

Unleashing Linear Optimizers for Group-Fair Learning and Optimization

Daniel Alabi

Harvard University, Cambridge, MA

ALABID@G.HARVARD.EDU

Nicole Immorlica

Adam Tauman Kalai

Microsoft Research New England, Cambridge, MA

NICIMM@MICROSOFT.COM

ADAM.KALAI@MICROSOFT.COM

Editors: Sebastien Bubeck, Vianney Perchet and Philippe Rigollet

Abstract

Most systems and learning algorithms optimize average performance or average loss – one reason being computational complexity. However, many objectives of practical interest are more complex than simply average loss. This arises, for example, when balancing performance or loss with fairness across people. We prove that, from a computational perspective, optimizing arbitrary objectives that take into account performance over a small number of groups is not significantly harder to optimize than average performance. Our main result is a polynomial-time reduction that uses a linear optimizer to optimize an arbitrary (Lipschitz continuous) function of performance over a (constant) number of possibly-overlapping groups. This includes fairness objectives over small numbers of groups, and we further point out that other existing notions of fairness such as individual fairness can be cast as convex optimization and hence more standard convex techniques can be used. Beyond learning, our approach applies to multi-objective optimization, more generally.

Keywords: optimization, fairness, learning, online algorithms, projection

1. Introduction

The objectives of today’s algorithms may be complex and non-linear. For example, a 2016 Bloomberg report demonstrated systematic bias in the roll-out of a major online retailer’s delivery service ([Ingold and Soper, 2016](#)). It argued that, in many major metropolitan areas like Chicago, Boston, and New York, the service area excluded predominantly black zip codes. We give reductions that might help such decision-makers account for complex factors such as fairness, from choosing facility locations, to advertising, to product selection and pricing. We show how existing systems that optimize “average performance” can be used to optimize more complex objective functions. Our work builds on a line of linear reductions in algorithms and fair machine learning (e.g., [Jaggi, 2013](#); [Kalai and Vempala, 2005](#); [Narasimhan et al., 2015](#); [Agarwal et al., 2017](#); [Dwork et al., 2018](#)).

Consider the following optimization problem:

$$\min_{c \in \mathcal{C}} f(\ell_1(c), \dots, \ell_K(c)) \quad (1)$$

Here \mathcal{C} is a (possibly constrained) set of choices and *loss vector* $\ell(c) \in [0, 1]^K$ represents costs or losses to K possibly-overlapping groups, such as racial and gender groups. As a simple optimization example, consider the classic Facility Location Problem (FLOP), in which an algorithm is given

a graph G and customer weights, and must select m locations for facilities amongst the nodes on a directed graph, so as to minimize the weighted average distance of nodes to the closest facility. FLOPs model a firm choosing a limited number of locations to serve a set of customers, where c could encode the m chosen facilities and $\ell_k(c)$ could encode the average distance over customers in that group. The customer weights can also be used to reweight groups, hence a standard FLOP algorithm can minimize $w \cdot \ell(c)$ for any given group weight vector w . But natural objectives f may be complex and non-convex, even without factoring in fairness. For instance, in a FLOP the marginal advantage of reducing group k 's average delivery time $l_k = \ell_k(c)$ from 2 hours to 1 hour may be more significant than reducing it from 48 hours to 47 hours, as would be captured by a concave function such as $f(l) = \sum_k \sqrt{l_k}$. Convexity is often required for algorithmic convenience, but may not be natural in any given application.

As an ML example, consider a loan approval decision where the choices are approval criteria, i.e., classifiers $c(x)$. While it may be illegal for the approval decisions to depend on sex, the choice of c may take into account resulting demographic differences. For simplicity, say there are two sexes and four groups based on the individual's sex (M or F) and whether $y = 1$ indicating they would repay or $y = 0$ indicating they would default. Let the loss vector $\ell(c) \in [0, 1]^4$ consist of the misclassification rates on the four groups, which are the F/M false-positive/negative rates. Each classifier corresponds to a feasible point in $[0, 1]^4$. For the moment, consider only classifiers with no false positives (no approvals that default), so the feasible set $F = \{\ell(c) | c \in \mathcal{C}\}$ can be visualized in the two dimensions of F/M false negative rates. The feasible set is the set of achievable F- and M-false-negative rates achievable by classifiers, e.g., each neural network corresponds to a feasible point. Since one may randomize amongst classifiers, one can further assume that the feasible set is convex.

Unfortunately the feasible set F is not directly accessible, e.g., we do not know if there is a neural network that achieves specified F/M false negative rates (or specified average group distances in FLOP). [Narasimhan et al. \(2015\)](#) showed how to use binary classification algorithms to optimize in any linear direction and consequently closely approximate F . In particular, they showed how to use a Frank-Wolfe style reduction to minimize convex objectives as well as ratio-of-linear objectives. [Agarwal et al. \(2017, 2018\)](#) provide a related reduction for fair ML with arbitrarily many groups and convex constraints. In this work, we consider the case of small numbers of groups, K , where we show we can optimize arbitrary (Lipschitz) continuous objectives.

For example, consider the following stylized objective that combines misclassification rate with a penalty for M-F disparity amongst approvals:

$$\min_{\substack{c \in \mathcal{C}: \\ \Pr[c(x)=1 \wedge y=0]=0}} \Pr[c(x) \neq y] + \left(\Pr[x_{\text{sex}} = \text{F} | c(x) = 1] - \Pr[x_{\text{sex}} = \text{M} | c(x) = 1] \right)^2 \quad (2)$$

This objective (and the false positive restriction) can be computed in terms of the loss vector $\ell(c)$. However, as illustrated in [Figure 1](#) (false positive rates not shown because they are all 0), this objective is not convex and can have local minima that are not global minima.

Fairness is complicated by many factors such as biased training labels and societal biases that cannot be corrected by loan approval alone. In [Figure 1](#), exactly equal F/M rates are only achieved by rejecting all loans, which may not be very useful to a loan-granting organization. Legal definitions of discrimination discuss process and other factors beyond approval rates ([Barocas and Selbst, 2016](#)). Nonetheless, given that algorithms are making such decisions, the flexibility to use existing algorithms to optimize complex non-linear objectives enables domain experts to choose objectives that may lead to desirable outcomes, accounting for biased data, fairness, and of course performance.




Figure 1: A contour plot of non-convex objective f from eq. (2) as a function of M/F false negative rates (with 0 false positives, and uniform $1/4$ probability of $\pm M/\pm F$). The set of false positive/negative rates achieved by classifiers (yellow) has a local (non-global) minimum.

Our reduction helps because it can reuse existing linear systems and can circumvent local optima that might trap more direct methods.

Algorithm 1 gives a (polynomial-time) reduction that can minimize (Lipschitz) continuous functions f (for constant K), using an underlying algorithm that can minimize linear loss $\sum w_k \ell_k(c)$, for arbitrary real-valued weights w_k . As mentioned, customer weights standard in FLOP algorithms (meant to indicate for instance the frequency of serving these customers) can be used to weight groups, and hence repeated calls to the FLOP algorithm on the same graph with different weights are used to efficiently minimize f . For binary classification, a corollary of our main theorem is the following:

Corollary 5 (Informal version) *Let M be an agnostic learner for family \mathcal{C} of binary classifiers $c : \mathcal{X} \rightarrow \{0, 1\}$. Let $\mathcal{X}_1, \dots, \mathcal{X}_K \subseteq \mathcal{X}$ be sensitive groups. Then any Lipschitz continuous function of the error rates and false positive and negative rates across the groups with Lipschitz-continuous constraints can be minimized over \mathcal{C} to within ϵ in time and calls to M that are $\text{poly}(1/\epsilon)$ for any constant K .*

Agnostic learning (Kearns et al., 1992) is a natural model of computationally efficient binary classification. The prior work of Narasimhan et al. (2015) and Agarwal et al. (2017, 2018), which motivates our work, shows how to optimize *convex* objectives (and ratio-of-linear objectives) with convex constraints on group error rates using an existing agnostic learner. The above corollary extends this to arbitrary continuous objectives, though it does require constant numbers of groups K . We also show (Corollary 9) how to use an agnostic learner to optimize a continuous and nondecreasing function of squared errors across groups – hence a classifier can be used for a form of fair regression.

NP-hard problem such as FLOP and many agnostic learning problems are intractable for some large instances, but algorithms may run quickly on “real-world” instances. For instance, as long as the FLOP algorithm runs quickly on a graph G , eq. (1) can be efficiently solved on G . Similarly, while neural network training algorithms are not guaranteed to find optimal solutions, in practice they often perform well.

Constraints. Note that general (continuous) constraint functions $g(l) \leq 0$ can easily be folded into the objective — if $R = \max_{l, l' \in F} f(l) - f(l')$ is the range of f and opt is the minimum value of

$f(l)$ over feasible $l = \ell(c)$ such that $g(l) \leq 0$ (assuming the constraint is satisfiable), then

$$f(l) + \Gamma \max(0, g(l)) \leq \text{opt} + \epsilon \quad \text{implies} \quad f(l) \leq \text{opt} + \epsilon \quad \text{and} \quad g(l) \leq (\epsilon + R)/\Gamma. \quad (3)$$

In our example (2), constraining the false positive rate to $\leq \delta$ (it cannot be constrained to exactly 0 with finite samples) can be achieved by adding $\Gamma \Pr[c(x) = 1 \wedge y = 0]$ to the objective, for $\Gamma \geq (R + \epsilon)/\delta$. A set of linear constraints $u_i \cdot l \leq b_i$ can be captured by $g(l) = \max_i u_i \cdot l - b_i$.

Convex objectives. If f happens to be a convex function, then Section 6 points out that approaches like that of Narasimhan et al. (2015) can be used to convert a weighted-sum minimizer of $\sum w_i \ell_i(c)$ to efficiently optimize (1) to within ϵ in time $\text{poly}(1/\epsilon, K)$, which means that we can handle large numbers of groups or even $K = n$, i.e., a group for each customer. For instance, a FLOP algorithm could be used to find a (distribution over) m facilities that minimize the maximum (expected) customer’s distance to its nearest facility, or the Generalized Gini Index considered in fairness (Weymark, 1981; Busa-Fekete et al., 2017) and which has been studied in FLOP in particular (Drezner et al., 2009). It can also find classifier parameters for individual fairness (also called metric fairness) (Dwork et al., 2012).

Organization. We next present related work followed by mathematical preliminaries. Section 4 presents our general reduction for continuous functions of losses across K groups. Section 5 shows how to apply this to learning. And Section 6 shows how to reduce convex objectives to linear optimization.

2. Related Work

Reductions in binary classification. Narasimhan et al. (2015) show reductions to binary classification albeit for performance measures that can be expressed as a ratio of linear functions or that are concave. Dwork et al. (2018) show how to reduce the problem of optimizing a group-fair loss (for disjoint groups) to standard binary classification, however their work requires that classifiers c have access to the sensitive attribute determining group membership. Their approach is simply to learn a different classifier for each group, whereas in our work the classifiers may or may not (or may have partial) access to sensitive attributes. Their work further requires f to be non-decreasing but, on the other hand, does not require f to be continuous. Agarwal et al. (2017, 2018) show how to reduce convex fairness constraints, such as demographic parity and equalized odds, to binary classification. Within classification, our work builds upon and generalizes these works in that we can handle arbitrary continuous functions of losses (for a constant number of possibly overlapping groups).

Multi-objective optimization. When faced with multiple objectives, no one single solution may be optimal (see, e.g., Zuluaga et al., 2013; Ravi et al., 1993; Grandoni et al., 2014). A solution is Pareto-optimal if no other solution is simultaneously better in all objectives. Often the goal is to recover the entire set of Pareto-optimal solutions. While computing the Pareto set (even membership) is NP-hard (Ravi et al., 1993), for some problems it can be succinctly approximated (Warburton, 1987; Papadimitriou and Yannakakis, 2000; Diakonikolas, 2011). Some work in multi-objective optimization addresses the problem as a convex programming problem, while other work such as Diakonikolas (2011) provides reductions to an underlying linear optimization problem. Our work follows in this spirit.

Optimizing non-decomposable ML losses. Misclassification rate is the most common performance measure studied in learning theory. However in cases of severe label imbalance, other evaluation metrics over the entire dataset may be used, but these metrics are often non-decomposable. For example, [Parambath et al. \(2014\)](#) show how to reduce the problem of minimizing F-measures such as the F1 score, the harmonic mean of precision and recall, to linear optimization. Our work generalizes this work.

Fairness in classification. The study of fairness in classification is an important emerging research area ([Dwork et al. \(2012\)](#); [Joseph et al. \(2016\)](#); [Kearns et al. \(2017\)](#)). Some impossibility results are known. For instance, except in cases of equal base rates of true classifications across groups or in very rare cases of perfect predictions, there is no well-calibrated and class-balanced classifier ([Kleinberg et al. \(2017\)](#); [Chouldechova \(2017\)](#)). Despite such results, we know that we can design multi-objective metrics that trade-off fairness for accuracy and vice-versa ([Woodworth et al. \(2017\)](#); [Menon and Williamson \(2017\)](#); [Dwork et al. \(2018\)](#); [Goh et al. \(2016\)](#)). We show that optimizing fair multi-objectives (for a constant number of possibly overlapping groups) is no harder than optimizing approximate linear optimizers. [Ogryczak et al. \(2014\)](#) survey some models and methods for achieving fairness, with a particular emphasis on the optimization of networks.

Frank-Wolfe and online optimization. The general approach of reducing optimization to linear optimization has been utilized in many areas of machine learning ([Jaggi, 2013](#); [Kalai et al., 2008](#)). For the convex optimization, we use the reduction of [Kakade et al. \(2009\)](#) as a sub-routine. We also leverage the fact that their reduction can handle minimizers with limited domains.

3. Mathematical Notation

Let $\Delta(S)$ denote the set of probability distributions over S and let T^S denote the set of functions $f : S \rightarrow T$, for sets S, T . For $U \subseteq S$, write $f(U) = \{f(u) \mid u \in U\}$. For $v, w \in \mathbb{R}^K$, we write $v \leq w$ if $v_i \leq w_i$ for all i , and for scalar $r \in \mathbb{R}$, the vector-scalar sum is $v + r = (v_1 + r, \dots, v_K + r) \in \mathbb{R}^K$. A function $f : \mathbb{R}^K \rightarrow \mathbb{R}$ is said to be nondecreasing if it is nondecreasing in each coordinate, or equivalently $f(v) \leq f(w)$ for all $v \leq w$. We say that function $f : \mathbb{R}^n \rightarrow \mathbb{R}^d$ is convex if it is convex in each coordinate, i.e., if $f(\lambda v + (1 - \lambda)w) \leq \lambda f(v) + (1 - \lambda)f(w)$ for $\lambda \in [0, 1]$.

For $v \in \mathbb{R}^d$, let $(v)_+ = (\max(0, v_1), \dots, \max(0, v_d)) \in \mathbb{R}^d$ and $\|v\|_\infty = \max_i v_i$. Let $\mathbb{R}_+ = [0, \infty)$. For set S , let $S^* = S \cup S \times S \cup \dots$ be arbitrary-length sequences of elements of S .

Also, for $L \geq 0$, $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is said to be L -Lipschitz continuous (or just L -Lipschitz) if $|f(x) - f(y)| \leq L\|x - y\|$ for all $x, y \in \mathbb{R}^d$. We say f is Lipschitz if it is 1-Lipschitz. Finally, let $[n] = \{1, 2, \dots, n\}$ and $\mathbb{1}[P]$ denote 1 if predicate P holds, otherwise 0.

4. Group Objectives

Let K be a constant number of groups, let \mathcal{C} be a set of choices, and let $\ell : \mathcal{C} \rightarrow [0, 1]^K$ be a loss function that determines the (bounded) average loss on each choice in c . We assume that \mathcal{C} is *closed under randomization* meaning that every $c_1, \dots, c_n \in \mathcal{C}$ can be identified with a choice $c \in \mathcal{C}$, written $c = \text{UniformDist}(\{c_1, \dots, c_n\})$, with $\ell(c) = (\ell(c_1) + \dots + \ell(c_n))/n$. This assumption is necessary for fairness, e.g., clearly randomization is necessary for fairness in locating a single facility on a two-node graph. Even for nonlinear c 's such as neural networks, it is possible to randomize between neural networks or average their outputs (which is very different than averaging the weights).

We also fix a comparison set $\bar{\mathcal{C}} \subseteq \mathcal{C}$, which may be equal $\bar{\mathcal{C}} = \mathcal{C}$ but more generally allows an algorithm to output a choice in a different set \mathcal{C} but only be compared to choices from $\bar{\mathcal{C}}$. This technicality allows us to capture so-called “improper” learning algorithms that learn a class $\bar{\mathcal{C}}$ by outputting classifiers in a different family, e.g., $\bar{\mathcal{C}}$ may be a class of bounded-size binary decision trees and \mathcal{C} may include the larger class of low-degree polynomials. We also assume that $\bar{\mathcal{C}}$ is closed under randomization, so our guarantees compared to the best $\bar{c} \in \bar{\mathcal{C}}$ also compare to probability distributions over choices.

For $w \in \mathbb{R}^K$ and $c \in \mathcal{C}$, we call $w \cdot \ell(c)$ the weighted loss of c . We require that we can efficiently (approximately) minimize $w \cdot \ell(c)$. In some applications, this can only be done for non-negative $w \geq 0$, and for those we will have an additional monotonicity requirement on f . To this end, consider two types of linear optimizers: (a) a general linear optimizer $M_\tau : \mathbb{R}^K \rightarrow \mathcal{C}$ which can take arbitrary real weights; and (b) a nonnegative linear optimizer $M_\tau : \mathbb{R}_+^K \rightarrow \mathcal{C}$, with its conditions in parentheses, which takes only nonnegative weights.

Definition 1 ((Nonnegative) Linear optimizer) *A (nonnegative) linear optimizer M_τ , for any tolerance $\tau > 0$, for any $w \in \mathbb{R}^K$ (with $w \geq 0$), outputs $M_\tau(w) \in \mathcal{C}$ which, to within τ , optimizes weighted loss, i.e.,*

$$w \cdot \ell(M_\tau(w)) \leq \min_{\bar{c} \in \bar{\mathcal{C}}} w \cdot \ell(\bar{c}) + \tau \|w\|$$

The tolerance τ is scaled by $\|w\|$ due to the fact that a scaling of w scales all $w \cdot \ell(c)$ by the same factor. Note that linear optimization does not need randomization: $\text{UniformDist}(\{c_1, \dots, c_n\})$ can always be replaced by the choice c_i that minimizes $w \cdot \ell(c_i)$. For instance, if the family \mathcal{C} consists of probability distributions over neural networks, then one could simply choose the most accurate neural network without degrading performance. For FLOP, the linear minimizer will output a fixed set of locations. Our algorithm, however, will randomize over small number of neural networks or a small number of t -tuples of locations.

Of course, we also must be able to compute the loss for each group, at least approximately:

Definition 2 (Loss assessor) *A loss assessor ℓ_τ , for any tolerance $\tau > 0$, computes a function $\ell_\tau : \mathcal{C} \rightarrow [0, 1]^K$ which approximates ℓ to within τ , i.e., $\|\ell_\tau(c) - \ell(c)\| \leq \tau$ for all $c \in \mathcal{C}$.*

Algorithm 1 uses a (nonnegative) linear optimizer and a loss assessor to minimize a (nondecreasing) Lipschitz continuous loss $f : [0, 1]^K \rightarrow \mathbb{R}$. It works by creating a grid G over $[0, 1]^K$, fine enough that the optimum is within $O(\epsilon)$ of some grid point. (Note that the optimum’s nearest grid point is nearly feasible.) We sort the nearly feasible grid points by $f(\cdot)$. For each such grid point r , in order from least to greatest $f(r)$, we see if we can find a $c \in \mathcal{C}$ with $\ell(c)$ within $O(\epsilon)$ of r by calling the minimizer $M(\ell(c) - r)$ and moving in that direction T times. Otherwise we move on to the next grid point. The algorithm is shown to stop on or before the grid point closest to the optimal feasible solution, and outputs a uniform distribution c over a set of $T = O(1/\epsilon^2 \log 1/\epsilon^2)$ choice settings.

Several notes about the algorithm are in order.

- While the algorithm outputs a convex combination of T choices, this can always be reduced to a combination of at most $K + 1$ choices (since we are in $[0, 1]^K$) using standard Carathodory’s Theorem (see, e.g., [Mulzer and Stein, 2014](#)).
- Equation (3) shows how to fold constraints $g(x) \leq 0$ for (nondecreasing) Lipschitz $g : \mathbb{R}^K \rightarrow \mathbb{R}$ into the objective.

- While the algorithm requires Lipschitz (i.e., $L = 1$) functions, it can be used for any L -Lipschitz f for $L > 1$ by simply applying the algorithm to Lipschitz function f/L with accuracy ϵ/L . The Lipschitz constant may depend on group demographics and on Γ (3). Further, Laplace smoothing (see, e.g. Schütze et al., 2008) can be used to make Lipschitz-continuous approximations to discontinuous objectives. For instance, the objective of eq. 2 is discontinuous at (1,1), but is continuous if one adds a constant to the counts of false positives/negatives in each group.

The theorem below provides guarantees for GROUPOPT (with guarantees for nonnegative linear optimizer $M_\tau : \mathbb{R}_+^K \rightarrow \mathcal{C}$ written in parentheses).

Algorithm 1 GROUPOPT: Minimizing group-loss f using linear optimizer

Input: accuracy $\epsilon > 0$, $f : [0, 1]^K \rightarrow \mathbb{R}$, loss assessor ℓ_τ , (nonnegative) linear optimizer M_τ

Output: $c \in \mathcal{C}$

Let $\beta = \frac{\epsilon}{5}$, $q = \frac{\beta}{\sqrt{K}}$, $\tau = \frac{\beta^2}{\sqrt{K}}$, $T = \frac{K}{\beta^2} \ln \frac{K}{\beta^2}$.

Create grid $G = \{0, q, 2q, 3q, \dots, \lfloor 1/q \rfloor q\}^K \subseteq [0, 1]^K$.

Check if f is nondecreasing coordinatewise on G . If so, let $N = 1$ else $N = 0$.

Sort points in G by $f(r)$ in increasing order

for r in G **do**

$c_1 = M_\tau(0)$ // any initial choice

for $t = 1$ to T **do**

$\hat{l}_t = \max \left(\frac{1}{t} (\ell_\tau(c_1) + \dots + \ell_\tau(c_t)), Nr \right)$

$c_{t+1} = M_\tau(\hat{l}_t - r)$

end

if $\|\hat{l}_T - r\| \leq 3\beta$ **then**

return $c = \text{UniformDist}(\{c_1, c_2, \dots, c_T\})$ // uniform probability distribution

end

end

Theorem 3 Let $K \geq 1$, $f : [0, 1]^K \rightarrow \mathbb{R}$ be (nondecreasing) Lipschitz-continuous functions, $\ell : \mathcal{C} \rightarrow [0, 1]^K$ be a group loss function with loss assessor ℓ_τ and (nonnegative) linear optimizer M_τ . For any $0 < \epsilon \leq 1$, $f(\ell(\text{GROUPOPT}(\epsilon, f, \ell_\tau, M_\tau))) \leq \min_{l \in \ell(\bar{\mathcal{C}})} f(l) + \epsilon$. For constant K , GROUPOPT makes $\text{poly}(1/\epsilon)$ queries to ℓ_τ and M_τ both with $\tau = \frac{\epsilon^2}{25\sqrt{K}}$, and uses $\text{poly}(1/\epsilon)$ additional runtime and calls to f .

Simply put, the computational complexity of optimizing continuous f is the same, up to polynomial factors, as linear optimization, for a constant number of groups K . As we shall see, since ℓ corresponds to the loss on a specific instance, this theorem implies that for those graphs or distributions on unlabeled data \mathcal{X} for which (weighted) linear optimization is possible, group optimization is possible as well. The proof of Theorem 3, which contains runtime bounds, is deferred to Appendix 8.1.

5. Group Learning Applications

In this section, we show how Theorem 3 can be applied to binary classification. In Section 5.2, we show that one can use a binary classifier to minimize arbitrary continuous group-loss objectives for classification. In Section 8.3 of the Appendix, we show that regression can be reduced to classification.

We begin by defining a standard notion of computationally efficient learning called *Agnostic learning* (Kearns et al., 1992). It is called agnostic because, unlike PAC learning (Valiant, 1984) which it generalizes, it makes no assumption on the distribution of data.

5.1. Agnostic Learning Definitions

Let \mathcal{X} be a set of feature vectors, $\mathcal{Z} \subseteq \mathcal{X} \times \{0, 1\}$ be the set of examples, and $\mu \in \Delta(\mathcal{Z})$ be a distribution over \mathcal{Z} . A *classifier* is a possibly-randomized function $c : \mathcal{X} \rightarrow \{0, 1\}$. Let $\mathcal{C} \subseteq \{0, 1\}^{\mathcal{X}}$ be a set of *legal* classifiers or predictors which can, for instance, enable restrictions on using protected attributes. Let $\bar{\mathcal{C}} \subseteq \mathcal{C}$ be a family of classifiers to be learned. Our framework is general and not restricted to the so-called *proper* case when $\bar{\mathcal{C}} = \mathcal{C}$. The error of a classifier c is

$$\text{err}_{\mu}(c) = \Pr_{x, y \sim \mu} [y \neq c(x)].$$

To fit in our framework, we assume the classes \mathcal{C} and $\bar{\mathcal{C}}$ are closed under randomization. For instance, if \mathcal{C} included rectangles in \mathbb{R}^d , it would also include probability distributions over rectangles. However, optimization with respect to this larger class is no harder because the most accurate single classifier (e.g., rectangle) is always as accurate as the most accurate distribution over rectangles. Hence, our requirement of agnostic learning is weak – one only needs to find a pure predictor, while our guarantees will be with respect to all distributions in $\bar{\mathcal{C}}$, e.g., mixtures over rectangles, which is necessary for fairness.

Definition 4 (Agnostic learning) *Polynomial-time algorithm $M : \mathcal{Z}^* \rightarrow \mathcal{C}$ is an agnostic learner for $\bar{\mathcal{C}} \subseteq \mathcal{C} \subseteq \{0, 1\}^{\mathcal{X}}$ ($[0, 1]^{\mathcal{X}}$) if there exists a polynomial $P(1/\epsilon, 1/\delta)$, such that for any $\epsilon, \delta > 0$, $\mu \in \Delta(\mathcal{Z})$, with probability $\geq 1 - \delta$ over $n \geq P(1/\epsilon, 1/\delta)$ examples $Z = \langle (x_1, y_1), \dots, (x_n, y_n) \rangle$ iid from μ , M 's output has error within ϵ of optimal, i.e.,*

$$\Pr_{Z \sim \mu^n} \left[\text{err}_{\mu}(M(Z)) \leq \min_{\bar{c} \in \bar{\mathcal{C}}} \text{err}_{\mu}(\bar{c}) + \epsilon \right] \geq 1 - \delta.$$

Note that an agnostic learning algorithm is often defined as taking ϵ, δ as input choices as well (Kearns et al., 1992), however, one can simply think of the algorithm as taking an arbitrary number of examples n as input and providing such a guarantee for a fixed schedule of ϵ, δ that are inverse polynomial functions of n (e.g. $\epsilon = \delta = n^{-1/4}$).

Distribution-specific agnostic learning refers to the case in which the above guarantee holds with respect to some specific distribution marginal on x . Common marginal distributions for which algorithms are analyzed include the uniform distribution on the hypercube $x \in \{0, 1\}^d$ or any Gaussian or log-concave distribution.

5.2. Group Agnostic Learning

We now consider agnostic learning for K groups $\mathcal{X}_1, \dots, \mathcal{X}_K \subseteq \mathcal{X}$, each with a group oracle, also denoted \mathcal{X}_k , that efficiently computes group membership. Note that classifiers do not directly have access to group membership – it is assumed to be available on training data but not for prediction on future examples. Typically, group membership is as simple as checking whether a protected attribute and/or label belong to a particular set. In our model, protected attributes are encoded in \mathcal{X} , but the family \mathcal{C} of legal classifiers may not be allowed to depend on them.

Let us illustrate on classification. Consider the following example for classification.

Example 1 $\mathcal{X} = \{-1, 1\}^2 \times \mathbb{R}^d$, where the first attribute x_1 encodes race, a binary protected attribute, meaning that it cannot be used for classification, while x_2 encodes gender, say also binary for illustration, which can be used for classification. The legal classifiers \mathcal{C} are all $c : \mathcal{X} \rightarrow [0, 1]$ such that $c(-1, x) = c(1, x)$ for any $x \in \{-1, 1\} \times \mathbb{R}^d$. Gender and race are needed for training data, while only gender is needed for classification. The family $\bar{\mathcal{C}}$ being learned could be linear thresholds $c_\theta(x) = \mathbb{1}[\theta \cdot x > 0]$ (with bounded norm $\|\theta\|$), with the additional requirement on $\theta \in \mathbb{R}^{2+d}$ that $\theta_1 = 0$.

To accommodate intersectionality (Crenshaw, 1989), one can create groups for each combination of group attributes. For instance, in the race/gender example, one could create 8 groups, one for each combination of race, gender, and label. More generally, given a constant number of subsets of \mathcal{X} , there are a constant number of intersections of groups and their complements together with binary labels.

$$\text{FPR}_k(c) = \Pr[c(x) = 1 \mid x \in \mathcal{X}_k \wedge y = 0]$$

$$\text{FNR}_k(c) = \Pr[c(x) = 0 \mid x \in \mathcal{X}_k \wedge y = 1]$$

For each $k \in [K], i \in \{0, 1\}$, define $p_k^i = \Pr[x \in \mathcal{X}_k \wedge y = i]$ and $p_k = p_k^0 + p_k^1 = \mu(\mathcal{X}_k \times \{0, 1\})$. For simplicity we assume that the algorithm is given exact knowledge of all p_k^i 's, but this assumption can be removed using standard estimation by Chernoff bounds. Note that with knowledge of p_k^i , the vector $(\text{FPR}(c), \text{FNR}(c)) \in [0, 1]^{2K}$ has sufficient statistics to compute all other quantities of interest such as the group error rates, precision, recall, and F1-score:

$$\text{ERR}_k(c) = \Pr[c(x) \neq y \mid x \in \mathcal{X}_k] = \frac{p_k^0 \text{FPR}_k(c) + p_k^1 \text{FNR}_k(c)}{p_k} \quad (4)$$

$$\text{PRE}_k(c) = 1 - \text{FNR}_k(c) \quad (5)$$

$$\text{REC}_k(c) = \frac{p_k^1 (1 - \text{FNR}_k(c))}{p_k^1 (1 - \text{FNR}_k(c)) + p_k^0 \text{FPR}_k(c)} \quad (6)$$

$$\text{F1}_k(c) = \frac{2 \text{PRE}_k(c) \text{REC}_k(c)}{\text{PRE}_k(c) + \text{REC}_k(c)} \quad (7)$$

The formal version of the corollary stated in the introduction is given below.

Corollary 5 *Let M be an efficient agnostic classification algorithm for $\bar{\mathcal{C}}$. Let $L > 0$ and $f : [0, 1]^{2K} \rightarrow \mathbb{R}$ be L -Lipschitz. For any $\epsilon, \delta > 0$, with probability $\geq 1 - \delta$, Algorithm 1 together with oracles to \mathcal{X}_k and knowledge of $p_k^i = \mu(\mathcal{X}_k \times \{i\})$ can be used to output $c \in \mathcal{C}$ such that,*

$$f(\text{FPR}(c), \text{FNR}(c)) \leq \min_{\bar{c} \in \bar{\mathcal{C}}} f(\text{FPR}(\bar{c}), \text{FNR}(\bar{c})) + \epsilon.$$

Let $p_{\min} = \min_{k,i} p_k^i$. For any constant K , the algorithm uses $\text{poly}(L, 1/\epsilon, 1/\delta, 1/p_{\min})$ total runtime, examples, calls to f , group membership oracles \mathcal{Z}_k , and agnostic learner M . M is called on examples with the same marginal distribution on \mathcal{X} as μ .

As seen from eq. (3), arbitrary nonlinear continuous constraints $g(\text{FPR}(c), \text{FNR}(c)) \leq \epsilon$ can be folded into the objective. The only subtlety in applying Algorithm 1 to this problem is in minimizing weighted combinations of group false positive and false negative rates using an agnostic learner. However, a standard randomized label-flipping trick is used to effectively reweight examples, even with negative weights.

5.2.1. SPECIFIC OBJECTIVES

In this section, we illustrate how existing objectives are group losses. Eq. (3) shows how a continuous constraint $g(l) \leq 0$ (which itself can capture multiple constraints) can be folded into a loss function. If g is satisfiable, then we find c with $g(c) \leq \epsilon$ in time $\text{poly}(1/\epsilon)$. In many learning applications, one cannot guarantee perfect satisfiability. For instance a constraint of equal error rates among men and women can only be approximately checked.

Equal false positive constraints could be modeled as $\sum_{k,k'} |\text{FP}_k(c) - \text{FP}'_{k'}(c)|$, as a maximum over deviations from the mean, or through other approaches. Equal error rates amongst groups could be modeled as $\sum_{k,k'} |\text{ERR}_k(c) - \text{ERR}_{k'}(c)|$, and eq. (4) shows how error rates can be written in terms of false positive and false negative rates.

Since it is known that one cannot simultaneously satisfy many natural notions of fairness (Chouldechova, 2017; Kleinberg et al., 2017), it may be desirable to specify an objective that trades off these various notions. Equations (4)-(7) show how various natural quantities such as error, precision, recall, and the F1 score can be written in terms of the false positive and false negative group rates.

One issue that arises is that the F1 score is discontinuous when the false negative rate is 1 because the recall and precision are 0. However, one can replace F1 by a version which is L -Lipschitz continuous and agrees with F1 outside of a small region in which the false negative rate is close to 1.

Consider the objective in the introduction, eq. (2) but more generally for K groups. It involves the following quantity,

$$\Pr[x \in \mathcal{X}_k \mid c(x) = 1] = \frac{p_k^0 \text{FPR}_k(c) + p_k^1 (1 - \text{FNR}_k(c))}{\sum_j p_j^0 \text{FPR}_j(c) + p_j^1 (1 - \text{FNR}_j(c))},$$

for each group k . Note that this is also discontinuous when the classifier is always 0, however replacing the denominator d by $\max(d, \delta)$ for a small δ makes the function continuous and does not change the objective except when the classifier is nearly trivial in the sense that $c(x) = 0$ almost all the time. A similar approximation can be done for other losses that are not Lipschitz-continuous.

6. Convex Losses

In this section, we use n to denote the number of components to reflect the fact that there may be many components, so $\ell : \mathcal{C} \rightarrow [0, 1]^n$. Narasimhan et al. (2015) show how to optimize convex performance functions of confusion matrices. Our algorithm and analysis in this section extends this to the case of multiple, possibly-overlapping, groups.

As mentioned, the reduction from a linear optimizer to a convex optimizer is a standard Frank-Wolfe style of reduction. In this section, rather than reinvent the wheel, we appeal to an existing result that meets our needs. To do so, we make some further simplifying assumptions. In this section we assume $\bar{\mathcal{C}} = \mathcal{C}$, i.e., the optimizer outputs an element of the comparison set. To use the results of [Kakade et al. \(2009\)](#), we also henceforth assume that $\ell(\mathcal{C}) = \{\ell(c) | c \in \mathcal{C}\}$ is a compact set. This assumption conveniently lets one assume that the minimum of f is attained for some choice $\bar{c} \in \mathcal{C}$.

Consider convex $f : [0, 1]^n \rightarrow [0, 1]$ with a gradient or subgradient oracle, denoted $\nabla f : [0, 1]^n \rightarrow \mathbb{R}^n$, that computes a subgradient at each point l . The goal is to minimize $f(\ell(c))$ over $c \in \mathcal{C}$, possibly subject to constraints. For differentiable f , the gradient is the unique tangent plane at every point. More generally, every convex function (differentiable or non-differentiable) has at least one sub-gradient at every point. Hence, we assume that we have a subgradient oracle $\nabla f : [0, 1]^n \rightarrow [0, 1]^n$ which is just the gradient of f if it is differentiable, and more generally satisfies,

$$\nabla f(l) \cdot (l' - l) \leq f(l') - f(l) \text{ for all } l, l' \in [0, 1]^n. \quad (8)$$

Corollary 7 adds convex constraints to the problem.

Observation 6 *Let $f : [0, 1]^n \rightarrow \mathbb{R}$ be a (nondecreasing) convex function with subgradient oracle $\nabla f : [0, 1]^n \rightarrow [-1, 1]^n$ ($\nabla f : [0, 1]^n \rightarrow [0, 1]^n$). For any $\epsilon > 0$, Algorithm 3.1 of [Kakade et al. \(2009\)](#) with (nonnegative) linear optimizer M_τ finds $c \in \mathcal{C}$ such that,*

$$f(\ell(c)) \leq \min_{\bar{c} \in \mathcal{C}} f(\ell(\bar{c})) + \epsilon.$$

The runtime and number of calls to M_τ , ℓ , and ∇f are $\text{poly}(1/\epsilon, n)$ with $\tau = \epsilon/6$.

Algorithm 3.1 of [Kakade et al. \(2009\)](#) is an online algorithm that can be used to minimize a sequence of functions f_1, \dots, f_T , with a regret bound ϵ meaning that the algorithm chooses a sequence of points c_1, c_2, \dots, c_T such that,

$$\frac{1}{T} \sum_{t=1}^T f_t(c_t) \leq \min_{\bar{c} \in \mathcal{C}} \frac{1}{T} \sum_{t=1}^T f_t(\bar{c}) + \epsilon.$$

Taking $f_1 = f_t = \dots = f$, convexity of f gives a bound on $f(c) \leq \min_{\bar{c} \in \mathcal{C}} f(\bar{c}) + \epsilon$ for $c = \text{UniformDist}(\{c_1, \dots, c_T\})$. However, each f_t in Algorithm 3.1 must be linear. To linearize, we take each f_t to be the tangent plane of f at c_{t-1} which is a lower-bound on f . The proof of Observation 6 is found in Section 8.4 of the appendix.

To fold constraints into the objective f , one needs to be able to compute subgradients. This is made explicit in the following corollary.

Corollary 7 *Let $f, g : [0, 1]^n \rightarrow \mathbb{R}$ be (nondecreasing) convex functions with subgradient oracles $\nabla f, \nabla g : [0, 1]^n \rightarrow [-1, 1]^n$ ($[0, 1]^n$) such that $\exists \bar{c} \in \mathcal{C}$ with $g(\ell(\bar{c})) \leq 0$. For any $\epsilon > 0$, Algorithm 3.1 of [Kakade et al. \(2009\)](#) with (nonnegative) linear optimizer M_τ can be used to find $c \in \mathcal{C}$ such that $g(c) \leq \epsilon$ and*

$$f(\ell(c)) \leq \min_{\bar{c} \in \mathcal{C}: g(\ell(\bar{c})) \leq 0} f(\ell(\bar{c})) + \epsilon.$$

The runtime and number of calls to M_τ , ℓ , and $\nabla f, \nabla g$ are $\text{poly}(1/\epsilon, n)$ and $\tau = \epsilon^2/(12\epsilon + 6n)$.

Proof Let $\lambda = \epsilon/(2\epsilon + n)$. We apply Observation 6 to minimize $h(l) = \lambda f(l) + (1 - \lambda) \max(0, g(l))$ to within $\epsilon' = \lambda\epsilon = \epsilon^2/(2\epsilon + n)$. From eq. (3), since the range of λf is at most $R \leq \lambda n$ (since the diameter of $[0, 1]^n$ is \sqrt{n} and $\|\nabla f(l)\| \leq \sqrt{n}$), we have that an ϵ' optimal minimizer of h have

$$g(l) \leq \frac{\epsilon' + \lambda n}{1 - \lambda} = \epsilon.$$

Note that $\mathbf{1}[g(l) \geq 0]\nabla g(l)$ is a subgradient of $\max(0, g(l))$ if $\nabla g(l)$ is a subgradient of $g(l)$. Thus we can efficiently compute subgradients ∇h as well. Clearly $\nabla h(l) \in [-1, 1]^n$ ($[0, 1]^n$) as well. ■

Using the same approach, constraints can be folded into the objectives in the other theorems and corollaries in this paper.

6.1. Applications

We now highlight two applications of the observation and corollary in the previous section. One convex loss would be,

$$f(l) = \lambda \sum_i l_i + (1 - \lambda) \max_i l_i,$$

which has a parameter $\lambda \in [0, 1]$. For $\lambda = 0$, this is the maximum in the spirit of max-min fairness Rawls (1971). For $\lambda = 1$, this is average performance. Since max is nondifferentiable, there are many subgradients $\nabla f : [0, 1]^n \rightarrow \mathbb{R}^n$, and any will do. One is

$$\nabla f_i(l) = \lambda + (1 - \lambda)\mathbf{1}\left[l_i = \max_j l_j\right].$$

Also note that this objective is nondecreasing. Hence, a FLOP algorithm could be used to minimize this objective which, for $\lambda \in [0, 1]$, trades off the average customer's distance to nearest facility with the maximum customer's (expected) distance to nearest facility.

Finally, consider the Gini index of inequality (Gini, 1912),

$$G(l) = \frac{\sum_{i,j} |l_i - l_j|}{2n \sum l_i} \in [0, 1].$$

While the Gini index is not convex itself, it is quasi-convex meaning that it's level sets are convex. In particular, for any given $\theta \in [0, 1]$, $G(l) \leq \theta$ is equivalent to the convex constraint,

$$\sum_{i,j} |l_i - l_j| - 2n\theta \sum l_i \leq 0.$$

This observation can be used to minimize any given quantity such as $\lambda \sum_i l_i + (1 - \lambda)G(l)$. In particular, one can solve the constrained optimization problem,

$$\min_{l \in F: \sum |l_i - l_j| - 2n\theta \sum l_i \leq 0} \sum_i l_i,$$

for a fine grid of $\theta \in [0, 1]$. Each constrained problem for $\theta \in [0, 1]$ can each be solved using Corollary 7, and the resulting solution can be checked to see if $G(l) \leq \theta$.

7. Conclusions

This paper enables application designers and data scientists to optimize complex continuous objectives, with possibly many local non-global minima, that account for differences across various groups. It is not our claim or belief that important issues such as fairness can be captured solely by such group statistics. However, with this ability, domain experts are hopefully better equipped to design objectives that, when optimized, lead to desirable outcomes.

Acknowledgments. We would like to thank Miro Dudk, Salil Vadhan, and the anonymous COLT reviewers for helpful comments.

References

- Alekh Agarwal, Alina Beygelzimer, Miroslav Dudk, and John Langford. A reductions approach to fair classification. In *FATML*, 2017.
- Alekh Agarwal, Alina Beygelzimer, Miroslav Dudk, John Langford, and Hanna Wallach. A reductions approach to fair classification. In *ICML*, 2018.
- Solon Barocas and Andrew D Selbst. Big data’s disparate impact. *Cal. L. Rev.*, 104:671, 2016.
- Róbert Busa-Fekete, Balázs Szörényi, Paul Weng, and Shie Mannor. Multi-objective bandits: Optimizing the generalized gini index. In *International Conference on Machine Learning*, pages 625–634, 2017.
- Alexandra Chouldechova. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big data*, 5(2):153–163, 2017.
- Kimberle Crenshaw. Demarginalizing the intersection of race and sex: A black feminist critique of antidiscrimination doctrine, feminist theory and antiracist politics. *U. Chi. Legal F.*, page 139, 1989.
- Ilias Diakonikolas. *Approximation of Multiobjective Optimization Problems*. PhD thesis, New York, NY, USA, 2011. AAI3454077.
- Tammy Drezner, Zvi Drezner, and Jeffery Guyse. Equitable service by a facility: Minimizing the gini coefficient. *Computers & Operations Research*, 36(12):3240–3246, 2009.
- Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference, ITCS ’12*, pages 214–226, New York, NY, USA, 2012. ACM.
- Cynthia Dwork, Nicole Immorlica, Adam Tauman Kalai, and Mark DM Leiserson. Decoupled classifiers for group-fair and efficient machine learning. In Sorelle A. Friedler and Christo Wilson, editors, *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, volume 81 of *Proceedings of Machine Learning Research*, pages 119–133, New York, NY, USA, 23–24 Feb 2018. PMLR.
- Corrado Gini. Variabilità e mutabilità. *Reprinted in Memorie di metodologica statistica (Ed. Pizetti E, Salvemini, T). Rome: Libreria Eredi Virgilio Veschi, 1912.*

- Gabriel Goh, Andrew Cotter, Maya Gupta, and Michael P Friedlander. Satisfying real-world goals with dataset constraints. In *Advances in Neural Information Processing Systems 29*, pages 2415–2423. 2016.
- Fabrizio Grandoni, R. Ravi, Mohit Singh, and Rico Zenklusen. New approaches to multi-objective optimization. *Math. Program.*, 146(1-2):525–554, 2014.
- David Ingold and Spencer Soper. Amazon doesnt consider the race of its customers. should it?, April 2016. URL <https://www.bloomberg.com/graphics/2016-amazon-same-day/>.
- Martin Jaggi. Revisiting frank-wolfe: Projection-free sparse convex optimization. In *ICML (1)*, pages 427–435, 2013.
- Matthew Joseph, Michael J. Kearns, Jamie H. Morgenstern, and Aaron Roth. Fairness in learning: Classic and contextual bandits. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 325–333, 2016.
- Sham M. Kakade, Adam Tauman Kalai, and Katrina Ligett. Playing games with approximation algorithms. *SIAM J. Comput.*, 39(3):1088–1106, 2009.
- Adam Tauman Kalai and Santosh Vempala. Efficient algorithms for online decision problems. *J. Comput. Syst. Sci.*, 71(3):291–307, 2005.
- Adam Tauman Kalai, Yishay Mansour, and Elad Verbin. On agnostic boosting and parity learning. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 629–638, 2008.
- Michael J. Kearns, Robert E. Schapire, and Linda M. Sellie. Toward efficient agnostic learning. In *Proceedings of the Fifth Annual Workshop on Computational Learning Theory, COLT '92*, pages 341–352, New York, NY, USA, 1992. ACM.
- Michael J. Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. *CoRR*, abs/1711.05144, 2017.
- Jon M. Kleinberg, Sendhil Mullainathan, and Manish Raghavan. Inherent trade-offs in the fair determination of risk scores. In *8th Innovations in Theoretical Computer Science Conference, ITCS 2017, January 9-11, 2017, Berkeley, CA, USA*, pages 43:1–43:23, 2017.
- Aditya Krishna Menon and Robert C. Williamson. The cost of fairness in classification. *CoRR*, abs/1705.09055, 2017.
- Wolfgang Mulzer and Yannik Stein. Computational aspects of the colorful carathéodory theorem. *arXiv preprint arXiv:1412.3347*, 2014.
- Harikrishna Narasimhan, Harish G. Ramaswamy, Aadirupa Saha, and Shivani Agarwal. Consistent multiclass algorithms for complex performance measures. In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, pages 2398–2407, 2015.

Włodzimierz Ogryczak, Hanan Luss, Michal Pióro, Dritan Nace, and Artur Tomaszewski. Fair optimization and networks: A survey. *J. Applied Mathematics*, 2014:612018:1–612018:25, 2014. URL <https://doi.org/10.1155/2014/612018>.

Christos H Papadimitriou and Mihalis Yannakakis. On the approximability of trade-offs and optimal access of web sources. In *Foundations of Computer Science, 2000. Proceedings. 41st Annual Symposium on*, pages 86–92. IEEE, 2000.

Shameem Puthiya Parambath, Nicolas Usunier, and Yves Grandvalet. Optimizing f-measures by cost-sensitive classification. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 2123–2131, 2014.

R. Ravi, Madhav V. Marathe, S. S. Ravi, Daniel J. Rosenkrantz, and Harry B. Hunt III. Many birds with one stone: multi-objective approximation algorithms. In *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing, May 16-18, 1993, San Diego, CA, USA*, pages 438–447, 1993.

John Rawls. A theory of justice. 1971.

Hinrich Schütze, Christopher D Manning, and Prabhakar Raghavan. *Introduction to information retrieval*, volume 39. Cambridge University Press, 2008.

Leslie G Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.

Arthur Warburton. Approximation of pareto optima in multiple-objective, shortest-path problems. *Operations Research*, 35(1):70–79, 1987.

John A Weymark. Generalized gini inequality indices. *Mathematical Social Sciences*, 1(4):409–430, 1981.

Blake Woodworth, Suriya Gunasekar, Mesrob I. Ohannessian, and Nathan Srebro. Learning non-discriminatory predictors. In *Proceedings of the 2017 Conference on Learning Theory*, volume 65, pages 1920–1953, Amsterdam, Netherlands, 07–10 Jul 2017.

Marcela Zuluaga, Guillaume Sergent, Andreas Krause, and Markus Püschel. Active learning for multi-objective optimization. In *Proceedings of the 30th International Conference on Machine Learning, ICML 2013, Atlanta, GA, USA, 16-21 June 2013*, pages 462–470, 2013.

8. Appendix

Below we present the proofs of our theorems.

8.1. Proof of Theorem 3

For some readers, intuition may be gained by first reading the proof for the case $N = 0$, i.e., where f may be decreasing in places so M_τ can take negative weights, and $\hat{l}_t = \frac{1}{t}(\ell_\tau(c_1) + \dots + \ell_\tau(c_t))$. We first show that for any grid point r , the inner loop of the algorithm does approximately minimize $\|\hat{l}_T - r\|$. If $\bar{C} = C$, this lemma essentially says that \hat{l}_T is very close to the projection of r onto $\ell(C)$.

Lemma 8 Fix any grid point r visited in the Algorithm 1 GROUPOPT. Let $s_t = \ell(c_t)$ and $l_t = \max(\frac{1}{t}(s_1 + \dots + s_t), Nr)$ for $t = 1, 2, \dots, T$. Then,

$$(l_t - r)^2 \leq \min_{\bar{c} \in \bar{\mathcal{C}}} (\ell(\bar{c}) - r)^2 + 4\tau\sqrt{K} + \frac{K(1 + \ln t)}{t}.$$

Proof Since $\max(Nr, \cdot)$ is a contraction map,

$$\begin{aligned} \|l_t - \hat{l}_t\| &= \left\| \max\left(Nr, \frac{1}{t}(s_1 + \dots + s_t)\right) - \max\left(Nr, \frac{1}{t}(\ell_\tau(c_1) + \dots + \ell_\tau(c_t))\right) \right\| \\ &\leq \left\| \frac{1}{t}(s_1 + \dots + s_t) - \frac{1}{t}(\ell(c_1) + \dots + \ell(c_t)) \right\| \\ &\leq \frac{1}{t}(\|s_1 - \ell(c_1)\| + \dots + \|s_t - \ell(c_t)\|) \leq \tau, \end{aligned} \quad (9)$$

using $\|s_i - \ell_\tau(c_i)\| \leq \tau$ by definition of ℓ_τ . We next claim that for $t > 1$,

$$(l_t - r)^2 \leq \left(\left(1 - \frac{1}{t}\right) (l_{t-1} - r) + \frac{1}{t}(s_t - r) \right)^2. \quad (10)$$

For $N = 0$, the above is an equality because $l_t = \left(1 - \frac{1}{t}\right) l_{t-1} + \frac{1}{t} s_t$. For $N = 1$, one can verify the above coordinate-wise using the fact that for any $x, y \in \mathbb{R}$,

$$(\max(x + y, 0))^2 \leq (\max(x + y, y))^2 = (\max(x, 0) + y)^2.$$

To get (10), sum over coordinates $k \in [K]$ and substitute $x = (s_{1k} + \dots + s_{(t-1)k} - (t-1)r_k)/t$ and $y = (1/t)(s_{tk} - r_k)$ for each k .

Let $u = \ell(\bar{c})$ for any $\bar{c} \in \arg \min_{\bar{c}} (\ell(\bar{c}) - r)^2$, i.e., u is the closest point to r in $\ell(\bar{\mathcal{C}})$. By definition of M_τ ,

$$\begin{aligned} (l_{t-1} - r)(s_t - u) &= (\hat{l}_{t-1} - r)(s_t - u) + (l_{t-1} - \hat{l}_{t-1})(s_t - u) \\ &\leq \tau \|\hat{l}_{t-1} - r\| + \|l_{t-1} - \hat{l}_{t-1}\| \cdot \|s_t - u\| \end{aligned}$$

But since $\|l_{t-1} - \hat{l}_{t-1}\| \leq \tau$ and $\hat{l}_t, l_t, r, u, s_t \in [0, 1]^K$, and the diameter of $[0, 1]^K$ is \sqrt{K} , the above is at most $2\tau\sqrt{K}$, or written differently:

$$(l_{t-1} - r)s_t \leq (l_{t-1} - r)u + 2\tau\sqrt{K}. \quad (11)$$

Next, (10) can be rewritten as,

$$\begin{aligned}
 (l_t - r)^2 &\leq \left(l_{t-1} - r + \frac{1}{t}(s_t - l_{t-1}) \right)^2 \\
 &= (l_{t-1} - r)^2 + \frac{2}{t}(l_{t-1} - r)(s_t - l_{t-1}) + \frac{1}{t^2}(s_t - l_{t-1})^2 \\
 &\leq (l_{t-1} - r)^2 + \frac{2}{t}(l_{t-1} - r)(s_t - l_{t-1}) + \frac{K}{t^2} && // \text{ since } s_t, l_{t-1} \in [0, 1]^K \\
 &\leq (l_{t-1} - r)^2 + \frac{2}{t}(l_{t-1} - r)(u - l_{t-1}) + \frac{2}{t}2\tau\sqrt{K} + \frac{K}{t^2} && // \text{ by (11)} \\
 &\leq \left(l_{t-1} - r + \frac{1}{t}(u - l_{t-1}) \right)^2 + \frac{2}{t}2\tau\sqrt{K} + \frac{K}{t^2} \\
 &= \left(\left(1 - \frac{1}{t} \right) (l_{t-1} - r) + \frac{1}{t}(u - r) \right)^2 + \frac{4\tau\sqrt{K}}{t} + \frac{K}{t^2} \\
 &\leq \left(1 - \frac{1}{t} \right) (l_{t-1} - r)^2 + \frac{1}{t}(u - r)^2 + \frac{4\tau\sqrt{K}}{t} + \frac{K}{t^2} && // \text{ by convexity of } (\cdot)^2
 \end{aligned}$$

Multiplying both sides by t ,

$$t(l_t - r)^2 \leq (t-1)(l_{t-1} - r)^2 + (u - r)^2 + 4\tau\sqrt{K} + \frac{K}{t}$$

To finish, the above implies, by induction on $\alpha_t = t(l_t - r)^2$ over t ,

$$\alpha_t = t(l_t - r)^2 \leq t(u - r)^2 + 4\tau t\sqrt{K} + K \left(1 + \frac{1}{2} + \dots + \frac{1}{t} \right).$$

(Note the base case $\alpha_1 = (l_1 - r)^2 \leq K$.) Dividing both sides by t and using the fact that $H_t \leq 1 + \ln(t)$ gives the lemma. \blacksquare

As mentioned, the above lemma is essentially a re-analysis of the Frank-Wolfe algorithm (Jaggi, 2013) that handles a comparison set \bar{C} .

Proof [Theorem 3] Let $\bar{l} \in \arg \min_{\bar{l} \in \ell(\bar{C})} f(\bar{l})$ be any optimum. Let \bar{r} be a grid point that is within $q\sqrt{K}/2 = \beta/2$ of \bar{l} , i.e., $\|\bar{r} - \bar{l}\| \leq \beta/2$, which must exist since the grid spacing is q .

We first argue that the outer for loop will not pass \bar{r} , i.e., it will return on or before the iteration $r = \bar{r}$. To see this, suppose it reaches $r = \bar{r}$. For l_t as defined in the statement of Lemma 8, we have that at the end of inner loop,

$$(l_T - \bar{r})^2 \leq \beta^2/4 + 4\tau\sqrt{K} + K(1 + \ln(T))/T \leq \beta^2 + 4\beta^2 + 2\beta^2 = 7\beta^2,$$

since $4\tau\sqrt{K} = 4\beta^2$ and by our choice of T , $K(1 + \ln(T))/T \leq 2\beta^2$. Again using $\|l_T - \hat{l}_T\| \leq \tau$ from (9), this implies,

$$\|\hat{l}_T - \bar{r}\| \leq \|l_T - \bar{r}\| + \|l_T - \hat{l}_T\| \leq \sqrt{7}\beta + \tau \leq 3\beta,$$

since $\tau \leq \beta/4$ for $\epsilon \leq 1$.

Hence we conclude that the algorithm will output some distribution on or before it reaches $r = \bar{r}$. Fix the r for which it returns and let c be the uniform distribution over c_1, \dots, c_T returned. Let s_t, l_t

be as defined in the statement of Lemma 8, so $\ell(c) = (s_1 + \dots + s_T)/T$ and $l_T = \max(Nr, \ell(c))$. Since the grid points are sorted, $f(r) \leq f(\bar{r})$.

To analyze $f(\ell(c))$, the definition of Lipschitz implies, $f(l_T) - f(r) \leq \|r - l_T\|$, which is $\leq \|\hat{l}_T - r\| + \|\hat{l}_T - l_T\| \leq 3\beta + \tau$. Now, if $N = 0$ then $\ell(c) = l_T$ and $f(\ell(c)) \leq f(l_T)$. If $N = 1$ then $\ell(c) \leq l_T$, and if f is nondecreasing then $f(\ell(c)) \leq f(l_T)$. By the Lipschitz condition, and again using the inequality $\|l_T - r\| \leq 3\beta + \tau$,

$$f(\ell(c)) \leq f(l_T) \leq f(r) + \|l_T - r\| \leq f(\bar{r}) + 3\beta + \tau \leq f(\bar{l}) + \|\bar{l} - \bar{r}\| + 3\beta + \tau.$$

Since $\|\bar{r} - \bar{l}\| \leq \beta/2$ and $3.5\beta + \tau < \epsilon$, this would seem to establish $f(\ell(c)) \leq f(\bar{l}) + \epsilon$ as needed. However, there is a subtle case not covered above in which $N = 1$ because f happens to be nondecreasing on the grid but f may be in fact decreasing in places. It is not hard to see that there must be another grid point $r' \leq r$ which is within $3.5\beta + \tau$ of $\ell(c)$, since every point is within $\beta/2$ of some grid point. Moreover, since f is nondecreasing on the grid, $f(r') \leq f(r)$ and the analysis goes through as above with an additional $\beta/2$, which is okay because $4\beta + \tau < \epsilon$.

To bound the runtime, note that $|G| \leq q^{-K} = (5/\epsilon)^K K^{K/2}$, $T = 25K/\epsilon^2 \ln(25K/\epsilon^2)$, and the number of iterations of the inner loop is at most $|G|T$. The number of calls to each oracle is $\leq |G|T$ (assuming we memoize calls to ℓ_τ). \blacksquare

8.2. Proof of Corollary 5

Proof WLOG $L = 1$. For $L > 1$, we can apply the theorem to f/L with accuracy ϵ/L , and f/L is Lipschitz. We apply Theorem 3 to the $K' = 2K$ groups $\mathcal{Z}_k^i = \mathcal{X}_k \times \{i\} \subseteq \mathcal{Z}$ for $k \in [K]$, $i \in \{0, 1\}$ and loss

$$\ell_k^i(c) = \Pr[c(x) \neq y \mid (x, y) \in \mathcal{Z}_k^i].$$

Note that $\ell_k^0(c) = \text{FPR}_k(c)$ and $\ell_k^1(c) = \text{FNR}_k(c)$. Hence, all that remains is to show how to implement ℓ_τ and M_τ for the $Q = \text{poly}(1/\epsilon)$ total calls to ℓ_τ and M_τ for classification and regression, for $\tau = \epsilon^2/(25\sqrt{K})$. By Chernoff bounds, with $\text{poly}(\log 1/\delta, 1/\epsilon, 1/p_{\min})$ examples, clearly we can estimate each of the T calls to ℓ_τ to within τ accuracy with probability $\geq 1 - \delta/Q$. This is because an expected $1/p_{\min}$ examples suffice to get an example from any group $\mathcal{Z}_{k,i}$.

We now describe how to simulate $M_\tau(w)$ for any $w = \langle w_k^i \rangle_{k,i} \in \mathbb{R}^{2K}$. Let normalization $W = \sum_{k,i} |w_k^i|/p_k^i$, and define $s(x, y, w) = s(z, w) = (1/W) \sum_{k:z \in \mathcal{Z}_k^i} w_k^i/p_k^i \in [-1, 1]$. In these terms, $w \cdot \ell(c) = \mathbb{E}[W s(z, w) | c(x) - y]$ and hence linear optimizer $M_\tau(w)$ must output c such that,

$$\mathbb{E}[s(x, y, w) | c(x) - y] \leq \min_{\bar{c} \in \bar{\mathcal{C}}} \mathbb{E}[s(x, y, w) | \bar{c}(x) - y] + \frac{\tau \|w\|}{W}. \quad (12)$$

We use a standard label-flipping trick to simulate positive or negative weights. In particular, consider the distribution μ' which samples x, y from μ and outputs x, y' where $y' = y$ with probability $(1 + s(x, y, w))/2$ and $y' = 1 - y$ otherwise. Then, for any $c \in \mathcal{C}$, the error with respect to μ' can be written as

$$\begin{aligned} \mathbb{E}_{x, y' \sim \mu'} [|c(x) - y'|] &= \mathbb{E}_{x, y \sim \mu} \left[\frac{1 + s(x, y, w)}{2} |c(x) - y| + \frac{1 - s(x, y, w)}{2} (1 - |c(x) - y|) \right] \\ &= \mathbb{E}_{x, y \sim \mu} \left[s(x, y, w) |c(x) - y| + \frac{1 - s(x, y, w)}{2} \right] \end{aligned}$$

This means that for any c, \bar{c} ,

$$\text{err}(c, \mu') - \text{err}(\bar{c}, \mu') = \mathbb{E}_{x, y \sim \mu} [s(x, y, w) |c(x) - y| - s(x, y, w) |\bar{c}(x) - y|].$$

So if we feed sufficiently many examples from μ' (which is easy to simulate given examples from μ) into M , with probability $\geq 1 - \delta/Q$, M will output c with $\text{err}(c, \mu') - \text{err}(\bar{c}, \mu') \leq \tau p_{\min}/\sqrt{K}$ which suffices for (12) because $\tau p_{\min}/\sqrt{K} \leq \tau \|w\|/W$ since $W \leq \|w\|_1/p_{\min} \leq \|w\|\sqrt{K}/p_{\min}$.

Because we are only changing labels, the underlying distribution on \mathcal{X} doesn't change. By the union bound, all Q calls to ℓ_τ and M_τ can be answered within τ tolerance with probability $\geq 1 - \delta$ as needed. The calls to the agnostic learning algorithm had accuracy $\epsilon' = \tau p_{\min}/\sqrt{K}$ and confidence $\delta' = \delta/Q$ which means that the total runtime is $\text{poly}(1/\epsilon, 1/p_{\min})$. \blacksquare

Algorithm 2 repeatedly calls a linear classifier to provably minimize squared error over the family of (bounded) linear combinations of classifiers. We also require nonnegative example weights, since minimizing squared error with negative weights requires exponential time in general.

8.3. Regression Group Loss

In this section, we show how agnostic learning can be used to minimize a nondecreasing continuous function f of squared error over groups.

Define the set of finite linear combinations of classifiers with coefficient bound B to be:

$$\mathcal{C}_B = \left\{ \sum_i \alpha_i c_i \mid (\alpha_i, c_i) \in \mathbb{R} \times \mathcal{C}, \sum_i |\alpha_i| \leq B \right\}.$$

We similarly define the comparison class $\bar{\mathcal{C}}_B \subseteq \mathcal{C}_B$ to the analogous set of bounded functions where all classifiers are in $\bar{\mathcal{C}}$. Even though \mathcal{C} consists of only binary classifiers, \mathcal{C}_B has real-valued predictors. Also note that most families of classifiers contain the constant 1 predictor $c(x) = 1$, in which case \mathcal{C}_B can fit a bias term.

Let ρ be a distribution over \mathcal{Z} . Let $\mathcal{Z}_1, \dots, \mathcal{Z}_K \subseteq \mathcal{Z}$ be K groups. For $h \in \mathcal{C}_B$, let,

$$\ell(h) = \left\langle \mathbb{E}_{x, y \sim \rho} [(h(x) - y)^2 \mid (x, y) \in \mathcal{Z}_k] \right\rangle_{k=1}^K.$$

Key to the reduction, Algorithm 2 solves regression using a classification algorithm. With this, Corollary 9 is a straightforward application of Theorem 3 for nondecreasing f , with a choice set of \mathcal{C}_B and comparison set of $\bar{\mathcal{C}}_B$.

Corollary 9 *Let M be an efficient agnostic learner for $\bar{\mathcal{C}}$. Let $B, L > 0$ and $f : \mathbb{R}_+^K \rightarrow \mathbb{R}$ be nondecreasing and L -Lipschitz. For any $\epsilon, \delta \in (0, 1)$, with probability $\geq 1 - \delta$, Algorithms 1 and 2 with knowledge of group probabilities $p_k = \rho(\mathcal{Z}_k)$ can be used to output $h \in \mathcal{C}_B$ such that,*

$$f(\ell(h)) \leq \min_{l \in \bar{\mathcal{C}}_B} f(l) + \epsilon.$$

Let $p_{\min} = \min_k p_k$ be the smallest group probability. For any constant K , the algorithm uses $\text{poly}(B, L, 1/\epsilon, 1/\delta, 1/p_{\min})$ total runtime, examples, calls to group membership oracles \mathcal{Z}_k , and calls to agnostic learner M .

Algorithm 2 CLASSYREGRESSION: Weighted regression using binary classification

Input: Agnostic learner M , weighted examples $Z \in ([0, 1] \times \mathcal{X} \times [-1, 1])^*$, $\epsilon > 0$, $B > 0$

Output: $h : \mathcal{X} \rightarrow \mathbb{R}$

Let $T = \frac{12B^2}{\epsilon} \ln \frac{12B^2}{\epsilon}$, batch size $n = \frac{|Z|}{2T}$.

Pick $c_1 \in \mathcal{C}$, $\alpha_1 = B$ // arbitrary initial classifier

for $t = 1$ **to** T **do**

$h_t = \frac{1}{t}(\alpha_1 c_1 + \dots + \alpha_t c_t)$ // $h_t : \mathcal{X} \rightarrow [-B, B]$

Take $2n$ fresh examples from Z , call them $(w_1, x_1, y_1), \dots, (w_{2n}, x_{2n}, y_{2n})$.

For $i = 1, \dots, n$, let $y'_i = 1$ with probability $\frac{1}{2} + \frac{w_i}{2(B+1)}(y_i - h_t(x_i))$, and $y'_i = 0$ otherwise.

Candidates $_t = \{(B, M(x_1, y'_1, \dots, x_n, y'_n)), (-B, M(x_1, 1 - y'_1, \dots, x_n, 1 - y'_n))\}$

Choose $(\alpha_{t+1}, c_{t+1}) \in \text{Candidates}_t$ with greater $\sum_{i=n+1}^{2n} \alpha_{t+1} c_{t+1}(x_i) w_i (y_i - h_t(x_i))$

end

return h_T

The simple proof follows from Theorem 3 and Lemma 10 which shows that Algorithm 2 reduces regression (with example weights) to agnostic learning.

Lemma 10 *Let σ be a distribution over $[0, 1] \times \mathcal{X} \times [-1, 1]$, M be an agnostic learner for $\bar{\mathcal{C}}, \mathcal{C}$. Then, for any $\epsilon, \delta \in (0, 1)$ and any $B > 0$, Algorithm 2 outputs a predictor $h \in \mathcal{C}_B$ such that, with probability $\geq 1 - \delta$,*

$$\mathbb{E}_{w, x, y \sim \sigma} [w(y - h(x))^2] \leq \min_{\bar{h} \in \bar{\mathcal{C}}_B} \mathbb{E}_{w, x, y \sim \sigma} [w(y - \bar{h}(x))^2] + \epsilon.$$

for poly($B, 1/\epsilon, 1/\delta$) examples and calls to M .

Proof Take $\bar{h} \in \bar{\mathcal{C}}_B$ that minimizes $\mathbb{E}[(\bar{h}(x) - y)^2]$. For the moment, fix any t . Let Candidates $_t = \{(B, c_a), (-B, c_b)\}$. For any $w, x, y \in [0, 1] \times \mathcal{X} \times \{0, 1\}$, imagine picking $y' = 1$ with probability $\frac{1}{2} + \frac{w}{2(B+1)}(y - h_t(x))$ and 0 otherwise, as in step t of the algorithm. Note that for any $c : \mathcal{X} \rightarrow \{0, 1\}$ and any w, x, y ,

$$\begin{aligned} \mathbb{E}_{y'} [|c(x) - y'|] &= \left(\frac{1}{2} + \frac{w}{2(B+1)}(y - h_t(x)) \right) (1 - c(x)) + \left(\frac{1}{2} - \frac{w}{2(B+1)}(y - h_t(x)) \right) c(x) \\ &= \left(\frac{1}{2} + \frac{w}{2(B+1)}(y - h_t(x)) \right) - \frac{w}{(B+1)}(y - h_t(x))c(x) \end{aligned}$$

Note the left term above is independent of h . Let $\bar{c} \in \bar{\mathcal{C}}$ be any comparison class classifier. Using this identity for $c \in \{c_a, \bar{c}\}$, and letting μ be the distribution of x, y' generated for $w, x, y \sim \sigma$,

$$\begin{aligned} \text{err}(c_a, \mu) - \text{err}(\bar{c}, \mu) &= \mathbb{E}_{x, y' \sim \mu} [|c_a(x) - y'| - |\bar{c}(x) - y'|] \\ &= \mathbb{E}_{w, x, y \sim \sigma} \left[\frac{w}{(B+1)}(y - h_t(x))(\bar{c}(x) - c_a(x)) \right] \\ &\leq \mathbb{E}_{w, x, y \sim \sigma} \left[\frac{w}{B}(y - h_t(x))(\bar{c}(x) - c_a(x)) \right] \end{aligned}$$

By the definition of agnostic learner, for $n = \text{poly}(T, 1/\delta)$, with probability $\geq 1 - \delta/(4T)$, this implies that M called on n examples from μ will have $\text{err}(c_a, \mu) - \text{err}(\bar{c}, \mu) \leq 1/2T$, which can be written as:

$$\mathbb{E}_{w,x,y \sim \sigma} [w(y - h_t(x))c_a(x)] \geq \max_{\bar{c} \in \bar{\mathcal{C}}} \mathbb{E}_{w,x,y \sim \sigma} [w(y - h_t(x))\bar{c}(x)] - \frac{B}{2T} \quad (13)$$

An entirely similar argument for the distribution over $(x, 1 - y')$ implies that,

$$\mathbb{E}_{w,x,y \sim \sigma} [-w(y - h_t(x))c_b(x)] \geq \max_{\bar{c} \in \bar{\mathcal{C}}} \mathbb{E}_{w,x,y \sim \sigma} [-w(y - h_t(x))\bar{c}(x)] - \frac{B}{2T} \quad (14)$$

By Chernoff bounds for $n = \text{poly}(B, 1/T, \log 1/\delta)$, with probability $\geq 1 - \delta/(2T)$, the choice for $(\alpha_{t+1}, c_{t+1}) \in \{(B, c_a), (-B, c_b)\}$ which has greater empirical correlation $\sum \alpha_{t+1} c_{t+1}(x_i) w_i(y_i - h_t(x_i))$ on n fresh iid examples from σ will be within B^2/T of the greater correlated of $\{Bc_a, -Bc_b\}$, i.e.,

$$\mathbb{E} [w(y - h_t(x))\alpha_{t+1}c_{t+1}(x)] \geq B \max(\mathbb{E}[w(y - h_t(x))c_a(x)], \mathbb{E}[-w(y - h_t(x))c_b(x)]) - \frac{B^2}{T}$$

Now, the above together with (13) and (14) and the fact that the maximum over a linear function always occurs at vertex, we get that

$$\begin{aligned} \mathbb{E} [w(y - h_t(x))\alpha_{t+1}c_{t+1}(x)] &\geq \max_{(\bar{\alpha}, \bar{c}) \in [-B, B] \times \bar{\mathcal{C}}} \mathbb{E} [w\bar{\alpha}\bar{c}(x)(y - h_t(x))] - \frac{B^2}{T} \\ &= \max_{\bar{h} \in \bar{\mathcal{C}}_B} \mathbb{E}_{x,y \sim \sigma} [w\bar{h}(x)(y - h_t(x))] - \frac{B^2}{T}. \end{aligned}$$

We now stop fixing t and look across t . For notational ease, let $s_t(x) = \alpha_t c_t(x)$ so the above implies:

$$\mathbb{E} [w(y - h_t(x))s_{t+1}(x)] \geq \max_{\bar{h} \in \bar{\mathcal{C}}_B} \mathbb{E}_{x,y \sim \sigma} [w(y - h_t(x))\bar{h}(x)] - \frac{B^2}{T} \quad (15)$$

In these terms, we also have $h_t = (1 - 1/t)h_{t-1} + (1/t)s_t$. In particular, similar to the proof of Lemma 8,

$$\begin{aligned} (h_t(x) - y)^2 &= \left(h_{t-1}(x) - y + \frac{1}{t}(s_t(x) - h_{t-1}(x)) \right)^2 \\ &= (h_{t-1}(x) - y)^2 + \frac{2}{t}(h_{t-1}(x) - y)(s_t(x) - h_{t-1}(x)) + \frac{1}{t^2}(s_t(x) - h_{t-1}(x))^2 \\ &\leq (h_{t-1}(x) - y)^2 + \frac{2}{t}(h_{t-1}(x) - y)(s_t(x) - h_{t-1}(x)) + \frac{4B^2}{t^2} \end{aligned}$$

We now use (15) (negated and for s_t) and the fact that $w \in [0, 1]$:

$$\begin{aligned} \mathbb{E}[w(h_t(x) - y)^2] &\leq \mathbb{E} \left[w(h_{t-1}(x) - y)^2 + \frac{2}{t}w(h_{t-1}(x) - y)(s_t(x) - h_{t-1}(x)) + \frac{4B^2}{t^2} \right] \\ &\leq \mathbb{E} \left[w(h_{t-1}(x) - y)^2 + \frac{2}{t}w(h_{t-1}(x) - y)(\bar{h}(x) - h_{t-1}(x)) + \frac{2B^2}{tT} + \frac{4B^2}{t^2} \right] \\ &\leq \mathbb{E} \left[w(h_{t-1}(x) - y)^2 + \frac{2}{t}w(h_{t-1}(x) - y)(\bar{h}(x) - h_{t-1}(x)) + \frac{6B^2}{t^2} \right]. \end{aligned}$$

We now complete the square in the first two terms above, giving,

$$\begin{aligned} \mathbb{E}[w(h_t(x) - y)^2] &\leq \mathbb{E}\left[w\left((h_{t-1}(x) - y) + \frac{1}{t}(\bar{h}(x) - h_{t-1}(x))\right)^2 + \frac{6B^2}{t^2}\right] \\ &\leq \mathbb{E}\left[w\left(\left(1 - \frac{1}{t}\right)(h_{t-1}(x) - y) + \frac{1}{t}(\bar{h}(x) - y)\right)^2 + \frac{6B^2}{t^2}\right]. \end{aligned}$$

Using the convexity of $(\cdot)^2$, this is at most

$$\begin{aligned} \mathbb{E}[(h_t(x) - y)^2] &\leq \mathbb{E}\left[\left(1 - \frac{1}{t}\right)(h_{t-1}(x) - y)^2 + \frac{1}{t}(\bar{h}(x) - y)^2 + \frac{6B^2}{t^2}\right] \\ t\mathbb{E}[(h_t(x) - y)^2] &\leq \mathbb{E}\left[(t-1)(h_{t-1}(x) - y)^2 + (\bar{h}(x) - y)^2 + \frac{6B^2}{t}\right] \end{aligned}$$

By induction on t with $\beta_t = t\mathbb{E}[(h_t(x) - y)^2]$,

$$\beta_T = T\mathbb{E}[(h_T(x) - y)^2] \leq T\mathbb{E}[(\bar{h}(x) - y)^2] + 6B^2\left(1 + \frac{1}{2} + \dots + \frac{1}{T}\right)$$

Finally, dividing by T , we get,

$$\mathbb{E}[(h_T(x) - y)^2] \leq \mathbb{E}[(\bar{h}(x) - y)^2] + \frac{6B^2}{T}H_T.$$

Using the fact that harmonic sums $H_T \leq 1 + \ln(T)$ gives the inequality in the proof for $T = \frac{12B^2}{\epsilon} \ln \frac{12B^2}{\epsilon}$ and $\epsilon \leq 1$.

For runtime, the two requirements on the batch size were $n = \text{poly}(T, 1/\delta)$ for the agnostic learner and $n = \text{poly}(B, 1/T, \log 1/\delta)$ for Chernoff bounds to determine which of the two predictors was better in each round. Hence, batch size $n = \text{poly}(B, 1/\epsilon, 1/\delta)$ suffices and the total number of examples and calls to M are $\leq 2Tn = \text{poly}(B, 1/\epsilon, 1/\delta)$. \blacksquare

Proof [Corollary 9] Note that the loss $\ell(h) \in [0, (B+1)^2]^K$ as defined since the total weight on classifiers is at most B and the labels are in $[0, 1]$. Theorem 3 can be trivially modified to handle $\ell : \mathcal{C} \rightarrow [0, (B+1)^2]^K$ and $f : [0, (B+1)^2]^K \rightarrow \mathbb{R}$ simply scaling down ℓ .

The proof is very similar to that of Corollary 5 above. Again let Q be the total number of calls to ℓ_τ and M_τ . First, WLOG $L = 1$, or else we can scale down f . Again, it is easy to estimate $\ell_\tau(h)$ using $\text{poly}(B, 1/\epsilon, \log(1/\delta), 1/p_{\min})$ examples by Chernoff bounds, with probability $\geq 1 - \delta/Q$.

To compute $M_\tau(w)$ for $w \in \mathbb{R}_+^K$, again define $W = \sum_k w_k/p_k$ and $s(z, w) = (1/W) \sum_{k: z \in \mathcal{Z}_k} w_k/p_k$. Now $s(z, w) \in [0, 1]$ and again $w \cdot \ell(h) = \mathbb{E}[W s(z, w)(h(x) - y)^2]$. Lemma 10 shows that Algorithm 2 can be used with M to find h with weighed error within $\epsilon' = \tau p_{\min}/\sqrt{K}$ for weight distribution $s(z, w)$, with probability $\geq 1 - \delta/Q$. This means that, compared to any $\bar{h} \in \bar{\mathcal{C}}_B$, h also satisfies,

$$w \cdot (\ell(h) - \ell(\bar{h})) = W \mathbb{E}[s(x, y, w)(h(x) - y)^2 - s(x, y, w)(\bar{h}(x) - y)^2] \leq W \frac{\tau p_{\min}}{\sqrt{K}} \leq \tau \|w\|,$$

which is what is needed to qualify M_τ as a linear optimizer. \blacksquare

Although not stated in Corollary 9, it is easy to see that the marginal distribution on \mathcal{X} for the examples fed into the agnostic learner is the same as that of μ .

8.4. Proof of Observation 6

Kakade et al. (2009) handles a more general problem in three ways. As discussed earlier, it handles an online sequence of functions. Second, it covers arbitrary restrictions on weights, not just non-negative. Third, it covers multiplicative α approximation algorithms. To address this last difference, we will first translate our M_τ with additive τ error into a α -multiplicative approximation algorithm by adding a constant to our loss, so that our additive τ error translates into a multiplicative $(1 + \tau)$ approximation algorithm.

Proof We define their set $\mathcal{W} = \{\frac{2}{3}\} \times [\frac{-1}{3n}, \frac{1}{3n}]^n$ ($\mathcal{W} = \{\frac{2}{3}\} \times [0, \frac{1}{3n}]^n$). Note that $\|w\| \leq 1$ for $w \in \mathcal{W}$. Our feasible choices are \mathcal{C} (which they call decisions \mathcal{S}). Define $\Phi : \mathcal{C} \rightarrow \{1\} \times [0, 1]^n$ by $\Phi(c) = (1, \ell(c))$. Clearly $\|\Phi(c)\| \leq \sqrt{1+n}$ and $\Phi(c) \cdot w \in [1/3, 1]$ for each $c \in \mathcal{C}, w \in \mathcal{W}$. As they further require (for convenience), we have assumed that $\ell(\mathcal{C})$, and hence $\Phi(\mathcal{C})$, is compact.

Define α -approximation algorithm $A : \mathcal{W} \rightarrow \mathcal{S}$ by $A(w) = A(2/3, w') = M_\tau(w)$. For convenience in this proof, for any $w \in \mathcal{W}$, we write $w' \in \mathbb{R}^n$ for the last n coordinates of w . Using the fact that $\Phi(c) \cdot w = 2/3 + \ell(c) \cdot w' \geq 1/3$ for any $w \in \mathcal{W}$, the definition of (nonnegative) linear optimizer guarantees that,

$$\Phi(A(w)) \cdot w \leq \min_{c \in \mathcal{C}} \left(\frac{2}{3} + \ell(c) \cdot w' \right) + \tau \|w'\| \leq (1 + 3\tau \|w'\|) \min_{c \in \mathcal{C}} \left(\frac{2}{3} + \ell(c) \cdot w' \right).$$

Since $\|w'\| \leq 1/3$, we have that A is a α approximation algorithm for $\alpha = 1 + \tau$.

We apply Algorithm 3.1 of Kakade et al. (2009) to the sequence $w_1 = (2/3, 0, 0, \dots, 0)$ and $w_{t+1} = (2/3, \nabla f(\ell(c_t))/3)$ where c_t is the choice output by the t th iteration of their algorithm for T iterations. We output $\text{UniformDist}\{c_1, \dots, c_T\}$. Since by assumption $\nabla f(\ell(c_t)) \in [-1, 1]^n$ ($[0, 1]^n$), $w_t \in \mathcal{W}_t$. Let \bar{c} be a minimizer of $f(\ell(c))$, which exists by compactness. Their Theorem 3.2 guarantees,

$$\frac{1}{T} \sum_{t=1}^T \Phi(c_t) \cdot w_t \leq (\alpha + 2) \sqrt{\frac{1+n}{T}} + \alpha \frac{1}{T} \sum_{t=1}^T \Phi(\bar{c}) \cdot w_t,$$

for any $\bar{c} \in \mathcal{C}$, using at most $4(\alpha + 2)^2 T$ calls to Φ and A (hence M_τ). Since $\frac{1}{T} \sum \Phi(\bar{c}) \cdot w_t \leq 1$ and $(1+n)/T \leq 1$, the above implies,

$$\frac{1}{T} \sum_{t=1}^T (\Phi(c_t) - \Phi(\bar{c})) \cdot w_t \leq (\alpha + 2) \sqrt{\frac{1+n}{T}} + \alpha - 1 = (3 + \tau) \sqrt{\frac{1+n}{T}} + \tau.$$

On the other hand, by convexity of f and the definition of Φ and w_t ,

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T (\Phi(c_t) - \Phi(\bar{c})) \cdot w_t &= \frac{1}{3T} \sum_{t=1}^T (\ell(c_t) - \ell(\bar{c})) \cdot \nabla f(\ell(c_t)) \\ &\geq \frac{1}{3T} \sum_t f(\ell(c_t)) - f(\ell(\bar{c})) \\ &\geq \frac{1}{3} \left(f \left(\frac{\ell(c_1) + \dots + \ell(c_T)}{T} \right) - f(\ell(\bar{c})) \right) \end{aligned}$$

The above is a bound on the output on $f(\ell(c))$ for $c = \text{UniformDist}\{c_1, \dots, c_T\}$. In particular, it shows $f(\ell(c)) - f(\ell(\bar{c})) \leq 3((3 + \tau) \sqrt{\frac{1+n}{T}} + \tau) = \epsilon$ for our choice of $T = 36(3 + \tau)^2(n + 1)/\epsilon^2 = \text{poly}(1/\epsilon, n)$ and $\tau = \epsilon/6$.

Two notes are in order. First, while the APPROX-PROJ algorithm of [Kakade et al. \(2009\)](#) is described as a randomized algorithm and Theorem 3.2 is stated in terms of expected performance, this is only a cosmetic difference. This is because, instead of outputting a probability distribution over choices, they output a choice at random from this distribution and bound its expected performance, which is the same by linearity as the expected performance of the probability distribution. Hence, one can run their algorithm and simply keep track of the choice distribution. Second, they define $\mathcal{W}_+ = \{aw \mid a \geq 0, w \in \mathcal{W}\}$ and require projection oracle onto \mathcal{W}_+ . It is not difficult to see that projection on to the set \mathcal{W}_+ for our \mathcal{W} , can be computed efficiently as required. ■