

Maximizing Welfare with Incentive-Aware Evaluation Mechanisms

Nika Haghtalab¹, Nicole Immorlica², Brendan Lucier² and Jack Z. Wang¹

¹Cornell University

²Microsoft Research

{nika,jackzwang}@cs.cornell.edu, {nicimm,brlucier}@microsoft.com

Abstract

Motivated by applications such as college admission and insurance rate determination, we propose an evaluation problem where the inputs are controlled by strategic individuals who can modify their features at a cost. A learner can only partially observe the features, and aims to classify individuals with respect to a quality score. The goal is to design an evaluation mechanism that maximizes the overall quality score, i.e., welfare, in the population, taking any strategic updating into account.

We further study the algorithmic aspect of finding the welfare maximizing evaluation mechanism under two specific settings in our model. When scores are linear and mechanisms use linear scoring rules on the observable features, we show that the optimal evaluation mechanism is an appropriate projection of the quality score. When mechanisms must use linear thresholds, we design a polynomial time algorithm with a (1/4)-approximation guarantee when the underlying feature distribution is sufficiently smooth and admits an oracle for finding dense regions. We extend our results to settings where the prior distribution is unknown and must be learned from samples.

1 Introduction

Automatic decision making is increasingly used to identify qualified individuals in areas such as education, hiring, and public health. This has inspired a line of work aimed at improving the performance and interpretability of classifiers for *identifying qualification and excellence* within a society given access to limited visible attributes of individuals. As these classifiers become widely deployed at a societal level, they can take on the additional role of *defining excellence and qualification*. That is, classifiers encourage people who seek to be “identified as qualified” to acquire attributes that are “deemed to be qualified” by the classifier. For example, a college admission policy that heavily relies on SAT scores will naturally encourage students to increase their SAT scores, which they might do by getting a better understanding of the core material, taking SAT prep lessons, cheating, etc.

Progress in machine learning has not fully leveraged classifiers’ role in incentivizing people to change their feature attributes, and at times has even considered it an inconvenience to the designer who must now take steps to ensure that their classifier cannot be “gamed” [Meir *et al.*, 2012; Chen *et al.*, 2018; Hardt *et al.*, 2016; Dekel *et al.*, 2010; Cai *et al.*, 2015]. One of the motivations for such *strategy-proof classification* is Goodhart’s law, which states “when a measure becomes a target, it ceases to be a good measure.” Taking Goodhart’s law to heart, one might view an individual’s attributes to be immutable, and any strategic response to a classifier (changes in one’s attributes) only serves to mask an agent’s true qualification and thereby degrade the usefulness of the classifier.

What this narrative does not address is that in many applications of classification, one’s qualifications can truly be improved by changing one’s attributes. For example, students who improve their understanding of core material truly become more qualified for college education. These changes have the potential to raise the overall level of qualification and excellence in a society and should be encouraged by a social planner. In this work, we focus on this powerful and under-utilized role of classification in machine learning and ask how to

design an evaluation mechanism on visible features that incentivizes individuals to improve a desired quality.

For instance, in college admissions, the planner might wish to maximize the “quality” of candidates. Quality is a function of many features, such as persistence, creativity, GPA, past achievements, only a few of which may be directly observable by the planner. Nevertheless the planner designs an admission test on visible features to identify qualified individuals. To pass the test, ideally candidates improve their features and truly increase their quality as a result. In another motivating example, consider a designer who wants to increase average driver safety, which can depend on many features detailing every aspect of a driver’s habits. The designer may only have access to a set of visible features such as age, driver training, or even driving statistics like acceleration/deceleration speed (as recorded by cars’ GPS devices). Then a scoring rule (that can affect a driver’s insurance premium) attempts to estimate and abstract a notion of “safety” from these features. Drivers naturally adjust their behavior to maximize their score. In

both cases, the mechanism does not just pick out high quality candidates or safe drivers in the underlying population, but it actually causes their distributions to change.

Modeling and Contributions. Motivated by these observations, we introduce the following general problem. Agents are described by feature vectors in a high-dimensional feature space, and can change their innate features at a cost. There is a *true* function mapping feature vectors to a true quality score (binary or real-valued). The planner observes a low-dimensional projection of agents’ features and chooses an evaluation mechanism, from a fixed class of mechanisms which map this low-dimensional space to *observed* scores (binary or real-valued). Agents get value from having a high observed score (e.g., getting admitted to university or having a low car insurance premium), whereas the planner wishes to maximize welfare, i.e., the average true score of the population.

To further study this model, we analyze the algorithmic aspects of maximizing welfare in two specific instantiations. In Section 3, we show that when the true quality function is linear and the mechanism class is the set of all linear functions, the optimal is a projection of the true quality function on the visible subspace. Our most interesting algorithmic results (Section 4), arise from the case when the true function is linear and mechanism class is the set of all linear thresholds. In this case, a simple projection does not work: we need to consider the distribution of agents (projected on the visible feature space) when choosing the mechanism. For this case, we provide polynomial time approximation algorithms for finding the optimal linear threshold. In Section 5, we also provide sample complexity guarantees for learning the optimal mechanism from samples only.

Prior Work. Our work builds upon the strategic machine learning literature introduced by Hardt et al. [2016]. As in our work, agents are represented by feature vectors which can be manipulated at a cost. Hardt et al. [2016] design optimal learning algorithms in the presence of these costly strategic manipulations. Hu et al. [2019] and Milli et al. [2019] extend Hardt et al. [2016] by assuming different groups of agents have different costs of manipulation and study the disparate impact on their outcomes. Dong et al. [2018] consider a setting in which the learner does not know the distribution of agents’ features or costs but must learn them through revealed preference. Importantly, in all these works, the manipulations do not change the underlying features of the agent and hence purely disadvantage the learning algorithm. Kleinberg and Raghavan [2019] introduce a different model in which manipulations do change the underlying features. Some changes are advantageous, and the designer chooses a rule that incentivizes these while discouraging disadvantageous changes. Their main result is that simple linear mechanisms suffice for a single known agent (i.e., known initial feature vector). In contrast, we study a population of agents with a known distribution of feature vectors and optimize over the class of linear, or linear threshold, mechanisms. Alon et al. [2020] also study extensions of Kleinberg and Raghavan [2019] to multiple agents. In that work, agents differ in how costly it is for them to manipulate their features but they all have the same

starting feature representation, but in our work, agents differ in their starting features while facing the same manipulation cost.

As noted by Kleinberg and Raghavan [2019], their model is closely related to the field of contract design in economics. The canonical principal-agent model (see, e.g., [Grossman and Hart, 1983; Ross, 1973]) involves a single principle and a single agent. There is a single-dimensional output, say the crop yield of a farm, and the principal wishes to incentivize the agent to expend costly effort to maximize output. However, the principle can only observe output, and the mapping of effort to output is noisy. Under a variety of conditions, the optimal contract in such settings pays the agent an affine function of output [Carroll, 2015; Dütting et al., 2019; Holmstrom and Milgrom, 1987], although the optimal contract in general settings can be quite complex [McAfee and McMillan, 1986]. Our model differs from this canonical literature in that both effort and output are multi-dimensional. In this regard, the work of Holmstrom and Milgrom [1991] is closest to ours. They also study a setting with a multi-dimensional feature space in which the principal observes only a low-dimensional representation. Important differences include that they only study one type of agent whereas we allow agents to have different starting feature vectors, and they assume transferable utility whereas in our setting payments are implicit and do not reduce the utility of the principal.

2 Preliminaries

The Model. We denote the true features of an individual by $\mathbf{x} \in \mathbb{R}^n$, where the feature space \mathbb{R}^n encodes all relevant features of a candidate, e.g., health history, biometrics, vaccination record, exercise and dietary habits. We denote by $f : \mathbb{R}^n \rightarrow [0, 1]$ the mapping from one’s true features to their true quality. For example, a real-valued $f(\mathbf{x}) \in [0, 1]$ can express the overall quality of candidates and a binary-valued $f(\mathbf{x})$ can determine whether \mathbf{x} meets a certain qualification level.

These true features of an individual may not be visible to the designer. Instead there is an $n \times n$ projection matrix P of rank k , i.e., $P^2 = P$, such that for any \mathbf{x} , $P\mathbf{x}$ represents the visible representation of the individual, such as vaccination record and health history, but not exercise and dietary habits. We define by $\text{Img}(P) = \{\mathbf{z} \in \mathbb{R}^n \mid \mathbf{z} = P\mathbf{z}\}$ the set of all feature representations that are visible to the designer. We denote by $g : \mathbb{R}^n \rightarrow \mathbb{R}$ a mechanism whose outcome for any individual \mathbf{x} depends only on $P\mathbf{x}$, i.e., the visible features of \mathbf{x} .¹ E.g., $g(\mathbf{x}) \in \mathbb{R}$ can express the payoff an individual receives from the mechanism, $g(\mathbf{x}) \in [0, 1]$ can express the probability that \mathbf{x} is accepted by a randomized mechanism g , or a binary-valued $g(\mathbf{x}) \in \{0, 1\}$ can express whether \mathbf{x} is accepted or rejected deterministically.

Let $\text{cost}(\mathbf{x}, \mathbf{x}')$ represent the cost that an individual incurs when changing their features from \mathbf{x} to \mathbf{x}' . We consider $\text{cost}(\mathbf{x}, \mathbf{x}') = c\|\mathbf{x} - \mathbf{x}'\|_2$ for some $c > 0$. Given mechanism g , the total payoff \mathbf{x} receives from changing its feature representation to \mathbf{x}' is $U_g(\mathbf{x}, \mathbf{x}') = g(\mathbf{x}') - \text{cost}(\mathbf{x}, \mathbf{x}')$. Let $\delta_g : \mathbb{R}^n \rightarrow \mathbb{R}^n$ denote the best response of an individual to

¹ Equiv. $g(\mathbf{x}) = g^{\parallel}(P\mathbf{x})$ for unrestricted $g^{\parallel} : \mathbb{R}^n \rightarrow \mathbb{R}$.

g , i.e., $\delta_g(\mathbf{x}) = \operatorname{argmax}_{\mathbf{x}'} U_g(\mathbf{x}, \mathbf{x}')$. We consider a distribution \mathcal{D} over feature vector in \mathbb{R}^n , representing individuals. Our goal is to design a mechanism g such that, when individuals from \mathcal{D} best respond to it, it yields the highest quality individuals on average. That is to find a mechanism $g \in \mathcal{G}$ that maximizes

$$\operatorname{Val}(g) = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [f(\delta_g(\mathbf{x}))]. \quad (1)$$

We often consider the gain in the quality of individuals compared to the average quality before deploying any mechanism, i.e., the baseline $\mathbb{E}[f(\mathbf{x})]$, defined by

$$\operatorname{Gain}(g) = \operatorname{Val}(g) - \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [f(\mathbf{x})]. \quad (2)$$

We denote the mechanism with highest gain by $g_{\text{opt}} \in \mathcal{G}$.

For example, f can indicate the overall health of an individual and \mathcal{G} the set of governmental policies on how to set insurance premiums based on the observable features of individuals, e.g., setting lower premiums for those who received preventative care in the previous year. Then, g_{opt} corresponds to the policy that leads to the healthiest society on average. In Sections 3 and 4, we work with different design choices for f and \mathcal{G} , and show how to find (approximately) optimal mechanisms.

Types of Mechanisms. More specifically, we consider a known true quality function f that is a *linear function* $\mathbf{w}_f \cdot \mathbf{x} - b_f$ for some vector $\mathbf{w}_f \in \mathbb{R}^n$ and $b_f \in \mathbb{R}$. Without loss of generality we assume that the domain and cost multiplier c are scaled so that \mathbf{w}_f is a unit vector.

In Section 3, we consider the class of *linear mechanisms* \mathcal{G}_{lin} , i.e., any $g \in \mathcal{G}_{\text{lin}}$ is represented by $g(\mathbf{x}) = \mathbf{w}_g \cdot \mathbf{x} - b_g$, for a scalar $b_g \in \mathbb{R}$ and vector $\mathbf{w}_g \in \operatorname{Img}(P)$ of length $\|\mathbf{w}_g\|_2 \leq R$ for some $R \in \mathbb{R}^+$. In Section 4, we consider the class of *linear threshold (halfspace) mechanisms* \mathcal{G}_{0-1} , where a $g \in \mathcal{G}_{0-1}$ is represented by $g(\mathbf{x}) = \operatorname{sign}(\mathbf{w}_g \cdot \mathbf{x} - b_g)$ for some unit vector $\mathbf{w}_g \in \operatorname{Img}(P)$ and scalar $b_g \in \mathbb{R}$.

Other Notation. $\|\cdot\|$ is the L_2 norm of a vector unless otherwise stated. For $\ell \geq 0$, the ℓ -margin density of a halfspace is $\operatorname{Den}_{\mathcal{D}}^{\ell}(\mathbf{w}, b) = \Pr_{\mathbf{x} \sim \mathcal{D}} [\mathbf{w} \cdot \mathbf{x} - b \in [-\ell, 0]]$. This is the total density \mathcal{D} assigns to points that are at distance at most ℓ from being included in the positive side of the halfspace. The *soft ℓ -margin density* of this halfspace is defined as

$$\operatorname{S-Den}_{\mathcal{D}}^{\ell}(\mathbf{w}, b) = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [(b - \mathbf{w} \cdot \mathbf{x}) \mathbf{1}(\mathbf{w} \cdot \mathbf{x} - b \in [-\ell, 0])].$$

We suppress ℓ and \mathcal{D} when the context is clear.

We assume that individuals have features that are within a ball of radius r of the origin, i.e., \mathcal{D} is only supported on $\mathcal{X} = \{\mathbf{x} \mid \|\mathbf{x}\| \leq r\}$. In Section 4, we work with distribution \mathcal{D} that is additionally σ -smooth. \mathcal{D} is σ -smoothed distribution if there is a corresponding distribution \mathcal{P} over \mathcal{X} such that to sample $\mathbf{x}' \sim \mathcal{D}$ one first samples $\mathbf{x} \sim \mathcal{P}$ and then $\mathbf{x}' = \mathbf{x} + N(0, \sigma^2 I)$. Smoothing is a common assumption in theory of computer science where $N(0, \sigma^2 I)$ models uncertainties in real-life measurements. To ensure that the noise in measurements is small compared to radius of the domain r , we assume that $\sigma \in O(r)$.

3 Linear Mechanisms

In this section, we show how the optimal linear mechanism in \mathcal{G}_{lin} is characterized by $g_{\text{opt}}(\mathbf{x}) = \mathbf{w}_g \cdot \mathbf{x}$ for \mathbf{w}_g that is (the largest vector) in the direction of $P\mathbf{w}_f$. This leads to an algorithm with $O(n)$ runtime for finding the optimal mechanism. At a high level, this result shows that when the true quality function f is linear, the optimal linear mechanism is in the direction of the closest vector to \mathbf{w}_f in the visible feature space. Indeed, this characterization extends to any true quality function that is a monotonic transformation of a linear function.

Theorem 1 (Linear Mechanisms). *Let $f(\mathbf{x}) = h(\mathbf{w}_f \cdot \mathbf{x} - b_f)$ for some monotonic function $h : \mathbb{R} \rightarrow \mathbb{R}$. Let $g(\mathbf{x}) = \frac{(P\mathbf{w}_f)_R}{\|P\mathbf{w}_f\|_2} \cdot \mathbf{x}$. Then, g has the optimal gain.*

It is interesting to note that designing an optimal *linear* mechanism (Theorem 1) does not use any information about the distribution of instances in \mathcal{D} , rather directly projects \mathbf{w}_f on the subspace of visible features. We see in Section 4 that this is not the case for *linear threshold* mechanisms, where the distribution of feature attributes plays a central role in choosing the optimal mechanism.

4 Linear Threshold Mechanisms

In this section, we consider the class of linear threshold mechanisms \mathcal{G}_{0-1} and explore the computational aspects of finding $g_{\text{opt}} \in \mathcal{G}_{0-1}$. Note that any $g(\mathbf{x}) = \operatorname{sign}(\mathbf{w}_g \cdot \mathbf{x} - b_g)$ corresponds to the transformation of a linear mechanism from \mathcal{G}_{lin} that only rewards those whose quality passes a threshold. As in the case of linear mechanisms, individuals only move in the direction of \mathbf{w}_g and every unit of their movement improves the payoff of the mechanism by $\mathbf{w}_g \cdot \mathbf{w}_f$. However, an individual's payoff from the mechanism improves only if its movement pushes the feature representation from the 0-side of a linear threshold to +1-side. Therefore only those individuals whose feature representations are close to the decision boundary move.

Lemma 1. *For $g(\mathbf{x}) = \operatorname{sign}(\mathbf{w}_g \cdot \mathbf{x} - b_g) \in \mathcal{G}_{0-1}$,*

$$\delta_g(\mathbf{x}) = \mathbf{x} - \mathbf{1}\left(\mathbf{w}_g \cdot \mathbf{x} - b_g \in \left[-\frac{1}{c}, 0\right]\right) (\mathbf{w}_g \cdot \mathbf{x} - b_g) \mathbf{w}_g.$$

This leads to very different dynamics and challenges compared to the linear mechanism case. For example, $\mathbf{w}_g \propto P\mathbf{w}_f$ no longer leads to a good mechanism as it may only incentivize a small fraction of individuals as shown in the following example.

Example 1. *As in Figure 1, consider $\mathbf{w}_f = \left(\frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}}\right)$ and P that projects a vector on its second and third coordinates. Consider a distribution \mathcal{D} that consists of $\mathbf{x} = (0, 0, x_3)$ for $x_3 \sim \operatorname{Unif}[-r, r]$. Any g only incentivizes individuals who are at distance at most $\ell = 1/c$ from the decision boundary, highlighted by the shaded regions. Mechanism $g(\mathbf{x}) = \operatorname{sign}(x_2 - \ell)$ incentivizes everyone to move ℓ unit in the direction of x_2 and leads to total utility of $\mathbb{E}[f(\delta_g(\mathbf{x}))] = \ell/\sqrt{3}$. But, $g'(\mathbf{x}) = \operatorname{sign}(\mathbf{w}'_g \cdot \mathbf{x} - b'_g)$ for unit vector $\mathbf{w}'_g \propto P\mathbf{w}_f$ only incentivizes $\sqrt{2}\ell/2r$ fraction of the individuals, on average each moves only $\ell/2 \cdot \sqrt{2/3}$ units in*




Figure 1: Mechanism $g'(\mathbf{x}) = \text{sign}(\mathbf{w}'_g \cdot \mathbf{x} - b_g)$ for any $\mathbf{w}'_g \propto P\mathbf{w}_f$ is far from optimal.

the direction of \mathbf{w}_f . Therefore, $\text{Gain}(g') \leq \frac{\ell^2}{2r\sqrt{3}} \ll \text{Gain}(g)$ when unit cost $c \gg \frac{1}{r}$.

Using the characterization of an individual's best-response to g from Lemma 1, for any true quality function $f(\mathbf{x}) = \mathbf{w}_f \cdot \mathbf{x} - b_f$ and $g(\mathbf{x}) = \text{sign}(\mathbf{w}_g \cdot \mathbf{x} - b_g)$, we have

$$\text{Gain}(g) = (\mathbf{w}_g \cdot \mathbf{w}_f) \cdot \text{S-Den}^{1/c}(\mathbf{w}_g, b_g). \quad (3)$$

While mechanisms with $\mathbf{w}_g \propto P\mathbf{w}_f$ can be far from the optimal, one can still achieve a $1/(4rc)$ -approximation to the optimal mechanism by using such mechanisms and optimizing over the bias term b_g .

Theorem 2 ($1/(4rc)$ -approximation). *Consider the polynomial time algorithm that returns the best g from $\mathcal{G} = \{\text{sign}(\frac{P\mathbf{w}_f}{\|P\mathbf{w}_f\|} \cdot \mathbf{x} - b_g) | b_g = i/2c, \forall i \in [[2rc] + 1]\}$. Then, $\text{Gain}(g) \geq \frac{1}{4rc} \text{Gain}(g_{\text{opt}})$.*

This is a good approximation when the cost unit c is not too large compared to the radius of the domain r . However, in most cases the cost to change ones' feature is much larger than a constant fraction of the radius. In this case to find $g_{\text{opt}} \in \mathcal{G}_{0-1}$, we need to simultaneously optimize the total density of instances that fall within the margin of the classifier, their average distance to the decision boundary, and the correlation between \mathbf{w}_g and \mathbf{w}_f .

One of the challenges involved here is finding a halfspace whose margin captures a large fraction of instances, i.e., has large margin density. Many variants of this problem have been studied in the past and are known to be hard. For example, the densest subspace, the densest halfspace, densest cube, and the densest ball are all known to be hard to approximate [Ben-David *et al.*, 2002; Hardt and Moitra, 2013; Johnson and Preparata, 1978]. Yet, finding dense regions is a routine unsupervised learning task for which existing optimization tools are known to perform well in practice. Therefore, we assume that we have access to such an optimization tool, which we call a *density optimization oracle*.

Definition 1 (Density Optimization Oracle). *Oracle \mathcal{O} takes any distribution \mathcal{D} , margin ℓ , a set $\mathcal{K} \subseteq \mathbb{R}^n$, takes $O(1)$ time and returns*

$$\mathcal{O}(\mathcal{D}, \ell, \mathcal{K}) \in \underset{\substack{\mathbf{w} \in \mathcal{K}, \|\mathbf{w}\|=1 \\ b \in \mathbb{R}}}{\text{argmax}} \text{Den}_{\mathcal{D}}^{\ell}(\mathbf{w}, b). \quad (4)$$

Algorithm 1 ($1/4 - \epsilon$) Approximation for \mathcal{G}_{0-1}

Input: σ -smoothed distribution \mathcal{D} with radius r before perturbation, Vector \mathbf{w}_f , Projection matrix P , Desired accuracy ϵ .

Output: a linear threshold mechanism g

Let $\epsilon' = \min\{\epsilon^4, \epsilon^2 \sigma^4 / r^4\}$ and $\mathcal{C} = \emptyset$.

for $\eta = 0, \epsilon' \|P\mathbf{w}_f\|, 2\epsilon' \|P\mathbf{w}_f\|, \dots, \|P\mathbf{w}_f\|$ **do**

 Let $\mathcal{K}_{\eta} = \text{Img}(P) \cap \{\mathbf{w} \mid \mathbf{w} \cdot \mathbf{w}_f = \eta\}$.

 Use the density optimization oracle to compute

$$(\mathbf{w}^{\eta}, b^{\eta}) \leftarrow \mathcal{O}(\mathcal{D}, \frac{1}{c}, \mathcal{K}_{\eta})$$

 Let $\mathcal{C} \leftarrow \mathcal{C} \cup \{(\mathbf{w}^{\eta}, b^{\eta}), (\mathbf{w}^{\eta}, b^{\eta} + \frac{1}{2c})\}$

end for

return $g(\mathbf{x}) = \text{sign}(\mathbf{w} \cdot \mathbf{x} - b)$, where

$$(\mathbf{w}, b) \leftarrow \underset{(\mathbf{w}, b) \in \mathcal{C}}{\text{argmax}} (\mathbf{w} \cdot \mathbf{w}_f) \text{S-Den}_{\mathcal{D}}^{1/c}(\mathbf{w}, b).$$

Another challenge is that $\text{Gain}(\cdot)$ is a non-smooth function. As a result, there are distributions for which small changes to (\mathbf{w}, b) , e.g., to improve $\mathbf{w} \cdot \mathbf{w}_f$, could result in a completely different gain. However, one of the properties of real-life distributions is that there is some amount of noise in the data, e.g., because measurements are not perfect. This is modeled by *smoothed distributions*, as described in Section 2. Smooth distributions provide an implicit regularization that smooths the expected loss function over the space of all solutions.

Using these two assumptions, i.e., access to a density optimization oracle and smoothness of the distribution, we show that there is a polynomial time algorithm that achieves a $1/4$ approximation to the gain of g_{opt} . At a high level, smoothness of \mathcal{D} allows us to limit our search to those $\mathbf{w}_g \in \mathbb{R}^n$ for which $\mathbf{w}_f \cdot \mathbf{w}_g = v$ for a small set of discrete values v . For each v , we use the density oracle to search over all \mathbf{w}_g such that $\mathbf{w}_f \cdot \mathbf{w}_g = v$ and return a candidate with high margin density. We show how a mechanism with high margin density will lead us to (a potentially different) mechanism with high soft-margin density and as a result a near optimal gain. This approach is illustrated in more detail in Algorithm 1.

Theorem 3 ((Main) $1/4$ -approximation). *Consider a distribution over \mathcal{X} and the corresponding σ -smoothed distribution \mathcal{D} for some $\sigma \in O(r)$. Then, for any small enough ϵ , Algorithm 1 runs in time $\text{poly}(d, 1/\epsilon)$, makes $O(\frac{1}{\epsilon^4} + \frac{r^2}{\epsilon^2 \sigma^2})$ oracle calls and returns $g \in \mathcal{G}_{0-1}$, such that*

$$\text{Gain}(g) \geq \frac{1}{4} \text{Gain}(g_{\text{opt}}) - \epsilon.$$

Proof Sketch of Theorem 3. Recall from Equation 3, $\text{Gain}(g_{\text{opt}})$ is the product of two values $(\mathbf{w}_{g_{\text{opt}}} \cdot \mathbf{w}_f)$ and $\text{S-Den}(\mathbf{w}_{g_{\text{opt}}}, b_{g_{\text{opt}}})$. The first lemma shows that we can do a grid search over the value of $(\mathbf{w}_{g_{\text{opt}}} \cdot \mathbf{w}_f)$. In other words, there is a predefined grid on the values of $\mathbf{w} \cdot \mathbf{w}_f$ for which there is a \mathbf{w} for which $\mathbf{w} \cdot \mathbf{w}_f \approx \mathbf{w}_{g_{\text{opt}}} \cdot \mathbf{w}_f$. This is demonstrated in Figure 2.




Figure 2: The $\mathbf{w}_g \cdot \mathbf{w}_f$ can be increased by any small amount ϵ to match a specific value, when \mathbf{w}_f is shifted by a $\leq \epsilon$ angle to vector \mathbf{w} . When a distribution is smooth, the margin density of mechanisms (\mathbf{w}, b) and (\mathbf{w}_g, b_g) are close when $\theta(\mathbf{w}, \mathbf{w}_g)$ and $|b - b_g|$ are small. The lightly shaded area shows that instances outside of the margin contribute to (soft-)margin density.

Lemma 2 (Discretization). *For any two unit vectors $\mathbf{w}_1, \mathbf{w}_2 \in \text{Img}(P)$, and any ϵ , such that $\mathbf{w}_1 \cdot \mathbf{w}_2 \leq 1 - 2\epsilon$, there is a unit vector $\mathbf{w} \in \text{Img}(P)$, such that $\mathbf{w} \cdot \mathbf{w}_2 = \mathbf{w} \cdot \mathbf{w}_1 + \epsilon$ and $\mathbf{w} \cdot \mathbf{w}_1 \geq 1 - \epsilon$.*

The second technical lemma shows that approximating \mathbf{w}_g by a close vector, \mathbf{w} , and b_g by a close scalar, b , only has a small effect on its soft margin density. That is, the soft margin density is *Lipschitz*. For this claim to hold, it is essential for the distribution to be smooth. The key property of a σ -smoothed distribution \mathcal{D} —corresponding to the original distribution \mathcal{P} —is that $\text{Den}_{\mathcal{D}}^{\ell}(\mathbf{w}_g, b_g)$ includes instances $\mathbf{x} \sim \mathcal{P}$ that are not in the margin of (\mathbf{w}_g, b_g) . As shown in Figure 2, these instances also contribute to the soft margin density of any other \mathbf{w} and b , as long as the distance of \mathbf{x} to halfplanes $\mathbf{w} \cdot \mathbf{x} = b$ and $\mathbf{w}_g \cdot \mathbf{x} = b_g$ are comparable. So, it is sufficient to show that the distance of any \mathbf{x} to two halfplanes with a small angle is approximately the same. Here, angle between two unit vectors is defined as $\theta(\mathbf{w}, \mathbf{w}') = \arccos(\mathbf{w} \cdot \mathbf{w}')$. Lemma 2 leverages this fact to prove that soft margin density is Lipschitz smooth.

Lemma 3 (Soft Margin Lipschitzness). *For any distribution over \mathcal{X} and its corresponding σ -smooth distribution \mathcal{D} , for $\epsilon \leq \frac{1}{3R^2}$, any $\nu < 1$, for $R = 2r + \sigma\sqrt{2\ln(2/\nu)}$, and any $\ell \leq O(R)$, if $\theta(\mathbf{w}_1, \mathbf{w}_2) \leq \epsilon\sigma^2$ and $|b_1 - b_2| \leq \epsilon\sigma$, we have*

$$\left| \text{S-Den}_{\mathcal{D}}^{\ell}(\mathbf{w}_1, b_1) - \text{S-Den}_{\mathcal{D}}^{\ell}(\mathbf{w}_2, b_2) \right| \leq O(\nu + \epsilon R^2).$$

The final step to bring these results together is to show that we can find a mechanism (between those on the grid) that has a large soft margin density. To do this, we show that the mechanism (with potentially small changes to it) that has the highest margin density has at least $1/4$ of the soft density of the optimal mechanism. The proof technique of this lemma involves analyzing whether more than half of instances within an ℓ -margin of a mechanism are at distance at most $\ell/2$ of the margin, in which case soft margin density of the mechanism is at least $1/4$ of its density. Otherwise, shifting the bias of the

mechanism by $\ell/2$ results in a mechanism whose soft margin is at least $1/4$ of the original margin density.

Lemma 4 (Margin Density Approximates Soft-Margin Density). *For any distribution \mathcal{D} over \mathbb{R}^n , a class of unit vectors $\mathcal{K} \subseteq \mathbb{R}^n$, margin $\ell > 0$, and a set of values N , let $(\mathbf{w}_{\eta}, b_{\eta}) = \mathcal{O}(\mathcal{D}, \ell, \mathcal{K} \cap \{\mathbf{w} \cdot \mathbf{w}_f = \eta\})$ for $\eta \in N$, and let $(\hat{\mathbf{w}}, \hat{b})$ be the mechanism with maximum $\eta \text{Den}_{\mathcal{D}}^{\ell}(\mathbf{w}_{\eta}, b_{\eta})$ among these options. Then,*

$$\max \left\{ \text{Gain}(\hat{\mathbf{w}}, \hat{b}), \text{Gain}(\hat{\mathbf{w}}, \hat{b} + \frac{\ell}{2}) \right\} \geq \frac{1}{4} \max_{\substack{\mathbf{w} \in \mathcal{K} \\ \mathbf{w} \cdot \mathbf{w}_f \in N \\ b \in \mathbb{R}}} \text{Gain}(\mathbf{w}, b).$$

Putting these lemmas together, one can show that Algorithm 1, finds a $\frac{1}{4}$ approximation. \square

5 Learning Optimal Mechanisms From Samples

Up to this point, we have assumed that distribution \mathcal{D} is known to the mechanism designer and all computations, such as measuring margin density and soft margin density, can be directly performed on \mathcal{D} . However, in most applications the underlying distribution over feature attributes is unknown or the designer only has access to a small set of historic data. In this section, we show that all computations that are performed in this paper can be done instead on a small set of samples that are drawn i.i.d. from distribution \mathcal{D} . Note that (since no mechanism is yet deployed at this point) these samples represent the initial non-strategic feature attributes.

We first note that the characterization of the optimal *linear* mechanism (Theorem 1) is independent of distribution \mathcal{D} , that is, even with no samples from \mathcal{D} we can still design the optimal linear mechanism. On the other hand, the optimal *linear threshold* mechanism (Theorem 3) heavily depends on the distribution of instances, e.g., through the margin and soft-margin density of \mathcal{D} . To address sample complexity of learning a linear threshold mechanism, we use the concept of *pseudo-dimension*. This is the analog of VC-dimension for real-valued functions.

Definition 2 (Pseudo-dimension). *Consider a set of real-valued functions \mathcal{F} on the instance space \mathcal{X} . A set of instances $x^{(1)}, \dots, x^{(n)} \in \mathcal{X}$ is shattered by \mathcal{F} , if there are values $v^{(1)}, \dots, v^{(n)}$ such for any subset $T \subseteq [n]$ there is a function $f_T \in \mathcal{F}$ for which $f(x^{(i)}) \geq v^{(i)}$ if and only if $i \in T$. The size of the largest set that is shattered by \mathcal{F} is the pseudo-dimension of \mathcal{F} .*

It is well-known that the pseudo-dimension of a function class closely (via upper and lower bounds) controls the number of samples that are needed for learning a good function. This is characterized by the uniform-convergence property as shown below.

Theorem 4 (Uniform Convergence [Pollard, 2011]). *Consider the class of functions \mathcal{F} such that $f : \mathcal{X} \rightarrow [0, H]$ with pseudo-dimension d . For any distribution \mathcal{D} and a set S of $m = \epsilon^{-2} H^2 (d + \ln(1/\epsilon))$ i.i.d. randomly drawn samples*




Figure 3: $\text{Gain}_{\mathbf{x}}(g)$ is a minimum of $h_g^1(\mathbf{x})$ and $h_g^2(\mathbf{x})$.

from \mathcal{D} , with probability $1 - \delta$, for all $f \in \mathcal{F}$

$$\left| \frac{1}{m} \sum_{x \in S} f(x) - \mathbb{E}_{x \sim \mathcal{D}} [f(x)] \right| \leq \epsilon.$$

As is commonly used in machine learning, uniform convergence implies that choosing any mechanism based on its performance on the sample set leads to a mechanism whose expected performance is within 2ϵ of the optimal mechanism for the underlying distribution.

Lemma 5 (Pseudo-dimension). *For each $g = \text{sign}(\mathbf{w}_g \cdot \mathbf{x} - b_g)$ and \mathbf{x} let $\text{Gain}_{\mathbf{x}}(g)$ be the gain of g just on instance \mathbf{x} and let $\text{Den}_{\mathbf{x}}^{\ell}(g) = \mathbf{1}(\mathbf{w}_g \cdot \mathbf{x} - b_g \in [-\ell, 0])$. The class of real-valued functions $\{\text{Gain}_{\mathbf{x}} \circ g\}_{g \in \mathcal{G}_{0,1}}$ has a pseudo-dimension that is at most $O(\text{rank}(P))$. Moreover, the class of functions $\{\mathbf{x} \rightarrow \text{Den}_{\mathbf{x}}^{\ell}(g)\}_{g \in \mathcal{G}_{0,1}}$ has a pseudo-dimension (equivalently VC dimension) $O(\text{rank}(P))$.*

Proof. For $g \in \mathcal{G}_{0,1}$, note that

$$\text{Gain}_{\mathbf{x}}(g) = \mathbf{1}\left(\mathbf{w}_g \cdot \mathbf{x} - b_g \in \left[-\frac{1}{c}, 0\right]\right) (b_g - \mathbf{w}_g \cdot \mathbf{x})(\mathbf{w}_f \cdot \mathbf{w}_g).$$

Note that we can write $\text{Gain}_{\mathbf{x}}(g) = \min\{h_1(\mathbf{x}), h_2(\mathbf{x})\}$ for $h_g^1(\mathbf{x}) := \mathbf{1}(\mathbf{w} \cdot \mathbf{x} - b \leq 0)(b - \mathbf{w} \cdot \mathbf{x})(\mathbf{w} \cdot \mathbf{w}_f)$ and $h_g^2(\mathbf{x}) := \mathbf{1}(\mathbf{w} \cdot \mathbf{x} - b \geq -\ell)(\mathbf{w} \cdot \mathbf{w}_f)$. As show in Figure 3, h_g^1 and h_g^2 are both monotone functions of $\mathbf{w} \cdot \mathbf{x} - b$.

Pollard [Pollard, 2011] shows that the set of all linear functions in a rank k subspace has pseudo-dimension k . Pollard [Pollard, 2011] also shows that the set of monotone transformations of these functions has pseudo-dimension $O(k)$. Therefore, the set of functions $\{h_g^1(\mathbf{x})\}_{g \in \mathcal{G}_{0,1}}$ and $\{h_g^2(\mathbf{x})\}_{g \in \mathcal{G}_{0,1}}$, each have pseudo-dimension $\text{Rank}(P)$. It is well-known that the set of all minimums of two functions from two classes each with pseudo-dimension d has a pseudo-dimension of $O(d)$. Therefore, $\{\mathbf{x} \rightarrow \text{Gain}_{\mathbf{x}}(g)\}_{g \in \mathcal{G}_{0,1}}$ has pseudo-dimension of at most $O(\text{rank}(P))$. A similar construction shows that $\{\mathbf{x} \rightarrow \text{Den}_{\mathbf{x}}^{\ell}(g)\}_{g \in \mathcal{G}_{0,1}}$ has VC dimension of at most $O(\text{rank}(P))$. \square

We give a bound on the number of samples sufficient for finding a near optimal mechanism in $\mathcal{G}_{0,1}$.

Theorem 5 (Sample Complexity). *For any small enough ϵ and δ , there is*

$$m \in O\left(\frac{1}{c^2 \epsilon^2} (\text{rank}(P) + \ln(1/\delta))\right)$$

such that for $S \sim \mathcal{D}^m$ i.i.d. samples with probability $1 - \delta$, $\hat{g} \in \mathcal{G}_{0,1}$ that optimizes the empirical gain on S has $\text{Gain}(\hat{g}) \geq \text{Gain}(g_{\text{opt}}) - \epsilon$. Furthermore, with probability $1 - \delta$ when Algorithm 1 is run on S , the outcome \hat{g} has $\text{Gain}(\hat{g}) \geq \frac{1}{4} \text{Gain}(g_{\text{opt}}) - \epsilon$.

6 Discussion

In this work, we focused on increasing the expected quality of the population, i.e., welfare, through classification. There are, however, other reasonable objectives that one may want to consider. In some cases, our goal may be to design evaluation mechanisms that lead to highest quality *accepted* individuals, not accounting for those who are rejected by the mechanism. It would be interesting to study these objectives under our model as well. An example of this is to consider, for a set of boolean valued functions \mathcal{G} , $\max_{g \in \mathcal{G}} \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [f(\delta_g(\mathbf{x})) \mid g(\mathbf{x}) = 1]$.

Throughout this work, we focused on the L_2 cost function (i.e., $\text{cost}(\mathbf{x}, \mathbf{x}') = c \|\mathbf{x} - \mathbf{x}'\|_2$). This specification models scenarios in which an individual can improve multiple features most efficiently by combining effort rather than working on each feature individually (L_1 norm). For example, simultaneously improving writing, vocabulary, and critical analysis of a student more by for example reading novels, is more effective than spending effort to improve vocabulary by, say, memorizing vocab lists, and then improving the other attributes. It would be interesting to analyze the algorithmic aspects of alternative cost functions, such as L_1 cost or even non-metric costs (e.g., ones with learning curves whereby the first 10% improvement is cheaper than the next 10%), and different costs for different types of individuals.

Finally, we have assumed we know the true mapping of features to qualities (i.e., f). In many settings, one might not know this mapping, or even the full set of features. Instead, the designer only observes the quality of individuals after they respond to their incentives (i.e., $(f(\delta_g(\mathbf{x})))$), and the projection of their new feature set (i.e., $P\delta_g(\mathbf{x})$). Existing works on Stackelberg games and strategic classification have considered the use of learning tools for designing optimal mechanisms without the knowledge of the underlying function [Blum *et al.*, 2014; Haghtalab *et al.*, 2016; Miller *et al.*, 2020]. It would be interesting to characterize how the nature of observations and interventions available in our setting specifically affects this learning process.

References

- [Alon *et al.*, 2020] Tal Alon, Magdalen Dobson, Ariel D Procaccia, Inbal Talgam-Cohen, and Jamie Tucker-Foltz. Multiagent evaluation mechanisms. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 2020.
- [Ben-David *et al.*, 2002] Shai Ben-David, Nadav Eiron, and Hans Ulrich Simon. The computational complexity of densest region detection. *Journal of Computer and System Sciences*, 64(1):22–47, 2002.
- [Blum *et al.*, 2014] Avrim Blum, Nika Haghtalab, and Ariel D. Procaccia. Learning optimal commitment to

- overcome insecurity. In *Proceedings of the 28th Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 1826–1834, 2014.
- [Cai *et al.*, 2015] Yang Cai, Constantinos Daskalakis, and Christos Papadimitriou. Optimum statistical estimation with strategic data sources. In *Proceedings of the 28th Conference on Computational Learning Theory (COLT)*, pages 280–296, 2015.
- [Carroll, 2015] Gabriel Carroll. Robustness and linear contracts. *American Economic Review*, 105(2):536–563, 2015.
- [Chen *et al.*, 2018] Yiling Chen, Chara Podimata, Ariel D Procaccia, and Nisarg Shah. Strategyproof linear regression in high dimensions. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, pages 9–26. ACM, 2018.
- [Dekel *et al.*, 2010] Ofer Dekel, Felix Fischer, and Ariel D. Procaccia. Incentive compatible regression learning. *Journal of Computer and System Sciences*, 76(8):759–777, 2010.
- [Dong *et al.*, 2018] Jinshuo Dong, Aaron Roth, Zachary Schutzman, Bo Waggoner, and Zhiwei Steven Wu. Strategic classification from revealed preferences. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, pages 55–70, 2018.
- [Dütting *et al.*, 2019] Paul Dütting, Tim Roughgarden, and Inbal Talgam-Cohen. Simple versus optimal contracts. In *Proceedings of the 20th ACM Conference on Economics and Computation (EC)*, pages 369–387, 2019.
- [Grossman and Hart, 1983] Sanford J Grossman and Oliver D Hart. An analysis of the principal-agent problem. *Econometrica*, 51(1):7–45, 1983.
- [Haghtalab *et al.*, 2016] Nika Haghtalab, Fei Fang, Thanh Hong Nguyen, Arunesh Sinha, Ariel D. Procaccia, and Milind Tambe. Three strategies to success: Learning adversary models in security games. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 308–314, 2016.
- [Hardt and Moitra, 2013] Moritz Hardt and Ankur Moitra. Algorithms and hardness for robust subspace recovery. In *Conference on Computational Learning Theory (COLT)*, pages 354–375, 2013.
- [Hardt *et al.*, 2016] Moritz Hardt, Nimrod Megiddo, Christos Papadimitriou, and Mary Wootters. Strategic classification. In *Proceedings of the 7th ACM Conference on Innovations in Theoretical Computer Science Conference (ITCS)*, pages 111–122, 2016.
- [Holmstrom and Milgrom, 1987] Bengt Holmstrom and Paul Milgrom. Aggregation and linearity in the provision of intertemporal incentives. *Econometrica*, 55(2):303–328, 1987.
- [Holmstrom and Milgrom, 1991] Bengt Holmstrom and Paul Milgrom. Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *Journal of Law, Economics, and Organization*, 7:24–52, 1991.
- [Hu *et al.*, 2019] Lily Hu, Nicole Immorlica, and Jennifer Wortman Vaughan. The disparate effects of strategic manipulation. In *Proceedings of the 2nd ACM Conference on Fairness, Accountability, and Transparency (FAT*)*, pages 259–268, 2019.
- [Johnson and Preparata, 1978] David S. Johnson and Franco P Preparata. *Theoretical Computer Science*, 6(1):93–107, 1978.
- [Kleinberg and Raghavan, 2019] Jon Kleinberg and Manish Raghavan. How do classifiers induce agents to invest effort strategically? In *Proceedings of the 20th ACM Conference on Economics and Computation (EC)*, pages 825–844, 2019.
- [Mcafee and McMillan, 1986] R Preston McAfee and John McMillan. Bidding for contracts: A principal-agent analysis. *The RAND Journal of Economics*, 17(3):326–338, 1986.
- [Meir *et al.*, 2012] Reshef Meir, Ariel D Procaccia, and Jeffrey S Rosenschein. Algorithms for strategyproof classification. *Artificial Intelligence*, 186:123–156, 2012.
- [Miller *et al.*, 2020] John Miller, Smitha Milli, and Moritz Hardt. Strategic classification is causal modeling in disguise. *arXiv preprint arXiv:1910.10362*, 2020.
- [Milli *et al.*, 2019] Smitha Milli, John Miller, Anca D Dragan, and Moritz Hardt. The social cost of strategic classification. In *Proceedings of the 2nd ACM Conference on Fairness, Accountability, and Transparency (FAT*)*, pages 230–239, 2019.
- [Pollard, 2011] D. Pollard. *Convergence of Stochastic Processes*. Springer Series in Statistics. 2011.
- [Ross, 1973] Stephen A Ross. The economic theory of agency: The principal’s problem. *American Economic Review*, 63(2):134–139, 1973.