cost*, where FPRs and FNRs are allowed to compensate one another according to a cost function (Pleiss et al. 2017). Finally, they propose a method that, using naive randomiza-

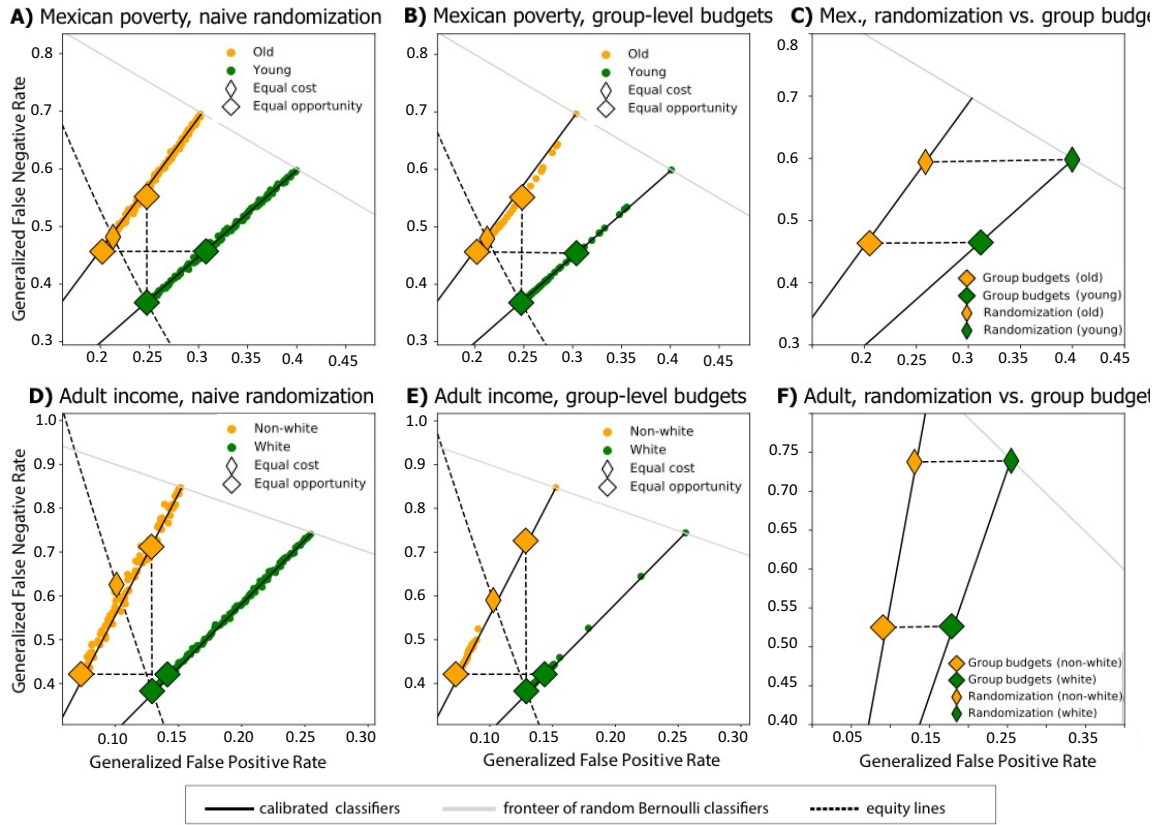


Figure 1: Achieving calibration and single error rate parity: classifiers with group-level information budgets vs. naive randomization. Rows correspond to analysis on two different datasets: the *Mexican poverty* and *adult income* datasets. Green and yellow colors correspond to error rates for two population subgroups (e.g., white and non-white individuals). Solid black lines represent the space of calibrated classifiers. Panels in the first column (A and D) show the generalized false positive and false negative rates (GFPR and GFNR) of classifiers that randomize an increasing proportion of individuals (0% to 100%), as in (Pleiss et al. 2017). In line with (Pleiss et al. 2017), naive randomization is able to achieve calibration and any single error parity constraint. Panels in the second column (B and E) show the same analysis for classifiers with group-level budgets. These classifiers are effective in achieving parity on either false positive or false negative rates, or *equal cost*, while maintaining calibration; yet without resorting to naive randomization. Finally, panels in the third column (C and F) compare the efficiency of both methods, by showing the best classifiers that achieve *equal opportunity* under an information budget restriction $\bar{b} = \frac{1}{n} \sum_{\forall i} b_i < b_{max}$. Classifiers with group-level budgets Pareto-dominate randomized classifiers by a wide margin, i.e. for the same information budget, both population subgroups are better off, incurring substantially lower false positive and false negative errors.

tion, is able to achieve parity on either error rate or *equal cost*. We compare our methods to these benchmarks in Section 5.

Objections to naive randomization

The above results on achieving *equal odds*, as well as on jointly achieving calibration and a single error parity measure, rely on naive randomization as means to fairness. Hence, they have been interpreted as unintuitive, discouraging, and unsettling (Hardt et al. 2016; Pleiss et al. 2017; Corbett-Davies and Goel 2018). Several objections have been put forth against the use of naive randomization to achieve classification parity. Among them:

Inefficiency. As pointed out by (Hardt et al. 2016; Pleiss et al. 2017; Corbett-Davies and Goel 2018), it is inefficient and

appears unintuitive to withhold information that is already in hand, by naively randomizing the classification of a subset of individuals.

Individual unfairness. Classifiers based on naive randomization, such as those in (Hardt et al. 2016; Pleiss et al. 2017; Corbett-Davies and Goel 2018), entail intra-group unfairness. Individuals who are randomized are not necessarily those with higher uncertainty but simply the ones who were unlucky, hence breaking ordinality between the probability of classification error and the underlying uncertainty.

Pareto sub-optimality and undesirability. Consider an unconstrained and unfair classifier \hat{Y}_U , which incurs higher errors on group A than group B ; and consider an alternative "fair" classifier \hat{Y}_F , where a percentage of individu-

als of group B are randomized to achieve parity in errors. Considering groups A and B as the system’s stakeholders, we note that the original unfair classifier \hat{Y}_U Pareto dominates the fair alternative \hat{Y}_F , i.e.: the disadvantaged group A will be indifferent, as its classification remained unchanged, while group B will strongly prefer \hat{Y}_U , the original classifier before accuracy was degraded by randomization. No group would prefer \hat{Y}_F .

3 Active Fairness

The present work explores *active feature acquisition* approaches for achieving fairness, where a decision-maker adaptively acquires information according to the needs of different groups or individuals, in order to balance disparities in classification performance. This section defines two such strategies, one that allocates group-level information budgets—constant for all members of a group—and one that allocates individual-level information budgets, which are computed dynamically at test time. Sections 5 and 6 demonstrate their use and advantages in attaining fairness.

Preliminaries. We denote data of each decision subject as a pair (x, y) , where x is a feature vector of dimensionality d , and y is an outcome of interest. Let $S = (x^i, y^i)_{i=1}^n$ denote a labeled dataset, and $A \subset S$ represent a population subgroup. Let $\hat{Y}(X)$ be a binary classifier. We denote by $FPR_A(\hat{Y})$ and $FNR_A(\hat{Y})$ the false positive and false negative rates of $\{(x, y) \in A\}$, and define disparity measures with respect to A in terms of the following FPR and FNR differences:

$$D_{FPR}^A = |FPR_A(\hat{Y}) - FPR_{A^c}(\hat{Y})|$$

$$D_{FNR}^A = |FNR_A(\hat{Y}) - FNR_{A^c}(\hat{Y})|$$

Equal opportunity—or FNR parity—with respect to A requires that $D_{FNR}^A = 0$, while *equal odds* requires that both $D_{FNR}^A = D_{FPR}^A = 0$ (Hardt et al. 2016).

3.1 Group-Level Information Budgets

Let b_A, b_B be the information budgets for population subgroups A, B . We define predictor h^g with group-level information budgets b_A, b_B by:

$$h^g(x_i) = \begin{cases} f(x_i^{(q_A)}) & \text{if } i \in A \\ f(x_i^{(q_B)}) & \text{if } i \in B \end{cases}$$

where q_A, q_B are feature sets that satisfy $|q_A| = b_A$ and $|q_B| = b_B$. Sections 5 and 6 show how decision-makers can achieve calibration and group-level equity by allocating budgets b_A, b_B .

3.2 Individual-Level Information Budgets

Beyond group-level budgets, an ADM may adaptively collect information of each decision subject until a confidence threshold is met, upon which a classification decision is made. Thereby, individual-level information budgets are set dynamically according to the needs of each decision subject, towards attaining equity.

In particular, Algorithm 1 specifies active inquiry at the individual level as the decision-making process that, given lower and upper probability thresholds $\alpha_l, \alpha_u \in (0, 1)$, and for each decision subject i , progressively expands the information set $x_i^{(q_i)}$ until either threshold is met, or the available feature set is exhausted. Together with the decision threshold, α_l and α_u control trade-offs between FPR and FNR. In line with related AFA methods (Gao and Koller), we apply early stopping to ensure we stop expanding the feature set if the classification confidence is no longer improving significantly. We estimate the parameter for early stopping using grid search while maximizing the AUC for a given budget.

We define predictor h^{ind} with individual-level information budgets as $h^{ind}(x_i) = f(x_i^{(q_i)})$, where q_i is the feature set according to active inquiry in Algorithm 1.

Algorithm 1: Active inquiry at the individual level

Input: data X , model f , probabilities (α_l, α_u) , decision threshold t ;
for $i = 1$ **to** $i = n$ **do**
 while $f(x_i^{(q_i)}) \leq \alpha_u$ **and** $f(x_i^{(q_i)}) \geq \alpha_l$ **and**
 $|q_i| < d$ **and not** e **do**
 $j' \leftarrow$ Get next best feature $j' \notin q_i$;
 $x_i^{(q_i)} \leftarrow x_i^{(q_i)} \cup x_{ij'}$;
 $e \leftarrow$ early_stopping($x_i^{(q_i)}, x_{ij'}$);
 end
 $\hat{y}_i = \mathbb{1}_{[t,1]}(f(x_i^{(q_i)}))$;
end
return $(q_i)_{i=1}^n, \hat{Y}$;

3.3 Random Forest Implementation

Implementation of active classifiers requires two elements: (1) a model f , able to estimate $P\{Y|X^{(q)}\}$ for arbitrary feature subsets $X^{(q)}$, with $q \in [0, d]$, and (2) a feature selection method for choosing expanding feature sets, either at the group- or individual-level.

Probabilistic model. We implement distribution-based classification with incomplete data based on a probabilistic random forest and extending related methods for dealing with incomplete data in trees (Quinlan 2014; Saar-Tsechansky and Provost 2007). In particular, when given an arbitrarily incomplete feature vector $x_i^{(q)}$, the algorithm traverses all possible paths of each tree according to the following rule: if value x_{ij} for the current decision node is available in q , the search follows the path according to the node’s decision function; otherwise, if the value is not available ($j \notin q$), the search follows both paths. We then compute classification probabilities as a weighted average of the leaf purity across all leaves landed on by the search. Finally, we compute the average predicted probability across all trees. Similar methods can be derived for adapting logistic regressions to admit arbitrarily incomplete feature vectors (Williams et al. 2005; Saar-Tsechansky and Provost 2007).

Static feature selection. We first consider a static feature ranking for guiding the acquisition of additional features in Algorithm 1, based on feature importance derived from the random forest inter-trees variability. Hence, under static feature selection, given feature ranking R , the group-level budget classifier uses the top- b_A variables in R for classification of any $i \in A$, and the top- b_B variables in R for any $i \in B$. Similarly, the individual-level budgets classifier collects the top- b_i features in R in order to classify each subject i .

Dynamic feature selection. In the same vein, we consider dynamic or personalized feature selection, given its potential for increased individual-level equity and overall cost-efficiency. For it we implemented a *greedy* feature selection algorithm, which, for each subject i , and at each feature collection iteration, searches for the feature $j' \notin q_i$ that maximizes the difference between the current predicted probability \hat{P} and the expected probability given that an additional feature j' is queried, given by:

$$j' = \arg \max_{\{j: j \notin q_i, j \in [0, d]\}} |\hat{P}\{y_i = 1 | x_i^{(q_i \cup j')}\} - \hat{P}\{y_i = 1 | x_i^{(q_i)}\}|$$

4 Datasets

We study these methods and compare them to randomization-based approaches on four real-world, public datasets. All results are computed using random 80%/20% train/test splits.

Mexican poverty. Targeted social programs are challenged with household poverty prediction in order to determine eligibility (Ibarrarán et al. 2017). This dataset is extracted from the Mexican household survey 2016, which contains ground-truth household poverty levels, as well as a series of visible household features on which inferences are based. The dataset comprises a sample of 70,305 households in Mexico, with 183 categorical and continuous features, related to households’ observable attributes and other socio-demographic features. Classification is binary according to the country’s official poverty line, with 36% of the households having the label poor. We study fairness across groups defined by a) young and old families, split by the mean (where 53% are young), and b) across families living in urban and rural areas (where 64% are urban).

Adult income. The Adult Dataset from UCI Machine Learning Repository (Lichman and others 2013) comprises 14 demographic and occupational attributes for 49,000 individuals, with the goal of classifying whether a person’s income is above \$50,000 (25% are above), and using ethnicity (whites v. non-whites) as sensitive attribute (where 86% are white).

German credit. The German Credit dataset from UCI Machine Learning Repository consists of 1000 instances, of which 70% correspond to credit-worthy applicants and 30% correspond to applicants to whom credit should not be extended. Each applicant is described by 24 attributes. The sensitive attribute describes whether people are below or above the mean age (60% is below).

Heart health prediction. The Heart Dataset from the UCI Machine Learning Repository contains 17 features from 906

adults. The target is to accurately predict whether or not an individual has a heart condition (54% has a heart condition). The sensitive attribute is whether people are below or above the mean age (46% is below).

5 Achieving Equal Opportunity & Calibration

This section demonstrates how an active strategy with group-level budgets can be used to achieve calibration and single error parity, resulting in a higher efficiency and without resorting to naive randomization.

We follow (Pleiss et al. 2017) and study predictive performance in terms of the generalized false positive (GFPR) and false negative rates (GFNR), appropriate for contexts where risk scores themselves are the outputs of the algorithm (as opposed to fully automated classification). We aim at designing classifiers that satisfy calibration and error parity. As shown by (Kleinberg, Mullainathan, and Raghavan 2016), the GFNR and GFPR of all calibrated classifiers for a given group A fall along the straight line with slope $(1 - \mu_A)/\mu_A$, where $\mu_A = P(Y = 1|A)$ is A ’s base-rate, and origin in the perfect classifier with $(GFPR, GFNR) = (0, 0)$.

Panels A and D in Figure 1 show the space of calibrated classifiers achievable by naive randomization (method in (Pleiss et al. 2017)), for the *Mexican poverty* and *adult income* datasets described in Section 4. These replicate results from (Pleiss et al. 2017), showing how naive randomization of individuals in the advantaged group can, by eroding prediction performance, achieve calibration as well as either parity in false positives, parity in false negatives (but not both), or an *equal cost* generalization.

Similarly, panels B and E in Figure 1 demonstrate how calibration and either of the three parity objectives can be achieved by adjusting information budgets according to the groups’ needs, without resorting to naive randomization. Moreover, the right column in Figure 1 shows that classifiers with group-level budgets achieve these fairness goals with much higher efficiency in terms of information cost. In particular, we set an overall information budget restriction for both types of classifiers, equal to the minimum budget required by the naive random classifier to achieve *equal opportunity*. It is observed in panels C and F that the classifiers with group-level budgets Pareto-dominate random classifiers by a wide margin, on both datasets, i.e.: for the same information budget, both population subgroups are better off, being exposed to substantially lower false positive and false negative errors.

6 Achieving Equal Odds

This section shows how active methods with group- and individual-level information budgets can be used to achieve parity in false positives and false negatives.

Figure 2 illustrates the achievable regions in FPR-FNR space for classifiers with group-level information budgets, for two subgroups in the *Mexican poverty dataset*. It is observed that urban households are more predictable than rural households (achievable regions closer to the origin). The yellow and purple areas comprise the achievable regions for




Figure 2: Achieving parity in false positives and false negatives (*equal odds*) via group-level information budgets. Results correspond to the *Mexican poverty dataset*. Achievable regions of classifiers for each population subgroup are plotted in blue and yellow. The outer and inner FPR-FNR curves of each achievable region correspond to classifiers using maximum and minimum information budgets. Points along the curve correspond to different values of the decision threshold. It is observed that active classifiers with group-level budgets achieve parity in both FNR and FPR (*equal odds*). Moreover, they provide *equal odds* solutions anywhere on the overlap of the achievable regions for both subgroups and thus along the entire FNR-FPR trade-off.

urban and rural groups. A substantial overlap is observed, showing a wide-ranged achievable region for *equal odds*.

In a similar way, we can obtain the achievable region of active classifiers with individual-level information budgets, by varying parameters $\alpha_l < \alpha_u \in [0, 1]$ of Algorithm 1 (see Section 3).

We ran experiments to compare the three methods—naive randomization (as in (Hardt et al. 2016)), group-level budgets, and individual-level budgets—and their performance in achieving *equal odds* solutions along the FNR-FPR trade-off. In particular, we introduce an information budget constraint $b = \frac{1}{n} \sum_{\mathcal{V}_i} b_i < b_{max}$, and compare solutions sets that satisfy it. Solutions of individual- and group-level classifiers are discrete, due to finite sample sizes and features dimensionality.

Figure 3 shows results for three real-world datasets: *Mexican poverty*, *German credit*, and *Heart health* datasets. We left out the *Adult Income* dataset used in Fig. 1 since there exists no overlap between achievable regions for both subgroups and therefore we cannot achieve *equal odds*. Points in the FNR-FPR space were filtered to include only classifier designs that satisfied *equal odds* and an overall information budget constraint b_{max} .

It is observed that both group-level and individual-level strategies yield *equal odds* solutions, covering a wide range along the FNR-FPR trade-off curve, and without resorting to naive randomization. Moreover, it is shown that both type of active classifiers are substantially more information-efficient than the randomized classifier—Pareto dominance

along most of the FNR-FPR trade-off curve—leading to lower false positive and false negative errors in budget-constrained environments. Finally, active classifiers with individual-level budgets tend to dominate classifiers with group-level budgets, due to their more efficient use of information by means of personalized inquiry.

7 Conclusions

We have proposed and demonstrated methods for simultaneously achieving equal opportunity and calibration, as well as for achieving equal odds. In contrast to prior work, the *active* framework does not rely on naive randomization to reach these fairness notions, avoiding several known disadvantages of randomized approaches. Instead, a decision-maker acquires partial information sets according to the needs of different groups or individuals, allocating resources equitably in order to achieve balance in predictive performance. By leveraging this additional degree of freedom, active approaches can outperform randomization-based classifiers previously considered optimal. Moreover, classifiers with individual-level budgets dominated their group-level counterparts. Finally, the extent to which the former can as well reduce intra-group unfairness is a relevant question left to future work.

More broadly, this work illustrates how, by jointly considering information collection, inference, and decision-making processes, we can design automated decision systems that more flexibly optimize social objectives, including fairness, accuracy, efficiency, and privacy. A natural direction for future work is to consider richer utility functions relevant to real-world decision systems. We expect future studies that generalize results here presented to contexts with varying feature costs; as well as to contexts with multi-stakeholder value functions, where the opportunity, privacy, and monetary costs that inquiry and decision-making bring to decision-subjects are jointly considered as part of the adaptive inquiry process.

Lastly, a relevant path forward is to allow observations with partial feature sets both during training and test phases. The current implementation of this work necessitates access to full-feature observations at training time. More efficient training and further model refinement could be achieved under schemes that can learn from partial feature vectors, or proactively collect features at training time; allowing to incorporate a wider set of features tailored to increase prediction accuracy over different types of individuals.

References

- Bechavod, Y., and Ligett, K. 2017. Learning fair classifiers: A regularization-inspired approach. *arXiv preprint*.
- Calders, T.; Kamiran, F.; and Pechenizkiy, M. Building classifiers with independency constraints. In *Data mining workshops, 2009. ICDMW'09. IEEE international conference on*.
- Chalfin, A.; Danieli, O.; Hillis, A.; Jelveh, Z.; Luca, M.; Ludwig, J.; and Mullainathan, S. Productivity and selection of human capital with machine learning. *American Economic Review* 106(5).
- Chouldechova, A.; Benavides-Prado, D.; Fialko, O.; and Vaithianathan, R. 2018. A case study of algorithm-assisted decision




Figure 3: Achievable space of equal odds classifiers. Comparison of active classifiers with group-level (green) and individual-level (yellow) information budgets, and naive randomization (blue, following (Hardt et al. 2016)). Classifiers are constrained by an information budget $\bar{b} \leq b_{max}$. Figure A corresponds to *Mexican poverty*, B to *German credit* (center) and C to *Heart health* under different budget constraints. It is observed that active classifiers yield *equal odds* solutions along the FPR-FNR trade-off, without resorting to randomization. Moreover, in a budget-constraint setting, the active classifiers are substantially more information-efficient. The individual-level budget Pareto-dominates the group-level classifiers which in turn dominate randomized classifiers in budget-constrained environments.

making in child maltreatment hotline screening decisions. In *Conference on Fairness, Accountability and Transparency*, 134–148.

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*.

Corbett-Davies, S.; Pierson, E.; Feller, A.; Goel, S.; and Huq, A. 2017. Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 797–806. ACM.

Flores, A. W.; Bechtel, K.; and Lowenkamp, C. T. 2016. False positives, false negatives, and false analyses: A rejoinder to machine bias: There’s software used across the country to predict future criminals. and it’s biased against blacks. *Fed. Probation* 80:38.

Gao, T., and Koller, D. Active classification based on value of classifier. In *Advances in Neural Information Processing Systems*.

Gillingham, P. 2015. Predictive risk modelling to prevent child maltreatment and other adverse outcomes for service users: Inside the black box of machine learning. *The British Journal of Social Work* 46(4):1044–1058.

Gorry, G., and Barnett, O. 1968. Sequential diagnosis by computer.

Hardt, M.; Price, E.; Srebro, N.; et al. 2016. Equality of opportunity in supervised learning. In *Advances in neural information processing systems*, 3315–3323.

Huang, C.-L.; Chen, M.-C.; and Wang, C.-J. 2007. Credit scoring with a data mining approach based on support vector machines. *Expert systems with applications* 33(4):847–856.

Ibarrarán, P.; Medellín, N.; Regalía, F.; Stampini, M.; Parodi, S.; Tejerina, L.; Cueva, P.; and Vásquez, M. 2017. How conditional cash transfers work.

Kanani, P., and Melville, P. 2008. Prediction-time active feature-value acquisition for cost-effective customer targeting. *Advances In Neural Information Processing Systems (NIPS)*.

Kleinberg, J.; Ludwig, J.; Mullainathan, S.; and Obermeyer, Z. Prediction policy problems. *American Economic Review* 105(5).

Kleinberg, J.; Mullainathan, S.; and Raghavan, M. 2016. Inherent trade-offs in the fair determination of risk scores. *arXiv preprint*.

Lichman, M., et al. 2013. Uci machine learning repository.

Liu, L.-P.; Yu, Y.; Jiang, Y.; and Zhou, Z.-H. Tefe: A time-efficient approach to feature extraction. In *Data Mining, 2008. ICDM’08. Eighth IEEE International Conference on*.

Louizos, C.; Swersky, K.; Li, Y.; Welling, M.; and Zemel, R. 2015. The variational fair autoencoder. *arXiv preprint arXiv:1511.00830*.

MacKay, D. J. 1992. Information-based objective functions for active data selection. *Neural computation* 4(4):590–604.

Pleiss, G.; Raghavan, M.; Wu, F.; Kleinberg, J.; and Weinberger, K. Q. 2017. On fairness and calibration. In *Advances in Neural Information Processing Systems*, 5680–5689.

Quinlan, R. 2014. *C4.5: programs for machine learning*. Elsevier.

Saar-Tsechansky, M., and Provost, F. 2007. Handling missing values when applying classification models. *Journal of machine learning research* 8(Jul):1623–1657.

Settles, B. 2012. Active learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 6(1):1–114.

Siegel, E. 2013. *Predictive analytics: The power to predict who will click, buy, lie, or die*. Wiley Hoboken.

Trapeznikov, K., and Saligrama, V. 2013. Supervised sequential classification under budget constraints. In *Artificial Intelligence and Statistics*, 581–589.

Williams, D.; Liao, X.; Xue, Y.; and Carin, L. 2005. Incomplete-data classification using logistic regression. In *Proceedings of the 22nd International Conference on Machine Learning*, 972–979.

Zafar, M. B.; Valera, I.; Rodriguez, M. G.; and Gummadi, K. P. 2015. Learning fair classifiers. *arXiv preprint arXiv:1507.05259*.

Žliobaitė, I. Measuring discrimination in algorithmic decision making. *Data Mining and Knowledge Discovery* 31(4).