
PAC Learning with Improvements

Idan Attias^{1,2} Avrim Blum² Keziah Naggita² Donya Saless² Dravyansh Sharma^{2,3} Matthew Walter²

Abstract

One of the most basic lower bounds in machine learning is that in nearly any nontrivial setting, it takes *at least* $1/\epsilon$ samples to learn to error ϵ (and more, if the classifier being learned is complex). However, suppose that data points are agents who have the ability to improve by a small amount if doing so will allow them to receive a (desired) positive classification. In that case, we may actually be able to achieve *zero* error by just being “close enough”. For example, imagine a hiring test used to measure an agent’s skill at some job such that for some threshold θ , agents who score above θ will be successful and those who score below θ will not (i.e., learning a threshold on the line). Suppose also that by putting in effort, agents can improve their skill level by some small amount r . In that case, if we learn an approximation $\hat{\theta}$ of θ such that $\theta \leq \hat{\theta} \leq \theta + r$ and use it for hiring, we can actually achieve error zero, in the sense that (a) any agent classified as positive is truly qualified, and (b) any agent who truly is qualified can be classified as positive by putting in effort. Thus, the ability for agents to improve has the potential to allow for a goal one could not hope to achieve in standard models, namely zero error. In this paper, we explore this phenomenon more broadly, giving general results and examining under what conditions the ability of agents to improve can allow for a reduction in the sample complexity of learning, or alternatively, can make learning harder. We also examine both theoretically and empirically what kinds of improvement-aware algorithms can take into account agents who have the ability to improve to a limited extent when it is in their interest to do so.

Authors are listed alphabetically. ¹University of Illinois at Chicago ²Toyota Technological Institute at Chicago ³Northwestern University. Correspondence to: Keziah Naggita <knaggita@ttic.edu>, Donya Saless <donya@ttic.edu>.

1. Introduction

There has been growing interest in recent years in machine learning settings where a deployed classifier will influence the behavior of the entities it is aiming to classify. For example, a classifier that maps loan applicants to credit scores and then uses a particular cutoff $\hat{\theta}$ to determine whether an applicant should receive a loan will induce those below the cutoff value to take actions to improve their score. This setting is called strategic classification (Hardt et al., 2016) or measure management (Bloomfield, 2016) when the actions taken do not truly improve the agent’s quality, and performative prediction (Perdomo et al., 2020) more generally. In this work, our focus is on the case that the improvements are real e.g., paying off high-interest credit card debt, taking a money management class, etc. for genuinely improving one’s loan application. That is, the agent responds to the classifier in order to potentially improve their classification (Kleinberg & Raghavan, 2020; Miller et al., 2020), changing its true features in the process. The classifier must take this “strategic improvement” response into account.

Unlike previous works on strategic improvement that focus extensively on efficiently incentivizing and maximizing agent improvement (e.g., Miller et al. 2020; Kleinberg & Raghavan 2020; Haghtalab et al. 2020; Shavit et al. 2020, among others), we aim to understand how an agent’s capacity for improvement impacts learnability, sample complexity, and algorithm design for accurate classification. One high-level take-away from our theoretical analysis and empirical results is that the ability of agents to improve favors algorithms that are more “conservative” in their decisions. This is both due to the reduced concern over false-negative errors (since agents in those regions may still be able to improve to be classified as positive) and the increased concern over false-positive errors (which may cause individuals to “improve” incorrectly).

To illustrate the potential reduction in sample complexity that result from agents’ ability to improve, one of the most basic lower bounds in machine learning is that in nearly any nontrivial setting, it takes *at least* $1/\epsilon$ samples to learn to error ϵ (and more, if the classifier being learned is complex). However, if agents have the ability to improve by a small amount, we may actually be able to achieve *zero* error by just being “close enough”. To the best of our knowledge, this has

not been previously observed in the strategic improvement literature. Returning to the loan example above, suppose that by putting in effort, agents can improve their credit score by some small amount r , and suppose we are in the realizable case that there is some true threshold θ such that agents with credit score above θ will be good customers and those who score below θ will not. In that case, if we learn an approximation $\hat{\theta}$ of θ such that $\theta \leq \hat{\theta} \leq \theta + r$ and use it as a cutoff to determine who should receive a loan, we can actually achieve zero error in that (a) any agent classified as positive is truly qualified, and (b) any agent who truly is qualified can get classified as positive by putting in effort. Thus, the ability for agents to improve can potentially allow for a goal one could not hope to achieve otherwise.

We also observe fundamental differences in the inherent learnability of concept classes, compared to both standard PAC learning where the agents cannot respond to the classifier, as well as the strategic PAC setting where the agent tries to deceive the classifier to obtain a more favorable classification. Somewhat surprisingly, learning with improvements can sometimes be easier than the standard PAC setting, and it can sometimes be harder than strategic classification. We show that proper learnability with improvements in the realizable setting is closely linked to the concept class being intersection-closed.

Concretely, our contributions are as follows:

- In Section 3, we show a separation between the standard PAC model and our model of PAC learning with improvements. Specifically, we show that a finite VC dimension is neither necessary nor sufficient for PAC learnability with improvements. We further show a similar separation from the more recently studied PAC model for strategic classification (Hardt et al., 2016; Sundaram et al., 2023).
- In Section 4, we study learnability of geometric concepts in \mathbb{R}^d . We show that any intersection-closed concept class is learnable under our model, and show that the generalization error can be smaller than the standard PAC setting for interesting cases including thresholds and high-dimensional rectangles. We also show that the intersection-closed property is essentially necessary for proper learnability in our setting.
- In Section 5, we study a graph model in which each node represents an agent and the improvement set of an agent is the set of its neighbors in the graph. We establish near-tight bounds on the number of labeled points the learner needs to see to learn a classifier which achieves zero error with high probability, given the ability to learn the labels of uniformly random nodes. We further show that it is possible to learn a “fairer” classifier that also enables improvement whenever it leads to a better classification for an agent. We also study a *teaching* setting where the teacher aims to find the smallest set of labels needed to

ensure that a risk-averse student achieves zero-error, and show that providing the labels for the dominating set of the positive subgraph (induced by the true positive nodes) is sufficient.

- In Section 6, we conduct experiments on three real-world and one fully synthetic binary classification tabular datasets to investigate how the error rate of a model function (h) decreases when test-set agents that it initially classified as negative improve. Our results indicate that while risk-averse models may start with higher error rates, their errors rapidly drop as the negatively classified test agents improve and the improvement budget (r) increases. A stricter penalty for false positives typically leads to more accurate positive classifications, resulting in greater gains from agent improvements. In most cases, test errors decline sharply, sometimes reaching zero (e.g., in Figure 10g).¹

Related Work. Learning in the presence of strategic (“gaming”), utility-maximizing agents has gained increasing attention in recent years (Hu et al. 2019; Milli et al. 2019; Braverman & Garg 2020; Ahmadi et al. 2021; Haghtalab et al. 2020, among others). Early research framed this problem as a Stackelberg competition (Hardt et al., 2016; Brückner & Scheffer, 2011), where negatively classified agents manipulate their features to obtain more favorable outcomes if the benefits outweigh the costs. Kleinberg & Raghavan (2020) extend this model by considering agents who can both manipulate and genuinely improve their features, proposing a mechanism that incentivizes authentic improvement. This model has been studied under a causal lens, where the learner may not a priori know which features correspond to manipulation or improvement. Strategic learning from observable data requires solving a causal inference problem in this setting (Miller et al., 2020), and the ability to test different decision rules can be helpful (Shavit et al., 2020). Ahmadi et al. (2022) consider a similar setting and propose classification models that balance maximizing true positives with minimizing false positives.

We extend this line of work by analyzing the inherent learnability of classes, the sample complexity of learning, and the ability to achieve zero-error classification, when agents can truly improve. Inherent learnability of concepts has been studied in the strategic manipulation setting (Sundaram et al., 2023; Cohen et al., 2024; Lechner & Uner, 2022), but not in the strategic improvement setting. In Section B.2, we show how our improvement setting differs from strategic manipulation with respect to learnability. The sample complexity of learning in the presence of purely improving agents has been studied by Haghtalab et al. (2020), but from a social welfare perspective where the goal is to maximize the true positives after improvement. In contrast, our primary focus is classi-

¹Our code is publicly available [here](#).

fication error, which is more sensitive to false positives. In Section 5.2, we show that these two objectives need not be in conflict and may be simultaneously optimized. Finally, the ability of a learner to achieve zero-error for non-trivial concept classes and distributions has not been previously observed in any strategic or non-strategic setting.

Our work also relates to research in reliable machine learning (Rivest & Sloan, 1988; El-Yaniv & Wiener, 2010), where a learner may abstain from classification to avoid mistakes, balancing coverage (the proportion of classified points) against error. In contrast, we strive for a zero false positive rate and minimal false negative rate, aligning with learning under one-sided error (Natarajan, 1987; Kalai et al., 2012). We include a more detailed discussion of the related work in Appendix A.

2. Formal Setting: PAC Learning with Improvements

Let \mathcal{X} denote the instance space consisting of agents with the ability to *improve*, as defined below. We restrict our attention to the case of binary classification, i.e., the label space is $\{0, 1\}$. Without loss of generality, we refer to label 0 as the *negative* class and label 1 as the *positive* class. Let $\Delta : \mathcal{X} \rightarrow 2^{\mathcal{X}}$ denote the *improvement function* that maps each point (agent) $x \in \mathcal{X}$ to a set of points $\Delta(x)$ (the *improvement set*) to which x can potentially *improve* itself in order to be classified positively. For example, if \mathcal{X} is a metric space, we can define $\Delta(x)$ as the ℓ_p -ball centered at x . Let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ denote the concept space. We will focus on the *realizable* setting, i.e. we assume the existence of an unknown (to the learner) target concept $f^* : \mathcal{X} \rightarrow \{0, 1\}$ that correctly labels all points in \mathcal{X} and satisfies $f^* \in \mathcal{H}$.

The intuition behind the model is as follows. The learner first publishes a classifier $h : \mathcal{X} \rightarrow \{0, 1\}$ (potentially based on some data sample labeled according to f^*). We will refer to this function h as the learner’s hypothesis. Each agent then reacts to h (Zrnic et al., 2021; Hardt et al., 2016)—if it was classified negatively by h , the agent attempts to find a point in its improvement set that is positively classified by h and moves to it. Note that the agents do not know the true function f^* and as a result cannot react with respect to the ground truth, only based on h .

We formalize this as the *reaction set* with respect to h ,

$$\Delta_h(x) = \begin{cases} \{x\} & \text{if } h(x) = 1, \\ \{x\} & \text{if } \{x' \in \Delta(x) \mid h(x') = 1\} = \emptyset, \\ \{x' \in \Delta(x) : h(x') = 1\} & \text{otherwise.} \end{cases}$$

In other words, if h classifies x as positive, the agent x stays in place and does not attempt to improve. If h classifies x as negative, there are two types of reactions. Either, there is no point in its improvement set that can improve the agent’s classification according to h and the agent again stays put.

Otherwise, the agent reacts and moves to be predicted positive by h . This corresponds to utility-maximizing agents that have a utility of 1 for being classified as positive, a utility of 0 for being classified as negative, and that incur a cost for moving, where $\Delta(x)$ corresponds to the points that x can move to at a cost less than 1.

We say that a test point x has been misclassified if there exists a point in the reaction set of x where h disagrees with f^* , formally,

$$\text{LOSS}(x; h, f^*) = \max_{x' \in \Delta_h(x)} \mathbb{I}[h(x') \neq f^*(x')]. \quad (1)$$

Remark 2.1. The formulation of the loss function in (1) allows for scenarios where an input x initially satisfies $f^*(x) = 1$ and $h(x) = 0$, but under $\Delta_h(x)$, it may transition to a setting where $f^*(x') = 0$ and $h(x') = 1$ for some $x' \in \Delta_h(x)$. Think of an example where there are two features such that improving one often comes at the expense of the other. For instance, consider the trade-off between strength and endurance in athletics. Let $f^*(x)$ represent a person’s endurance (e.g., marathon running capability), and $h(x)$ represent their strength (e.g., sprinting power). Focusing on increasing $h(x)$ through strength training enhances power, but this often comes at the expense of endurance, thus reducing $f^*(x)$. This reflects the natural conflict between optimizing for one feature while sacrificing the other.

In words, this corresponds to an assumption that agents will improve to a point in their reaction set while breaking ties adversarially, or equivalently, that they will break ties in favor of points x' for which $f^*(x') = 0$. This assumption is natural if we want our positive results to be robust to unknown tie-breaking mechanisms, and would also hold if improving to points $x' \in \Delta_h(x)$ whose true label according to f^* is negative is less effort than improving such points whose true label is positive. Note that this loss function favors classifiers that label uncertain points as negative rather than positive. For example, if $\{x \mid h(x) = 1\} \subseteq \{x \mid f^*(x) = 1\}$ then h may still have zero loss if all points x in the difference have at least one point $x' \in \Delta(x)$ for which $h(x') = 1$. The fact that true positives might need to put in effort to improve in order to be classified as positive (or that some negative points are not able to improve themselves to be classified as positive by h even if they would have been able to do so with respect to f^*) does not count as an error in our setting.

See Section 4.1 for a concrete example.

Analogous to standard PAC learning, we assume the learner has access to a finite set of samples $S \in \mathcal{X}^m$ drawn randomly according to some fixed but unknown distribution \mathcal{D} over \mathcal{X} , and labeled by f^* . The learner’s population loss is given by $\text{LOSS}_{\mathcal{D}}(h, f^*) = \mathbb{P}_{x \sim \mathcal{D}} [\text{LOSS}(x; h, f^*)]$. This is formalized in the following.

Definition 2.2 (PAC Learning with improvements). Algorithm \mathcal{A} PAC-learns with improvements a concept class \mathcal{H} with respect to improvement function Δ and data distribution \mathcal{D} using sample size $M := M(\epsilon, \delta, \Delta, \mathcal{H}, \mathcal{D})^2$, if for any $f^* \in \mathcal{H}$, any $\epsilon > 0$ and $\delta > 0$, the following holds. Algorithm \mathcal{A} , with access to a sample $S \stackrel{\text{i.i.d.}}{\sim} \mathcal{D}^M$ labeled according to f^* , produces with probability at least $1 - \delta$ a hypothesis $h : \mathcal{X} \rightarrow \{0, 1\}$ with $\text{LOSS}_{\mathcal{D}}(h, f^*) \leq \epsilon$. We further say that \mathcal{A} learns \mathcal{H} w.r.t. Δ and \mathcal{D} with zero-error with sample size M if for any $\delta > 0$, given $S \stackrel{\text{i.i.d.}}{\sim} \mathcal{D}^M$ labeled by f^* , it returns h with $\text{LOSS}_{\mathcal{D}}(h, f^*) = 0$ with probability at least $1 - \delta$. We will also consider distribution-independent learning, where the guarantee should hold for all distributions \mathcal{D} and proper learning where we require $h \in \mathcal{H}$.

Note that in our learning with improvements setting zero-error can be achieved by learning (with high probability) from a finite sample in several interesting cases, which is impossible to achieve in the standard PAC model.

3. Separating PAC Learning with Improvements and the Standard PAC Model

In this section, we prove that learning with improvements diverges from the traditional behavior of the standard PAC model for binary classification. In the PAC model, the learnability of a concept class is equivalent to the class having a finite VC dimension. However, in our setting, where agents can improve, this condition is neither necessary nor sufficient for learnability. Concretely, we demonstrate that a class with an infinite VC dimension can still be learnable with improvements. We also provide examples of concept classes with finite VC dimensions and corresponding improvement sets that cannot be learned in our framework.

Theorem 3.1. *Finite VC dimension is neither necessary nor sufficient for PAC learnability with improvements.*

Proof. The proof is in Examples 1 and 2 below. \square

Example 1 (Finite VC dimension is not necessary for learnability with improvements). Consider any class \mathcal{H} of infinite VC-dimension, and define $\Delta(x) = \mathcal{X}$ for all examples $x \in \mathcal{X}$. We can learn this class \mathcal{H} with respect to this improvement function Δ with sample complexity $M(\epsilon, \delta) = \frac{1}{\epsilon} \ln(\frac{1}{\delta})$ for any data distribution \mathcal{D} as follows. First, draw a sample S of size $M(\epsilon, \delta)$. Next, if all examples in S are negative, then output the “all-negative” classifier; otherwise, select any positive example $x^* \in S$ and output the classifier $h(x) = \mathbb{I}[x = x^*]$. Note that in the latter case, the hypothesis h has error zero, because all agents will improve to x^* . Therefore, if $\mathbb{P}_{x \sim \mathcal{D}}[f^*(x) = 1] > \epsilon$, then h will have zero error with probability at least $1 - \delta$, whereas

²We say the sample complexity of \mathcal{A} is the smallest such M .

if $\mathbb{P}_{x \sim \mathcal{D}}[f^*(x) = 1] \leq \epsilon$, then h will have error at most ϵ with probability 1.

Example 2 (Finite VC dimension is not sufficient for learnability with improvements). Let the instance space \mathcal{X} be $[0, 1]$, let $\mathcal{H} = \{h_{abcd} : h_{abcd}(x) = 1 \text{ iff } x \in [a, b) \cup (c, d]\}$ (i.e., \mathcal{H} is the class of unions of two intervals, where to make the example easier, we define the intervals to be half-open), and let \mathcal{D} be the uniform distribution over $[0, 1]$. We define Δ as follows. For $x \in [0, 1/4) \cup (3/4, 1]$ let $\Delta(x) = [0, 1]$; for $x \in [1/4, 3/4]$, let $\Delta(x) = \{x\}$.

We claim that no algorithm with finite training data can guarantee an expected error of less than $1/4$, even though the class is easily PAC-learnable without improvements.

Consider a target function defined as the union of two intervals $[1/4, b) \cup (b, 3/4]$ where the number b was randomly chosen in $[1/4, 3/4]$. With probability 1, the learner will not see the point b in its training data, so it learns nothing from its training data about the location of b . Finally, if the learner outputs a classifier whose positive region has probability mass $\leq 1/4$, then its error rate is at least $1/4$ because the positive examples cannot move so at least half of their probability mass will get misclassified. On the other hand, if the learner outputs a classifier whose positive region has probability mass greater than $1/4$, then it has at least a 50% chance of including a negative point in its positive region (it will surely include a negative point if it is not contained in $[1/4, 3/4]$ and has at least a 50% chance of doing so otherwise, since b was uniformly chosen from $[1/4, 3/4]$). If the classifier has a negative point in its positive region, then it will have an error rate at least 50%, because all the negatives in $[0, 1/4)$ and $(3/4, 1]$ will move to a false positive (here we use that agents break ties adversarially). So, either way, its expected error is 25%.

Union of two intervals is arguably the simplest class that is not intersection-closed (Definition 4.3). Indeed, we show in Section 4.2 that such an example could not be possible for intersection-closed classes. As another example of the separation of our model from the standard PAC setting, we show it is possible that for some target function f^* the learner using a hypothesis space $\tilde{\mathcal{H}}$ achieves no error in the standard PAC setting, but it is impossible to avoid a constant error rate when learning the same target with improvements using the same hypothesis space $\tilde{\mathcal{H}}$.

Theorem 3.2. *Consider a hypothesis class $\tilde{\mathcal{H}}$ and target function $f^* \notin \tilde{\mathcal{H}}$. Let $d(f^*, \tilde{\mathcal{H}})$ denote the error of the best hypothesis in $\tilde{\mathcal{H}}$ w.r.t. the target f^* . It is possible to have $d(f^*, \tilde{\mathcal{H}}) = 0$ in the standard PAC setting (there exists $h \in \tilde{\mathcal{H}}$ that achieves error 0) but $d(f^*, \tilde{\mathcal{H}}) = 1/2$ for PAC learning with improvements (i.e., all $h \in \tilde{\mathcal{H}}$ have error at least $1/2$).*

Proof. See Example 3 in Appendix B.1. \square

We further show a separation of our model from strategic classification in Appendix B.2. Remarkably, learning with improvements may be either easier or harder than strategic classification depending upon the setting.

4. PAC Learning of Geometric Concepts

In this section, we first demonstrate the gain of the learner when agents can improve for the natural class of thresholds on the real line, where agents can move by a distance of at most r . We then study intersection-closed classes. In particular, we derive sample complexity bounds for the class of axis-aligned hyperrectangles, where the improvement sets are the ℓ_∞ balls. We further establish negative results for proper learners in the absence of the intersection-closed property. Lastly, we study the class of homogeneous halfspaces under the uniform distribution over the unit ball, where agents can improve by adjusting their angle. Complete proofs for this section are located in Appendix C.

We will use $(a)_+$ to denote $\max\{a, 0\}$.

4.1. Warm-up: Zero Error for Learning Thresholds

Let $\mathcal{H} = \{h_t : t \in \mathbb{R}\}$ be the class of one-sided threshold functions, where $h_t(x) = \mathbb{I}\{x \geq t\}$. The improvement set of x is simply the closed ball centered at x with radius r , i.e., $\Delta(x) = \{x' \mid |x - x'| \leq r\}$. Suppose the data distribution \mathcal{D} is uniform over $[0, 1]$, and labels are generated according to a target threshold $h_{t^*} \in \mathcal{H}$ for some $t^* \in [0, 1]$. Let $S = \{(x_i, y_i)\}_{i=1}^m$ be the set of training samples, where $x_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{D}$ and $y_i = h_{t^*}(x_i)$.

There are several options for choosing a threshold that achieves zero empirical error on S , as shown by the shaded area in Figure 1. Due to the asymmetry of the loss function (Eqn. 1), we choose the rightmost threshold consistent with S . This is the most “conservative” option, as any x that improves up to this threshold is guaranteed to be positive with respect to the unknown ground-truth h_{t^*} . This is a property that would not necessarily hold for lower thresholds. We define this threshold with respect to S as follows,

$$t_{S^+} = \begin{cases} \min(S^+), & \text{if } S^+ \neq \emptyset, \\ 1, & \text{if } S^+ = \emptyset, \end{cases}$$

where $S^+ = \{x_i \in S : y_i = 1\}$ is the set of positive examples in S . The hypothesis h_{S^+} is defined as $h_{S^+} = \mathbb{I}\{x \geq t_{S^+}\}$. Notice that using classifier h_{S^+} will induce agents (at test time) $x \in [t_{S^+} - r, t_{S^+})$ to improve to be classified as positive by h_{S^+} , which will be a correct classification since $t_{S^+} \geq t^*$.

Theorem 4.1 (Thresholds, uniform distribution). *Let the improvement set Δ be the closed ball with radius r , $\Delta(x) = \{x' \mid |x - x'| \leq r\}$. Let \mathcal{D} be the uniform distribution on*




Figure 1. Learning thresholds with improvements.

$[0, 1]$. For any $\epsilon, \delta \in (0, 1/2)$, with probability $1 - \delta$,

$$\text{LOSS}_{\mathcal{D}}(h_{S^+}, h_{t^*}) \leq (\epsilon - r)_+,$$

with sample complexity $M = O\left(\frac{1}{\epsilon} \log \frac{1}{\delta}\right)$.

Proof. See Appendix C.1 □

Note that the population error is improved from ϵ (in the standard PAC setting) to $\epsilon - r$ for the same sample size, and we can achieve zero error as long as we set $\epsilon \leq r$.

Theorem C.1 in Appendix C.2 proves a similar result for arbitrary distribution \mathcal{D} , where instead of getting $\epsilon - r$ population error, the reduction in the error is replaced by the following distribution-dependent quantity

$$p(h_{S^+}; h_{t^*}, \mathcal{D}, r) = \mathbb{P}_{x \sim \mathcal{D}} [x \in [t_{S^+} - r, t_{S^+}]] . \quad (2)$$

Note that the class of thresholds is closed under intersection: $\bigcap_{i=1}^n h_{t_i} = h_{\max\{t_1, t_2, \dots, t_n\}}$. In the following, we extend the analysis to such concept classes, more generally.

4.2. Intersection Closed Classes

The learnability of intersection-closed concept classes in the standard PAC model has been extensively studied (Helmbold et al., 1990; Auer, 1997; Auer & Cesa-Bianchi, 1998; Auer & Ortner, 2007; Darnstädt, 2015). In this section, we study the learnability with improvements of these classes. We start with the following definitions.

Definition 4.2 (Closure operator of a set). For any set $S \subseteq \mathcal{X}$ and any concept class $\mathcal{H} \subseteq 2^{\mathcal{X}}$, the *closure of S with respect to \mathcal{H}* , denoted by $\text{CLOS}_{\mathcal{H}}(S) : 2^{\mathcal{X}} \rightarrow 2^{\mathcal{X}}$, is defined as the intersection of all functions in \mathcal{H} that contain S , that is, $\text{CLOS}_{\mathcal{H}}(S) = \bigcap_{h \in \mathcal{H}, S \subseteq h} h$. In words, the closure of S is the smallest concept in \mathcal{H} which contains S . If $\{h \in \mathcal{H} : S \subseteq h\} = \emptyset$, then $\text{CLOS}_{\mathcal{H}}(S) = \mathcal{X}$.

Definition 4.3 (Intersection-closed classes). A concept class $\mathcal{H} \subseteq 2^{\mathcal{X}}$ is *intersection-closed* if for all finite $S \subseteq \mathcal{X}$, $\text{CLOS}_{\mathcal{H}}(S) \in \mathcal{H}$. In words, the intersection of all functions in \mathcal{H} containing an arbitrary subset of the domain belongs to \mathcal{H} . For finite concept classes, an equivalent definition states that for any $h_1, h_2 \in \mathcal{H}$, the intersection $h_1 \cap h_2$ is in \mathcal{H} as well (Natarajan, 1987).

Many natural concept classes are intersection-closed, for example, axis-parallel d -dimensional hyperrectangles, intersections of halfspaces, k -CNF boolean functions, and subspaces of a linear space.

The *Closure algorithm* is a learning algorithm that generates a hypothesis by taking the closure of the positive examples in a given dataset, and negative examples do not influence the generated hypothesis. The hypothesis returned by this algorithm is always the smallest hypothesis containing all of the positive examples seen so far in the training set.

Definition 4.4 (Closure algorithm [Natarajan, 1987](#); [Helm-bold et al., 1990](#)). Let $f^* \in \mathcal{H}$ and let $S = \{(x_1, y_1 = f^*(x_1)), \dots, (x_m, y_m = f^*(x_m))\}$ be a set of labeled examples, where each $x_i \in \mathcal{X}$. The hypothesis h_S^c produced by the closure algorithm is defined as:

$$h_S^c(x) = \begin{cases} 1, & \text{if } x \in \text{CLOS}_{\mathcal{H}}(\{x_i \in S : y_i = 1\}), \\ 0, & \text{otherwise.} \end{cases}$$

Here, $\text{CLOS}_{\mathcal{H}}(\{x_i \in S : y_i = 1\})$ denotes the closure of the set of positive examples in S with respect to \mathcal{H} .

The closure algorithm learns intersection-closed classes with VC dimension d with an optimal sample complexity of $\Theta(\frac{1}{\epsilon}(d + \log \frac{1}{\delta}))$ ([Auer & Ortner, 2007](#); [Darnstädt, 2015](#)).

We apply the closure algorithm for learning with improvements. In order to quantify the improvement gain of the returned hypothesis, we define the *improvement region* of h as the set of points that can improve from a negative classification to a (correct) positive classification by h .

Definition 4.5 (Improvement Region). The improvement region of hypothesis $h \subseteq f^*$, w.r.t. f^* and Δ is

$$\text{IR}(h; f^*, \Delta) := \{x : h(x) = 0, \exists x' \in \Delta(x) : h(x') = f^*(x') = 1\}. \quad (3)$$

The gain from improvements is the probability mass of the improvement region under \mathcal{D} : $\mathbb{P}_{x \sim \mathcal{D}}[x \in \text{IR}(h; f^*, \Delta)]$.

Note that for the class of thresholds, the closure algorithm returns exactly the hypothesis h_{S^+} , and the probability mass of the improvement region is $p(h_{S^+}; h_{t^*}, \mathcal{D}, r)$ (cf. Eqn. 2).

Axis-Aligned Hyperrectangles in $[0, 1]^d$. An axis-aligned hyperrectangle classifier assigns a value of 1 to a point if and only if the point lies within a specific rectangle. Formally, let $a = (a_1, \dots, a_d), b = (b_1, \dots, b_d) \in [0, 1]^d$ where $a_i \leq b_i$ for $i \in \{1, \dots, d\} := [d]$. A hyperrectangle $R_{(a,b)} = \prod_{i \in [d]} [a_i, b_i]$ classifies a point $x = (x_1, \dots, x_d)$ as: $R_{(a,b)}(x) = \mathbb{I}\{x_i \in [a_i, b_i], \forall i \in [d]\}$.

In the following, we show that the closure algorithm learns with improvements the concept class $\mathcal{H}_{\text{rec}} = \{R_{(a,b)} : a, b \in [0, 1]^d\}$.

Theorem 4.6 (Axis-aligned Hyperrectangles). *Let the improvement set Δ be the closed ℓ_∞ ball with radius r , $\Delta(x) = \{x' \mid \|x - x'\|_\infty \leq r\}$. Let R_S^c be the rectangle returned by the closure algorithm given $S \stackrel{\text{i.i.d.}}{\sim} \mathcal{D}^m$, and R^* be the target rectangle. For any distribution \mathcal{D} , for any $\epsilon, \delta \in (0, 1/2)$, with probability $1 - \delta$,*

$$\text{LOSS}_{\mathcal{D}}(R_S^c, R^*) \leq (\epsilon - \mathbb{P}_{x \sim \mathcal{D}}[x \in \text{IR}(R_S^c; R^*, \Delta)])_+,$$

with sample complexity $M = O(\frac{1}{\epsilon}(d + \log \frac{1}{\delta}))$.

When \mathcal{D} is the uniform distribution on $[0, 1]^2$, we can get the following expression. Denote by l_1 and l_2 the width and height (respectively) of the rectangle R_S^c . Then,

$$\mathbb{P}_{x \sim \mathcal{D}}[x \in \text{IR}(R_S^c; R^*, \Delta)] = 2r(l_1 + l_2) + 4r^2.$$

Note that, as opposed to the simple case of thresholds, the improvement region for hyperrectangles depends on the geometry of the target concept.

Arbitrary Intersection-closed Classes. We will now show that any intersection-closed concept class with a finite VC dimension is PAC learnable with improvements w.r.t. any improvement function Δ .

Theorem 4.7. *Let \mathcal{H} be an intersection-closed concept class on instance space \mathcal{X} . There is a learner that PAC-learns with improvements \mathcal{H} with respect to any improvement function Δ and any data distribution \mathcal{D} given a sample of size $O(\frac{1}{\epsilon}(d_{\text{VC}}(\mathcal{H}) + \log \frac{1}{\delta}))$, where $d_{\text{VC}}(\mathcal{H})$ denotes the VC-dimension of \mathcal{H} .*

Proof Sketch. Note that the closure algorithm outputs a hypothesis h^c which exactly classifies the positive agreement region as positive. By known results for intersection-closed concept classes, we have that $O(\frac{1}{\epsilon}(d_{\text{VC}}(\mathcal{H}) + \log \frac{1}{\delta}))$ examples are enough to guarantee that the error of h^c is less than ϵ in the standard PAC learning setting. Note that any point that h^c classifies as positive is positive according to f^* . Now if the negative points move in response to h^c and get positively classified by h^c , they must become truly positive. Thus, for the stated sample size, the error can only decrease with improvements below ϵ . This is quantified by the probability mass under \mathcal{D} of the points in the improvement region, as defined in Eqn. 3. Intuitively, if \mathcal{H} is an intersection-closed class in \mathbb{R}^d and Δ is the ℓ_∞ ball of radius r , then the improvement region constitutes a strip of width r along the boundary of h^c (a $(d - 1)$ -dimensional hypersurface). \square

We can also establish the following negative result which indicates the hardness of proper learning in the absence of the intersection-closed property.

Theorem 4.8. *Let \mathcal{H} be any concept class on a finite instance space \mathcal{X} such that at least one point $x' \in \mathcal{X}$ is classified negative by all $h \in \mathcal{H}$ (i.e. $\{x \mid h(x) = 0 \text{ for all } h \in \mathcal{H}\}$).*

$\mathcal{H}\} \neq \emptyset$), and suppose $\mathcal{H} \upharpoonright_{\mathcal{X} \setminus \{x'\}}$ is not intersection-closed on $\mathcal{X} \setminus \{x'\}$. Then there exists a data distribution \mathcal{D} and an improvement function Δ such that no proper learner can PAC-learn with improvements \mathcal{H} w.r.t. Δ and \mathcal{D} .

Note that we have an additional requirement that all classifiers in the concept space agree on some negative point (intuitively, agents who should never achieve positive classification). Consider the following simple example where this condition does not hold, the concept class is not intersection-closed, and learnability is possible in our setting.

Suppose $\mathcal{X} = \{x_1, x_2\}$ and $\mathcal{H} = \{h_1, h_2\}$ with $h_1(x_1) = 1, h_1(x_2) = 0$ and $h_2(x) = 1 - h_1(x)$ for either $x \in \mathcal{X}$. Clearly $h_1 \cap h_2 \notin \mathcal{H}$ and \mathcal{H} is not intersection-closed, yet knowledge of a single label tells us the target concept.

4.3. Halfspaces on the Unit Ball

We now consider the problem of learning homogeneous halfspaces with respect to the uniform distribution on the unit ball (or any spherically-symmetric distribution), when agents have the ability to improve by an angle of r .

Theorem 4.9. *Consider the class of d -dimensional halfspaces passing through the origin, i.e., $\mathcal{H} = \{x \mapsto \text{sign}(w^T x) : w \in \mathbb{R}^d\}$. Suppose \mathcal{X} is the surface of the origin-centered unit sphere in \mathbb{R}^d for $d > 2$, and \mathcal{D} is the uniform distribution on \mathcal{X} . For each point $x \in \mathcal{X}$, define its neighborhood $\Delta(x) = \{x' \mid \arccos(\langle x, x' \rangle) \leq r\}$. For any $\delta \in (0, 1/2)$, and training sample $S \stackrel{\text{i.i.d.}}{\sim} \mathcal{D}^m$ of size $\tilde{O}\left(\frac{d + \log \frac{1}{\delta}}{r}\right)$, with probability $1 - \delta$, $\text{LOSS}_{\mathcal{D}}(\text{POS}_{\mathcal{H}}(S), f^*) = 0$, where $\text{POS}_{\mathcal{H}}(S)$ is the intersection of the positive regions of all $h \in \mathcal{H}$ consistent with the training set S .*

Remark 4.10. It is worth noting that while the aforementioned result relies on the classifier $\text{POS}(\mathcal{H}_S)$, which is a fairly complex function, a similar guarantee can be achieved using a linear classifier (though non-homogeneous, so it is still not “proper”). Specifically, by obtaining a sufficiently large sample, one can construct a homogeneous linear classifier whose angle with respect to the target is at most $\frac{\pi}{2}$. We can then shift that classifier by $r/2$ (so it is no longer homogeneous) to ensure its positive region is contained inside the positive region for f^* .

5. Zero-error Learning in the Graph Model

In this section, we will consider a general discrete model for studying classification of agents with the ability to improve. The agents are located on the nodes of an undirected graph, and the edges determine the improvement function, i.e. the agents can move to neighboring nodes in order to potentially improve their classification. Note that the graph nodes correspond to an arbitrary discrete instance space \mathcal{X} . Remarkably,

zero error may be attained even in this general setting. All proofs in this Section are deferred to Appendix D.

Formally, let $G = (V, E)$ denote an undirected graph. The vertex set $V = \{x_1, x_2, \dots, x_n\}$ represents a fixed collection of n points corresponding to a finite instance space \mathcal{X} . The edge set $E \subseteq V \times V$ captures the adjacency information relevant for defining the improvement function. More precisely, for a given vertex $x \in V$, the improvement set of x is given by its neighborhood in the graph, i.e. $\Delta(x) = \{x' \in V \mid (x, x') \in E\}$ ³. Let $f^* : V \rightarrow \{0, +1\}$ represent the target labeling (or partition) of the vertices in the graph G . Assume that the concept space \mathcal{H} is the set of all possible labelings of the graph, which is finite.

5.1. Near-tight Sample Complexity for Zero-error

Our first result is to show that we can obtain zero-error in the learning with improvements setting, when the data distribution \mathcal{D} is given by a uniform distribution over V , and obtain near-tight bounds on the sample complexity. Our learner in this case is the conservative classifier $h \in \mathcal{H}$ that classifies exactly the positive points seen in the sample as positive, and the remaining points as negative. Even though we allow f^* to be an arbitrary labeling in \mathcal{H} , we do not need to see all the labels to learn an h that achieves zero error w.r.t. f^* . Intuitively, this is because for any positively labeled node x it is sufficient to see the label of x or that of one of its neighbors since the agents can move to a neighbor predicted positive by h . We further show that no algorithm can achieve a better sample complexity, up to some logarithmic factors.

Theorem 5.1. *Let $G = (V, E)$ be an undirected graph with $n = |V|$ vertices, and let $f^* : V \rightarrow \{0, +1\}$ denote the ground truth labeling function. Let d_{\min}^+ denote the minimum degree of the vertices in G^+ , the induced subgraph of G on the vertices $x \in V$ with $f^*(x) = 1$. Assume that the data distribution \mathcal{D} is uniform on V . For any $\delta > 0$, and training sample $S \stackrel{\text{i.i.d.}}{\sim} \mathcal{D}^m$ of size $m = O\left(\frac{n(\log n + \log \frac{1}{\delta})}{d_{\min}^+ + 1}\right)$, there exists a learner that achieves zero generalization error, i.e. learns a hypothesis h such that $\text{LOSS}_{\mathcal{D}}(h, f^*) = 0$, with probability at least $1 - \delta$ over the draw of S . Moreover, there exists a graph G for which any learner that achieves zero generalization error must see at least $\Omega\left(\frac{n}{d_{\min}^+ + 1} \log \frac{n}{d_{\min}^+ + 1}\right)$ labeled points in the training sample, with high constant probability.*

We note that the value of d_{\min}^+ (and, therefore our bound on the sample complexity) is generally not known to the learner in advance, and it depends on the graph structure as well as the (unknown) target labeling f^* . It is an interesting open

³Our results readily extend to $\Delta(x) = \{x' \in V \mid d_G(x, x') \leq r\}$, where d_G denotes the shortest path metric on G , by applying our arguments to G^r , the r^{th} power of G (see appendix).

question to design a learner that can determine whether a sample of sufficient size has been collected to guarantee zero-error. For the special case of the complete graph, our sample complexity bound becomes $m = \tilde{O}\left(\frac{n}{n^+}\right)$, where n^+ is the number of nodes labeled positive by f^* .

5.2. Enabling Improvement Whenever It Helps

Note that our loss function $\text{LOSS}(x; h, f^*)$ penalizes the learner for mistakes w.r.t. the target f^* after the agents have potentially reacted to h . However, we say nothing about whether a negative point x that truly has the ability to improve and get positively classified (i.e. $f^*(x') = 1$ for some $x' \in \Delta(x)$) will also be able to do so under our published classifier h . We will now consider an alternative measure of the performance of h (conceptually captures recall in the improvement setting) which measures the probability mass of the points for which we fail to enable improvement even though it is possible under f^* . Formally, we define

$$\text{LOSS}_{\mathcal{D}}^{\text{E}}(h, f^*) = \mathbb{P}_{x \sim \mathcal{D}} [f^*(x) = 0 \wedge \mathbb{I}[\Delta_{f^*}(x) = \{x\}] = \mathbb{I}[\Delta_h(x) = \{x\}]] \quad (4)$$

That is, we wish to ensure that an agent x with $f^*(x) = 0$ and an option to improve to a truly positive point in its reaction set w.r.t. f^* , will also see some option to improve and get positively classified according to h . In Theorem D.4 in the Appendix, we obtain near-tight sample complexity bounds for learning a hypothesis h that simultaneously guarantees that $\text{LOSS}_{\mathcal{D}}^{\text{E}}(h, f^*) = 0$ and $\text{LOSS}_{\mathcal{D}}(h, f^*) = 0$ when the data distribution is uniform over V .

5.3. Teaching a Risk-averse Student

The theory of teaching (Goldman & Kearns, 1995) studies the size of the smallest set of labeled examples needed to guarantee that a unique function in the concept space is consistent with the set (for labels according to any target concept in the space). If the teacher (that knows f^*) provides this labeled set, then a student that can do consistent learning (find a concept consistent with training data) will learn the target concept f^* . In our learning with improvements over the graph setting, it is natural to consider a simple variant where the student outputs the most risk-averse concept that only labels positive points seen in the labeled set received from the teacher as positive. Here we will consider the question of the minimum number of labeled examples the teacher needs to provide to the risk-averse student to achieve zero-error in our setting.

Let G^+ denote the induced subgraph of G on $V^+ = \{x \in V \mid f^*(x) = +1\}$, the nodes labeled positive by the target concept f^* . We show that it is sufficient for the teacher to present the labels of a *dominating set* of G^+ (Definition

D.2) for the risk-averse student to learn a zero-error classifier h . This observation also motivates and helps establish our learning result (Theorem 5.1).

Theorem 5.2. *Let $G = (V, E)$ be an undirected graph, and f^* be the target labeling. Let G^+ denote the induced subgraph on the vertices $x \in V$ with $f^*(x) = 1$, and S^+ denote the dominating set of G^+ . Then $\text{LOSS}(x; h_{S^+}, f^*) = 0$ for any $x \in V$, where $h_{S^+}(x) = \mathbb{I}\{x \in S^+\}$.*

Proof. See Appendix D.4. \square

6. Evaluation

Below and in Appendix E, we present the setup and results of our evaluation of improvement-aware algorithms, focusing on practical risk-aversion strategies, specifically loss-based and threshold-based methods. These strategies align with our theory, which predicts optimal performance for risk-averse classifiers when agents improve under a limited budget r . We also investigate whether, and under what conditions, model error can be driven to zero when agents can improve within an ℓ_∞ ball of radius r . Achieving zero error is a notable property in the learning with improvements setting, even for broad concept classes (cf. Section 5, where the instance space is discrete and concepts may be arbitrary functions). Our code is available [here](#).

Datasets. We use three real-world tabular datasets: the Adult UCI dataset (Becker & Kohavi, 1996), the OULAD and Law School datasets (Le Quy et al., 2022a), and a synthetic 8-dimensional binary classification dataset with class separability 4 and minimal outliers, generated using Scikit-learn’s `make_classification` function (Pedregosa et al., 2011). In each case we train a zero-error model f^* on the entire dataset, which we treat as the true labeling function for our experiments.

Let $\mathcal{S}_T = \{(x, y) \mid x \in \mathbb{R}^d, y \in \{0, 1\}\}$ represent the dataset (e.g., Adult), where x is the feature vector and $y = f^*(x)$ is the label. For all experiments, we split \mathcal{S}_T into training $\mathcal{S}_{\text{train}}$ (70%) and testing $\mathcal{S}_{\text{test}}$ (30%) subsets. Further dataset details, including improvement features and class distributions, are provided in Appendix E.1.

Classifiers. For each full dataset \mathcal{S}_T , we trained a zero-error model f^* using decision trees. We trained the decision-maker model $h : \mathbb{R}^d \rightarrow \{0, 1\}$, taking the form of a two-layer neural network, on $\mathcal{S}_{\text{train}}$ with tuned hyperparameters. To assess the loss function’s impact on error drop rate when agents improve, we trained both a standard model with binary cross-entropy (\mathcal{L}_{BCE}) loss and a risk-averse model with weighted-BCE ($\mathcal{L}_{\text{wBCE}}$) loss.

$$\mathcal{L}_{\text{wBCE}} = -\frac{1}{n} \sum_{i=1}^n \left[w_{\text{FP}}(1 - y_i) \log(1 - \hat{y}_i) + w_{\text{FN}} y_i \log(\hat{y}_i) \right] \quad (5)$$

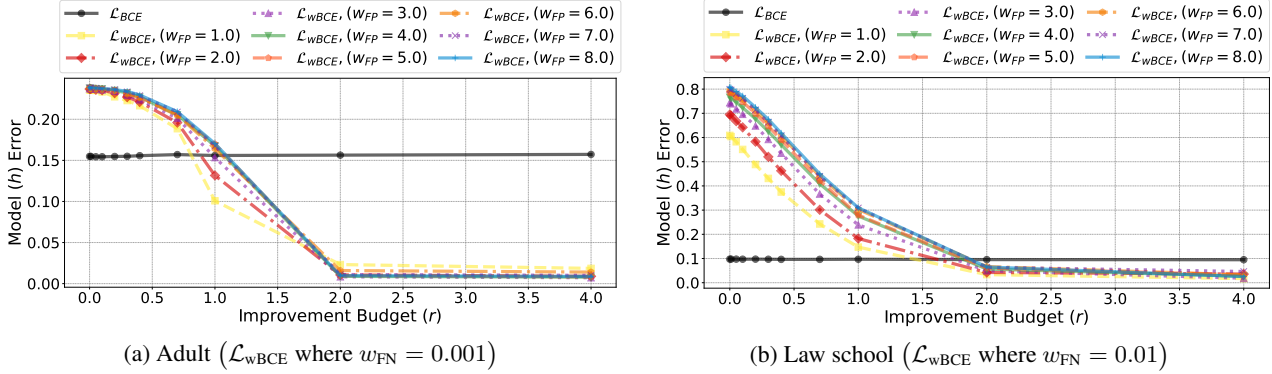


Figure 2. We compare the performance gains when agents improve to the risk-averse ($\mathcal{L}_{\text{wBCE}}, \frac{w_{\text{FP}}}{w_{\text{FN}}} > 1, w_{\text{FP}} = \{i\}_{i=1}^8$) and the standard ($\mathcal{L}_{\text{BCE}}, w_{\text{FP}} = w_{\text{FN}} = 1$) models across the Adult and Law School datasets using a fixed classification threshold of 0.5. Higher improvement budgets (r) and greater risk-aversion (high $\frac{w_{\text{FP}}}{w_{\text{FN}}}$) accelerate error reduction. See Figure 10 (Appendix) for full results.

where, $n = |\mathcal{S}_{\text{train}}|$, $y \in \{0, 1\}$ is the true label, $\hat{y} \in (0, 1]$ is the model prediction, and w_{FP} and w_{FN} are the false positive and false negative weights, respectively. Setting $w_{\text{FP}} = w_{\text{FN}} = 1$ in $\mathcal{L}_{\text{wBCE}}$ recovers \mathcal{L}_{BCE} .

Beyond the $\mathcal{L}_{\text{wBCE}}$ loss function, which penalizes false positives more heavily, we explore another form of risk-averse classification by applying a higher threshold of 0.9 (instead of the usual 0.5) to the sigmoid output of the final layer. For further details, refer to Appendix E.2.

Improvement. Given a trained model function h , a data sample $x \in \mathcal{S}_{\text{test}}$ with an undesirable model outcome $h(x) = 0$, and a subset of improvable features along with a predefined improvement budget r , we use Projected Gradient Descent (Madry et al., 2018) to compute the minimal change within the budget r required to transform x into a positive outcome $h(x') = 1$. Specifically, we aim to find:

$$x' = \text{Proj}_{\Delta(x)}(x_t) + \alpha \cdot \text{sign}(\nabla_{x_t} \mathcal{L}(h(x_t), h(x))) \quad (6)$$

such that $h(x') = 1$. Here, $\nabla \mathcal{L}(h(x_t), h(x))$ represents the gradient of the loss function (BCE or wBCE), t the current iteration, and α the step size. $\text{Proj}_{\Delta(x)}$ denotes the projection of x_t onto the ℓ_∞ ball of radius r centered at x , $\Delta(x) = \{x_t \in \mathbb{R}^d : \|x_t - x\|_\infty \leq r\}$. This ensures that updates remain within the r -ball constraint. A successful improvement occurs when a negatively classified sample x transforms within the specified budget r into x' such that $h(x') = 1$ and $f^*(x') = 1$ (see Appendix E.3).

Results. We highlight three key findings from our evaluations (see also Appendix E.4). First, risk-averse (wBCE-trained) models consistently outperform standard (BCE-trained) models in reducing overall error as the improvement budget increases (Figures 2, 8, and 10). While error gains relative to BCE-trained models tend to cancel out, wBCE-trained models retain low false positive rates

after agent movement and exhibit a marked decline in false negatives as the improvement budget increases (Appendix, Figure 9). Among risk-averse strategies, loss-based risk aversion, where $\mathcal{L}_{\text{wBCE}}$ -trained models use $\frac{w_{\text{FP}}}{w_{\text{FN}}} > 1$ outperforms threshold-based approaches that classify agents as positive only if predicted probability exceeds 0.9 (Figure 2). Second, modest improvement budgets ($r \leq 2.0$) lead to substantial error reduction, but returns diminish for ($r > 2.0$), especially with a decision threshold of 0.5. Third, dataset class separability (Appendix, Figures 6 and 7) significantly affects the optimal level of risk aversion and the relationship between improvement budget and error reduction.

In summary, risk-averse models initially incur higher errors but achieve rapid error reduction as agents improve and r increases. A stricter false-positive penalty improves the positive agreement region, reducing test error, sometimes to zero (e.g., Appendix, Figure 10g).

7. Discussion

We propose a novel model for learning with strategic agents where the agents are allowed to improve. Surprisingly, we are able to achieve zero error (with high probability) by designing appropriate risk-averse learners for several well-studied concept classes, including a fairly general discrete graph-based model. We show that the VC dimension of the concept class is not the correct combinatorial dimension to capture learnability in the context of improvements. We further show that the intersection-closed property is sufficient, and in a certain sense necessary for proper learning with respect to arbitrary improvement sets. We leave open the questions of fully characterizing proper PAC learning with improvements and characterizing improper PAC learnability with improvements in terms of the concept class and the improvement sets available to the agents.

Acknowledgments

This work was supported in part by the National Science Foundation under grants CCF-2212968, ECCS-2216899, and ECCS-2217023, by the Simons Foundation under the Simons Collaboration on the Theory of Algorithmic Fairness, and by the Office of Naval Research MURI Grant N000142412742.

Impact Statement

Our work centers on individuals’ capability to strategically improve in response to decision-making algorithms, and the effect this has on the kinds of guarantees that one can hope to prove for machine learning algorithms. Our findings can potentially help inform decision-makers to design better algorithms. However, our main contributions are primarily theoretical in nature, and we anticipate no direct social impact of our work. We nonetheless believe there is a valuable opportunity to leverage our findings in ways that support both decision-makers and the individuals subject to these algorithms. To realize this potential, future research must critically assess the ethical and practical consequences of deploying improvement-aware algorithms, particularly in high-stakes domains.

References

- Ahmadi, S., Beyhaghi, H., Blum, A., and Naggita, K. The Strategic Perceptron. In *Proceedings of the ACM Conference on Economics and Computation (EC)*, pp. 6—25, 2021.
- Ahmadi, S., Beyhaghi, H., Blum, A., and Naggita, K. On Classification of Strategic Agents Who Can Both Game and Improve. In *Proceedings of the Symposium on Foundations of Responsible Computing (FORC)*, 2022.
- Ahmadi, S., Beyhaghi, H., Blum, A., and Naggita, K. Setting Fair Incentives to Maximize Improvement. In *Proceedings of the Symposium on Foundations of Responsible Computing (FORC)*, 2023.
- Ashtiani, H., Pathak, V., and Urner, R. Adversarially robust learning with tolerance. In *International Conference on Algorithmic Learning Theory*, pp. 115–135. PMLR, 2023.
- Ashtiani, H., Pathak, V., and Urner, R. Simplifying Adversarially Robust PAC Learning with Tolerance. In *Conference on Learning Theory (to appear)*. PMLR, 2025.
- Attias, I. and Hanneke, S. Adversarially robust PAC learnability of real-valued functions. In *International Conference on Machine Learning*, pp. 1172–1199. PMLR, 2023.
- Attias, I., Hanneke, S., and Mansour, Y. A characterization of semi-supervised adversarially robust PAC learnability. *Advances in Neural Information Processing Systems*, 35: 23646–23659, 2022a.
- Attias, I., Kontorovich, A., and Mansour, Y. Improved generalization bounds for adversarially robust learning. *Journal of Machine Learning Research*, 23(175):1–31, 2022b.
- Auer, P. Learning nested differences in the presence of malicious noise. *Theoretical Computer Science*, 185(1): 159–175, 1997.
- Auer, P. and Cesa-Bianchi, N. Online learning with malicious noise and the closure algorithm. *Annals of mathematics and artificial intelligence*, 23:83–99, 1998.
- Auer, P. and Ortner, R. A new PAC bound for intersection-closed concept classes. *Machine Learning*, 66(2):151–163, 2007.
- Awasthi, P., Balcan, M.-F., and Long, P. M. The power of localization for efficiently learning linear separators with noise. *Proceedings of the forty-sixth annual ACM Symposium on Theory of Computing (STOC)*, 2013.
- Balcan, M.-F. and Haghtalab, N. Noise in Classification. In *Beyond the Worst-Case Analysis of Algorithms*, 2020.
- Balcan, M.-F., Blum, A., Hanneke, S., and Sharma, D. Robustly-reliable learners under poisoning attacks. In *Conference on Learning Theory*, pp. 4498–4534. PMLR, 2022.
- Balcan, M.-F., Blum, A., Sharma, D., and Zhang, H. An analysis of robustness of non-Lipschitz networks. *Journal of Machine Learning Research*, 24(1), January 2023a. ISSN 1532-4435.
- Balcan, M.-F., Hanneke, S., Pukdee, R., and Sharma, D. Reliable learning in challenging environments. *Advances in Neural Information Processing Systems*, 36:48035–48050, 2023b.
- Becker, B. and Kohavi, R. Adult. <https://OPTdoi.org/10.24432/C5XW20>, 1996.
- Bloomfield, R. J. What Counts and What Gets Counted. Available at SSRN 2427106, 2016.
- Blum, A. and Saless, D. Regularized Robustly Reliable Learners and Instance Targeted Attacks. *arXiv preprint arXiv:2410.10572*, 2024.
- Braverman, M. and Garg, S. The Role of Randomness and Noise in Strategic Classification. In *Proceedings of the Symposium on Foundations of Responsible Computing (FORC 2020)*, 2020.

- Brückner, M. and Scheffer, T. Stackelberg games for adversarial prediction problems. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pp. 547–555. Association for Computing Machinery, 2011.
- Bshouty, N. H. and Burroughs, L. Maximizing Agreements with One-Sided Error with Applications to Heuristic Learning. *Machine Learning*, 59(1):99–123, 2005.
- Chawla, N. V., Bowyer, K. W., Hall, L. O., and Kegelmeyer, W. P. SMOTE: synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16(1): 321–357, Jun 2002.
- Chen, Y., Liu, Y., and Podimata, C. Learning strategy-aware linear classifiers. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS ’20, Red Hook, NY, USA, 2020. Curran Associates Inc. ISBN 9781713829546.
- Cohen, J., Rosenfeld, E., and Kolter, Z. Certified Adversarial Robustness via Randomized Smoothing. *International Conference on Machine Learning*, 2019.
- Cohen, L., Mansour, Y., Moran, S., and Shao, H. Learnability gaps of strategic classification. In *The Thirty Seventh Annual Conference on Learning Theory*, pp. 1223–1259. PMLR, 2024.
- Cullina, D., Bhagoji, A. N., and Mittal, P. PAC-learning in the presence of evasion adversaries. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS’18, pp. 228–239, Red Hook, NY, USA, 2018. Curran Associates Inc.
- Darnstädt, M. The optimal PAC bound for intersection-closed concept classes. *Information Processing Letters*, 115(4):458–461, 2015.
- Davies, D. L. and Bouldin, D. W. A Cluster Separation Measure. *IEEE Trans. Pattern Anal. Mach. Intell.*, 1(2): 224–227, Feb 1979.
- Dong, J., Roth, A., Schutzman, Z., Waggoner, B., and Wu, Z. S. Strategic Classification from Revealed Preferences. In *Proceedings of the ACM Conference on Economics and Computation (EC)*, 2018.
- El-Yaniv, R. and Wiener, Y. On the Foundations of Noise-free Selective Classification. *Journal of Machine Learning Research*, 11(53):1605–1641, 2010.
- Geifman, Y. and El-Yaniv, R. Selective Classification for Deep Neural Networks. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017.
- Goldman, S. and Kearns, M. On the Complexity of Teaching. *Journal of Computer and System Sciences*, 50(1):20–31, 1995.
- Haghtalab, N., Immorlica, N., Lucier, B., and Wang, J. Z. Maximizing Welfare with Incentive-Aware Evaluation Mechanisms. In Bessiere, C. (ed.), *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pp. 160–166. International Joint Conferences on Artificial Intelligence Organization, 7 2020. doi: 10.24963/ijcai.2020/23. Main track.
- Hardt, M., Megiddo, N., Papadimitriou, C., and Wootters, M. Strategic Classification. In *Proceedings of the ACM Conference on Innovations in Theoretical Computer Science*, pp. 111–122, 2016.
- Helmbold, D., Sloan, R., and Warmuth, M. K. Learning nested differences of intersection-closed concept classes. *Machine Learning*, 5(2):165–196, jun 1990.
- Hendrickx, K., Perini, L., Van der Plas, D., Meert, W., and Davis, J. Machine learning with a reject option: a survey. *Machine Learning*, 113(5):3073–3110, mar 2024.
- Hu, L., Immorlica, N., and Vaughan, J. W. The Disparate Effects of Strategic Manipulation. In *Proceedings of the Conference on Fairness, Accountability, and Transparency, FAT* ’19*, pp. 259–268, New York, NY, USA, 2019. Association for Computing Machinery.
- Kalai, A. T., Kanade, V., and Mansour, Y. Reliable agnostic learning. *Journal of Computer and System Sciences*, 78 (5):1481–1495, Sep 2012.
- Kleinberg, J. and Raghavan, M. How Do Classifiers Induce Agents to Invest Effort Strategically? *ACM Trans. Econ. Comput.*, 8(4), Oct 2020.
- Le Quy, T., Roy, A., Iosifidis, V., Zhang, W., and Ntoutsi, E. A survey on datasets for fairness-aware machine learning. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 12(3), 2022a.
- Le Quy, T., Roy, A., Iosifidis, V., Zhang, W., and Ntoutsi, E. Fairness datasets. https://github.com/tailequy/fairness_dataset/commit/f8726608fe1e7a77a3d3f14d7244e6b46e77aebd, 2022b.
- Lechner, T. and Urner, R. Learning losses for strategic classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pp. 7337–7344, 2022.
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. Towards Deep Learning Models Resistant to Adversarial Attacks. In *International Conference on Learning Representations*, 2018.

- Miller, J., Milli, S., and Hardt, M. Strategic Classification is Causal Modeling in Disguise. In III, H. D. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 6917–6926. PMLR, 13–18 Jul 2020.
- Milli, S., Miller, J., Dragan, A. D., and Hardt, M. The Social Cost of Strategic Classification. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pp. 230–239, 2019.
- Montasser, O., Hanneke, S., and Srebro, N. VC Classes are Adversarially Robustly Learnable, but Only Improperly. In Beygelzimer, A. and Hsu, D. (eds.), *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pp. 2512–2530. PMLR, 25–28 Jun 2019.
- Montasser, O., Hanneke, S., and Srebro, N. Reducing adversarially robust learning to non-robust PAC learning. *Advances in Neural Information Processing Systems*, 33: 14626–14637, 2020.
- Montasser, O., Hanneke, S., and Srebro, N. Adversarially robust learning with unknown perturbation sets. In *Conference on Learning Theory*, pp. 3452–3482. PMLR, 2021.
- Natarajan, B. K. On learning boolean functions. In *Proceedings of the nineteenth annual ACM symposium on Theory of computing (STOC)*, pp. 296–304, 1987.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Perdomo, J., Zrnic, T., Mendler-Dünner, C., and Hardt, M. Performative Prediction. In *Proceedings of the 37th International Conference on Machine Learning*, pp. 7599–7609, 2020.
- Raman, V., Subedi, U., and Tewari, A. On proper learnability between average-and worst-case robustness. *Advances in Neural Information Processing Systems*, 36: 13115–13126, 2023.
- Rivest, R. L. and Sloan, R. Learning complicated concepts reliably and usefully. In *Proceedings of the Seventh AAAI National Conference on Artificial Intelligence*, pp. 635–640, 1988.
- Rousseeuw, P. J. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20:53–65, 1987.
- Shahapure, K. R. and Nicholas, C. Cluster Quality Analysis Using Silhouette Score. In *2020 IEEE 7th International Conference on Data Science and Advanced Analytics (DSAA)*, pp. 747–748, 2020.
- Shavit, Y., Edelman, B., and Axelrod, B. Causal Strategic Linear Regression. In III, H. D. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 8676–8686. PMLR, 13–18 Jul 2020.
- Sundaram, R., Vullikanti, A., Xu, H., and Yao, F. PAC-learning for strategic classification. *The Journal of Machine Learning Research*, 24(1), Jan 2023.
- Zhang, H. and Conitzer, V. Incentive-Aware PAC Learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(6):5797–5804, May 2021.
- Zrnic, T., Mazumdar, E., Sastry, S. S., and Jordan, M. I. Who leads and who follows in strategic classification? In *Proceedings of the 35th International Conference on Neural Information Processing Systems, NIPS ’21*, Red Hook, NY, USA, 2021. Curran Associates Inc.

A. Additional Related Work

Classification of gaming agents. [Hardt et al. \(2016\)](#) formalized the concept of strategic behavior, often referred to as “gaming,” where test-set agents who are negatively classified intentionally modify their features—within the bounds of a separable cost function—without altering their target label, to deceive the model into classifying them as positive. They theoretically and empirically showed that their strategy-robust algorithm outperforms the standard SVM algorithm under gaming. However, as the extent of gaming increases, overall model accuracy declines. [Dong et al. \(2018\)](#) also study a Stackelberg equilibrium where agents strategically respond to classification learners. However, unlike [Hardt et al. \(2016\)](#), their model assumes that the learner lacks direct knowledge of the agents’ utility functions and instead infers them through observed revealed preferences. Additionally, agents arrive sequentially, and only the true negatives strategically respond to the learner. The learner’s objective is to minimize the Stackelberg regret. [Chen et al. \(2020\)](#) also study a learner whose goal is to minimize the Stackelberg regret, where gaming agents arrive sequentially. However, unlike [Dong et al. \(2018\)](#) which assumes a convex loss function, they deal with a less smooth agent utility function and learner loss function. They propose the Grinder algorithm, which adaptively partitions the learner’s action space based on the agents’ responses. Performative prediction ([Perdomo et al., 2020](#)) considers a setting that involves a repeated interaction between the classifier and the agents, and as a result the underlying distribution of the gaming agents may change over time.

Classification of agents that can both game and improve. Unlike earlier works in the strategic classification literature, which primarily focus on settings where agents engage in gaming behavior, [Kleinberg & Raghavan \(2020\)](#) examine a scenario where agents can genuinely improve. In this context, the agent can modify their observable features and true label to achieve a positive model outcome. The authors demonstrate that a learner employing a linear mechanism can encourage rational agents, who optimize their allocation of effort, to prioritize actions that result in meaningful improvement. They show how to achieve this by selecting an evaluation rule that incentivizes a desirable effort profile.

[Ahmadi et al. \(2022\)](#), like [Kleinberg & Raghavan \(2020\)](#), consider the agents’ potentially truthful and actionable responses to the model. However, the primary objective of [Ahmadi et al. \(2022\)](#) is to maximize true positive classifications while minimizing false positives. Notably, for the linear case, they show that the resulting classifier can become non-convex, depending on the agents’ initial positions.

On the other hand, [Ahmadi et al. \(2023\)](#) design reachable sets of target levels such that they can incentivize effort-bounded agents within each group to improve optimally.

Theoretical guarantees of incentive-aware or incentive-compatible classifiers. [Zhang & Conitzer \(2021\)](#) show that the vanilla ERM principle fails under strategic manipulation (gaming), even in simple scenarios that would otherwise be straightforward without gaming. To address this, they propose the concepts of incentive-aware and incentive-compatible ERM, theoretically analyzing the corresponding classifiers, their sample complexity, and the impact on learnability of the VC dimension of the concept hypothesis class. Finally, they extend their analysis to ERM-based learning in environments with transitive strategic manipulation.

Given adversarial data points wanting to receive an incorrect label, [Cullina et al. \(2018\)](#) theoretically show that the sample complexity of PAC-learning a set of halfspace classifiers does not increase in the presence of adversaries bounded by convex constraint sets and that the adversarial VC dimension can be arbitrarily larger or smaller than the standard VC dimension. [Sundaram et al. \(2023\)](#) provide theoretical guarantees for an offline, full-information strategic classification framework where data points have distinct preferences over classification outcomes (+ or −) and incur varying manipulation costs, modeled using seminorm-induced cost functions. They propose a PAC-learning framework for strategic linear classifiers in this setting, providing a detailed analysis of their statistical and computational learnability. Additionally, they extend the concept of the adversarial VC dimension ([Cullina et al., 2018](#)) to this strategic context. They also show, among other things, that employing randomized linear classifiers can substantially improve accuracy compared to deterministic methods.

Reliable machine learning. The concept of risk aversion in our work is closely related to selective classification or machine learning with a reject option ([El-Yaniv & Wiener, 2010](#); [Geifman & El-Yaniv, 2017](#); [Hendrickx et al., 2024](#)), where the classifier balances the trade-off between risk and coverage, opting to abstain from making predictions when it is likely to make mistakes. Similarly, risk-averse classification aligns with aspects of learning with one-sided error ([Natarajan, 1987](#); [Bshouty & Burroughs, 2005](#)), particularly positive reliable learners ([Kalai et al., 2012](#)), which aim to achieve zero false positive errors while minimizing false negatives. Prior work has shown connections between strategic classification and adversarial learning (e.g. [Sundaram et al. 2023](#)), but it remains an interesting open question if similar connections can be

established between learning with improvements and reliable learning in the presence of adversarial attacks (Balcan et al., 2022; 2023b; Blum & Saless, 2024).

Adversarial robustness. Our results indicate that the properties of the concept class that govern learnability with improvements are different from those established for adversarial robustness. In particular, finite VC dimension is not sufficient for PAC learnability with improvements (Example 2) as opposed to adversarial robustness, where it is sufficient for (improper) learning (Montasser et al., 2019). Furthermore, our risk-averse learning algorithm is very different from techniques proposed in the adversarial literature, including robust Empirical Risk Minimization (Cullina et al., 2018; Attias et al., 2022b), boosting algorithms whose generalization is based on sample compression schemes (Montasser et al., 2019; 2020; Attias et al., 2022a; Attias & Hanneke, 2023) and approaches developed for neural networks (Madry et al., 2018; Cohen et al., 2019; Balcan et al., 2023a). Similarly, our algorithm for learning halfspaces on a unit ball differs from classical approaches for learning from malicious noise (Awasthi et al., 2013; Balcan & Haghtalab, 2020). However, due to analogies between the frameworks, it is meaningful to ask future research questions along the lines of the directions pursued in the adversarial robustness literature—for example, extensions to unknown or uncertain improvement sets (Montasser et al., 2021; Lechner & Urner, 2022), and learning with tolerance (Ashtiani et al., 2023; Raman et al., 2023; Ashtiani et al., 2025).

B. Separating PAC Learning with Improvements from the Standard and Strategic PAC Models

Here, we further prove that learning with improvements diverges from the behavior of the standard PAC model for binary classification, and also from the more recently studied PAC learning model for strategic classification (Hardt et al., 2016; Sundaram et al., 2023).

B.1. Separation from Standard PAC Learning Model when $\tilde{\mathcal{H}} \subset \mathcal{H}$

Example 3 (Error gap when the learner’s hypothesis space $\tilde{\mathcal{H}}$ is a strict subset of the concept space \mathcal{H} that contains f^*). Let $\mathcal{X} = [-1, 1]$ and \mathcal{H} denote the set of concepts including unions of up to k open intervals. The set of possible improvements for any point $x \in \mathcal{X}$ is given by $\mathcal{X} \cap \mathbb{Q}$, where \mathbb{Q} denotes the set of rational numbers. Suppose the data distribution is uniform over \mathcal{X} . We set the target concept f^* as follows

$$f^*(x) = \begin{cases} 0, & \text{if } x < 0, \text{ or } x \in \mathbb{Q}, \\ 1, & \text{otherwise,} \end{cases}$$

Note that rationals are dense in $[0, 1]$ and the set of all rationals have a Lebesgue measure zero. Thus, on any finite sample $S \in \mathcal{X}^m$, any sampled point x will have a positive label according to f^* iff $x \geq 0$ (with probability 1). In the standard PAC learning setting, the classifier $\tilde{f} = \mathbb{I}\{x \in (0, 1)\}$ achieves zero error w.r.t. the target f^* . This is because the misclassification error for points in \mathbb{Q} is zero.

In our setting where agents have the ability to improve, for an $h \in \tilde{\mathcal{H}}$ which predicts any point x' in \mathbb{Q} as positive, all negative agents in $[-1, 0)$ can move to such a point x' and be falsely classified as positive. This corresponds to a lower bound of $\frac{1}{2}$ on the error. Since rationals are dense in the reals, any open interval which h classifies as positive must contain a point in \mathbb{Q} . On the other hand, if h classifies no point as positive, then error rate is again $\frac{1}{2}$ as all the positive points are misclassified.

B.2. Comparison with the PAC Model for Strategic Classification

We first observe that the strategic classification loss can be obtained by a subtle modification to our loss function (1),

$$\text{LOSS}^{\text{STR}}(x; h, f^*) = \max_{x' \in \Delta_h(x)} \mathbb{I}[h(x') \neq f^*(x)].$$

Intuitively, for a negative point with $f^*(x) = 0$, $\Delta_h(x)$ here denotes the set of points that the agent x can “pretend” to be in order to potentially deceive the classifier h into incorrectly classifying the agent positive. Since the movement within $\Delta_h(x)$ is viewed as a manipulation by the agent x , the prediction on the strategically perturbed point is compared with the original label of x , i.e. $f^*(x)$.

Prior work has shown that learnability in the strategic classification setting is captured by the *strategic VC dimension* (SVC) introduced by Sundaram et al. (2023). We state below the definition of SVC, adapted to our improvement-set based setting above which is a special case of the general cost functions studied under the strategic classification setting.

Definition B.1 (Strategic VC dimension, Sundaram et al. 2023). Define the n -th shattering coefficient of a strategic classification problem as

$$\sigma_n(\mathcal{H}, \Delta) = \max_{(x_1, \dots, x_n) \in \mathcal{X}^n} |\{(h(x'_1), \dots, h(x'_n)) : h \in \mathcal{H}, x'_i \in \Delta_h(x_i)\}|.$$

Then $\text{SVC}(\mathcal{H}, \Delta) = \sup\{n \geq 0 : \sigma_n(\mathcal{H}, \Delta) = 2^n\}$.

A natural question to ask is whether learning with improvements is “easier” than strategic classification. That is, if a concept space \mathcal{H} is learnable w.r.t. Δ and \mathcal{D} in the strategic classification setting, then is it also learnable with improvements? Interestingly, we answer this question in the negative. More precisely, we show that finite SVC (which is known from prior work to be a sufficient condition for strategic PAC learning) is actually not a sufficient condition for PAC learnability with improvements.

Theorem B.2. *Finite strategic VC dimension (Sundaram et al., 2023) does not necessarily imply PAC learnability with improvements.*

Proof. Let the instance space \mathcal{X} be $[0, 1]$, let $\mathcal{H} = \{h_{abcd} : h_{abcd}(x) = 1 \text{ iff } x \in [a, b) \cup (c, d]\}$, and let \mathcal{D} be the uniform distribution over $[0, 1]$. We define Δ as follows. For $x \in [0, 3/4]$ let $\Delta(x) = \mathcal{B}(x, 1/4) = (x - 1/4, x + 1/4) \cap [0, 1]$; for $x \in [3/4, 1]$, let $\Delta(x) = \{x\}$.

We claim that no algorithm with finite training data can guarantee an expected error of less than $1/16$ for the above when learning with improvements, even though the class is PAC-learnable in the strategic classification setting. To see the latter, note that $\text{SVC}(\mathcal{H}, \Delta) \leq 4$. Indeed, consider the points $(0, 1/4, 1/2, 3/4, 1) \in \mathcal{X}^5$. Notice the (strategic) labeling $(1, 0, 1, 0, 1)$ cannot be achieved for any $h \in \mathcal{H}$, which establishes the claim.

Now consider a target function defined as the union of two intervals $[1/2, b) \cup (b, 1]$ where the number b was randomly chosen in $[3/4, 1]$. The learner will not see the point b given a finite training set, so it learns nothing about the location of b (almost surely). Now, we consider two cases. Either, the learner outputs a classifier whose positive region has probability mass at most $1/16$ over the interval $[3/4, 7/8]$. Then its error rate is at least $1/16$ because the positive examples in $[3/4, 7/8]$ cannot move so at least half of their probability mass will get misclassified. On the other hand, if the learner outputs a classifier whose positive region has probability mass greater than $1/16$ on the interval $[3/4, 7/8]$, then it has at least a 50% chance of including the negative point b in its positive region (over the random choice of the target function). If the classifier has a negative point in $[3/4, 7/8]$ that is incorrectly predicted to be positive, then it will have an error rate at least $1/16$, because all the positives in $[5/8, 3/4]$ will move to a false positive (here we use that agents break ties adversarially, see also Remark 2.1). So, either way, its expected error is $\geq 1/16$. \square

On the other hand, it is not too hard to come up with examples where it is easier to learn in the improvements setting when compared to the strategic setting. Example 4 shows that it is possible to learn perfectly with improvements (with zero error) in a setting where avoiding a large constant error is unavoidable in the strategic classification setting.

Example 4 (Learnability with improvements may be easier than strategic classification). Define $\Delta(x) = \mathcal{X}$ for all examples $x \in \mathcal{X}$. Suppose the “all-negative” classifier $h_-(x) = 0$, the “all-positive” classifier $h_+(x) = 1$, and all “singleton-positive” classifiers $h_{x^*}(x) = \mathbb{I}[x = x^*]$ lie in the concept space \mathcal{H} . Select any $f^* \in \mathcal{H}$ and any data distribution \mathcal{D} over \mathcal{X} such that $\mathbb{P}_{x \sim \mathcal{D}}[f^*(x) = 0] = \mathbb{P}_{x \sim \mathcal{D}}[f^*(x) = 1] = \frac{1}{2}$. Now with $O(\log \frac{1}{\delta})$ examples, the learner sees a positive example, say x^+ , in its training set with probability $\geq \delta$. Outputting h_{x^+} achieves zero-error in the learning with improvements setting, as all negative points can improve to x^+ . In contrast, a learner in the strategic classification setting must suffer an error of at least $1/2$ here. Indeed, either the learner outputs h_- and suffers an error of $1/2$ on the positive points. Or, the learner selects an h that labels at least one point as positive and incurs an error $1/2$ on the negative points, all of which successfully deceive the learner.

Furthermore, let’s consider an improvement function Δ that takes into account f^* , such that $f^*(x') = 1$ for $x' \in \Delta(x)$ for any $x \in \mathcal{X}$. That is the improvement function is in a certain sense consistent with f^* , guaranteeing positive classification after any move. In this setting, any classifier h will have lower error in the improvements setting compared to strategic classification. This is because a negative point that moves and becomes positive is an error in strategic learning but the point

would have genuinely improved in this case. Contrasting this with Remark 2.1, when Δ does not satisfy the above property, we note that the reason it is possible to do worse in the improvement setting (e.g. Theorem B.2) is because some true positive examples can potentially become negative when moving in response to a false positive for the learner's hypothesis h .

C. Missing Proofs from Section 4

We include below missing proofs from Section 4.

C.1. Proof of Theorem 4.1: Learning Thresholds with the Uniform Distribution

Proof. Let $S \stackrel{\text{i.i.d.}}{\sim} \mathcal{D}^m$, where \mathcal{D} is the uniform distribution over $[0, 1]$. By using a standard calculation of the sample complexity of thresholds,

$$\begin{aligned} \mathbb{P}_{S \sim \mathcal{D}^m} \left[\mathbb{P}_{x \sim \mathcal{D}} [h_{t^*}(x) \neq h_{S^+}(x)] > \epsilon \right] &\leq \prod_{i=1}^m \mathbb{P}[x_i \notin [t^*, t^* + \epsilon]] \\ &\leq (1 - \epsilon)^m \\ &\leq e^{-\epsilon m} \\ &\leq \delta, \end{aligned} \tag{7}$$

where the last inequality holds for $m \geq \frac{1}{\epsilon} \log \frac{1}{\delta}$. Since whenever h_{S^+} classifies a point as positive, h_{t^*} also classifies it as positive, any negative point that improves in response to h_{S^+} must move to a true positive point, and the error can only decrease in the improvements setting for the choice of h_{S^+} .

Now, since we allow improvements of distance r , the points in the interval $[t_{S^+} - r, t_{S^+}]$ that would have been classified negatively without improvement are able to improve under h_{S^+} (and indeed improve to be positive with respect to h_{t^*}) and are thus classified correctly. The points on which h_{S^+} makes mistakes are those in the interval $[t^*, t_{S^+} - r]$. Since \mathcal{D} is uniform, our previous inequality implies that with probability at least $1 - \delta$ we have $t_{S^+} \leq t^* + \epsilon$. This implies the that error is at most $\max(\epsilon - r, 0)$ with probability $1 - \delta$ as desired. \square

C.2. Learning Thresholds with An Arbitrary Distribution

Theorem C.1 (Thresholds, arbitrary distribution). *Let the improvement set Δ be the closed ball with radius r , $\Delta(x) = \{x' \mid |x - x'| \leq r\}$. For any distribution \mathcal{D} , and any $\epsilon, \delta \in (0, 1/2)$, with probability $1 - \delta$,*

$$\text{LOSS}_{\mathcal{D}}(h_{S^+}, h_{t^*}) \leq (\epsilon - p(h_{S^+}; h_{t^*}, \mathcal{D}, r))_+,$$

where

$$p(h_{S^+}; h_{t^*}, \mathcal{D}, r) = \mathbb{P}_{x \sim \mathcal{D}} [x \in [t_{S^+} - r, t_{S^+}]],$$

with sample complexity $M = O\left(\frac{1}{\epsilon} \log \frac{1}{\delta}\right)$.

Proof. Let t_0 be such that $\mathbb{P}_{x \sim \mathcal{D}} [x \in [t^*, t_0]] = \epsilon$. By following the same derivation as in Eqn. 7 and replacing $t^* + \epsilon$ with t_0 , we get that

$$\mathbb{P}_{S \sim \mathcal{D}^m} \left[\mathbb{P}_{x \sim \mathcal{D}} [h_{t^*}(x) \neq h_{S^+}(x)] \leq \epsilon \right],$$

with probability $1 - \delta$ for $m \geq \frac{1}{\epsilon} \log \frac{1}{\delta}$.

The points in the interval $[t_{S^+} - r, t_{S^+}]$ are able to improve under h_{S^+} and thus classified correctly. The gain to the error of h_{S^+} would be the probability mass of points in \mathcal{D} that fall into this interval, defined as

$$p(h_{S^+}; h_{t^*}, \mathcal{D}, r) = \mathbb{P}_{x \sim \mathcal{D}} [x \in [t_{S^+} - r, t_{S^+}]].$$

The points on which h_{S^+} makes mistakes are those in the interval $[t^*, t_{S^+} - r]$.

We conclude that with probability at least $1 - \delta$ we have $t_{S^+} \leq t_0$, and given the improvement of points in $[t^*, t_{S^+} - r]$ we have an error at most $\max(\epsilon - p(h_{S^+}; h_{t^*}, \mathcal{D}, r), 0)$ with probability $1 - \delta$ as desired. \square

C.3. Proof of Theorem 4.6

Proof. Let $R_S^c = \text{CLOS}_{\mathcal{H}_{\text{rec}}}(\{x_i \in S : y_i = 1\})$ be the output of the closure algorithm. For the standard PAC setting, we have

$$\mathbb{P}_{S \sim \mathcal{D}^m} \left[\mathbb{P}_{x \sim \mathcal{D}} [R_S^c(x) \neq R^*(x)] \leq \epsilon \right],$$

with probability $1 - \delta$ for $m \geq \Omega\left(\frac{1}{\epsilon} \left(d + \log \frac{1}{\delta}\right)\right)$, see [Darnstädt \(2015\)](#). Since whenever R_S^c classifies a point as positive, R^* also classifies it as positive, any negative point that improves in response to R_S^c must move to a true positive point and the error can only decrease in the improvements setting for the choice of R_S^c .

Now, in order to quantify the gain in error from the improvements, we define the “outer boundary strip” of R_S^c . Let the rectangle defined by $R_S^c = \prod_{i \in [d]} [a_i, b_i]$. The points that are able to improve under R_S^c are exactly fall into the outer boundary strip of size r , defined as

$$\text{BS}(R_S^c, r) = \prod_{i \in [d]} [a_i + r, b_i + r] \setminus \prod_{i \in [d]} [a_i, b_i].$$

Note that this is exactly the improvement region of R_S^c : $\text{BS}(R_S^c, r) = \text{IR}(R_S^c; R^*, \Delta)$. Under general distribution \mathcal{D} , the probability mass of the improvement region is

$$\mathbb{P}_{x \sim \mathcal{D}} [x \in \text{IR}(R_S^c; R^*, \Delta)] = \mathbb{P}_{x \sim \mathcal{D}} [x \in \text{BS}(R_S^c, r)],$$

since R_S^c is the smallest rectangle that fits S , these points that are able to improve under R_S^c indeed improve to be positive with respect to R^* . This implies that

$$\text{LOSS}_{\mathcal{D}}(R_S^c, R^*) \leq \max(\epsilon - \mathbb{P}_{x \sim \mathcal{D}} [x \in \text{IR}(R_S^c; R^*, \Delta)], 0).$$

Now, for the uniform distribution, we can compute an exact expression of the improvement region. Let $l_i = b_i - a_i$, then

$$\mathbb{P}_{x \sim \mathcal{D}} [x \in \text{IR}(R_S^c; R^*, \Delta)] = \prod_{i \in [d]} (l_i + 2r) - \prod_{i \in [d]} l_i.$$

For $d = 2$, we get

$$\begin{aligned} \mathbb{P}_{x \sim \mathcal{D}} [x \in \text{IR}(R_S^c; R^*, \Delta)] &= (l_1 + 2r)(l_2 + 2r) - l_1 l_2 \\ &= 2r(l_1 + l_2) + 4r^2. \end{aligned}$$

\square

C.4. Proof of Theorem 4.7

Proof. Let $S \sim \mathcal{D}^m$ and h_S^c denote the classifier learned by the closure algorithm (Definition 4.4). For some $m = O\left(\frac{1}{\epsilon} (d_{\text{VC}}(\mathcal{H}) + \log \frac{1}{\delta})\right)$, we know from prior work ([Auer & Ortner, 2007](#); [Darnstädt, 2015](#)) that h_S^c satisfies, with probability at least $1 - \delta$,

$$\mathbb{P}_{x \sim \mathcal{D}} [h_S^c(x) \neq f^*(x)] \leq \epsilon,$$

for any target concept $f^* \in \mathcal{H}$.

Now for any $x \in \mathcal{X}$ if $h_S^c(x) = 1$, the improvement loss $\text{LOSS}(x; h_S^c, f^*) = \mathbb{I}[h_S^c(x) \neq f^*(x)] = 0$ since $\Delta_{h_S^c}(x) = \{x\}$ and $f^*(x) = 1$ since h_S^c is obtained using the closure algorithm. If $h_S^c(x) = 0$ and if $\Delta_{h_S^c}(x) \neq \{x\}$, for any point $x' \in \Delta_{h_S^c}(x)$, we have $h_S^c(x') = 1 = f^*(x')$ and therefore $\text{LOSS}(x; h_S^c, f^*) = 0$. So the only points for which h_S^c can make a mistake are points where $h_S^c(x) = 0$ and $\Delta_{h_S^c}(x) = \{x\}$, i.e. the points do not move in reaction to h_S^c . This implies h_S^c must disagree with f^* on these points also in the PAC setting. But the probability mass of these points is at most ϵ as noted above. \square

C.5. Proof of Theorem 4.8

Proof. For any concept $h \in \mathcal{H}$, let \mathcal{X}_h^+ denote the set of points $\{x \in \mathcal{X} \mid h(x) = 1\}$ positively classified by h . Since $\mathcal{H}' := \mathcal{H} \setminus \{x'\}$ is not intersection-closed, there must exist a set $S \subseteq \mathcal{X} \setminus \{x'\}$ such that $\text{CLOS}_{\mathcal{H}'}(S) \notin \mathcal{H}'$. For the uniformly negative point x' , we have its improvement set as $\Delta(x') = \mathcal{X} \setminus S$. For points in $\text{CLOS}(S)$ we have the improvement set as the empty set. We set the data distribution \mathcal{D} as the uniform distribution over $\text{CLOS}(S) \cup \{x'\}$. Let $h_1 \in \mathcal{H}$ be a minimally consistent classifier w.r.t. S , i.e. if $h' \in \mathcal{H}$ and $\mathcal{X}_{h'}^+ \subseteq \mathcal{X}_{h_1}^+$, then $h' = h_1$. By choice of S , there is a point $x_1 \in \mathcal{X}_{h_1}^+ \setminus \text{CLOS}_{\mathcal{H}'}(S)$. By the definition of closure of S , there must exist $h_2 \in \mathcal{H}$ consistent with S (assumed minimally consistent WLOG) such that $h_2(x_1) = 0$. Also, since h_1 was chosen to be minimally consistent, there must exist $x_2 \in \mathcal{X}_{h_2}^+$ such that $h_1(x_2) = 0$. We will set the target concept f^* to one of h_1 or h_2 .

Now any learner that picks a concept not consistent with S will clearly suffer a constant error on the points in $\text{CLOS}(S)$ which are incorrectly classified as negative and not allowed to improve. Suppose therefore that the learner selects a hypothesis h consistent with S . Let \tilde{h} denote a classifier which is minimally consistent with S and $\mathcal{X}_{\tilde{h}}^+ \subseteq \mathcal{X}_h^+$ (\tilde{h} could possibly be the same as h). If $\tilde{h} = h_1$ (resp. $\tilde{h} = h_2$), the learner suffers a constant error as x' can improve to the false positive x_1 (resp. x_2) when the target concept is h_2 (resp. h_1). Else, there must exist $\tilde{x} \in \mathcal{X}_{\tilde{h}}^+$ such that $h_1(\tilde{x}) = 0$, since h_1 was chosen to be minimally consistent (and likewise for h_2). $h(\tilde{x}) = 1$ in this case, and x' can now falsely “improve” to \tilde{x} . Since the learner has no way of knowing from the sample whether the target is h_1 or h_2 , it must suffer a constant error for any h it selects from \mathcal{H} . \square

D. Missing Proofs and Additional Definitions from Section 5

We include below additional definitions and complete proofs for results in Section 5.

D.1. Additional Definitions

Definition D.1. The shortest path metric $d_G : V \times V \rightarrow [0, |V| + 1]$ is defined as follows:

$$d_G(x, x') = \begin{cases} \min \{k \mid \exists (x_0 = x, x_1, \dots, x_k = x') \subseteq V, (x_{i-1}, x_i) \in E \forall i \in [k]\}, & \text{if a path exists,} \\ |V| + 1, & \text{if no path exists.} \end{cases}$$

Here, k is the length of the shortest path between x and x' in terms of the number of edges. If there is no path connecting x and x' , the distance is defined as $|V| + 1$. The shortest path metric satisfies the following properties:

- **Non-negativity:** $d_G(x, x') \geq 0$ for all $x, x' \in V$, with $d_G(x, x) = 0$.
- **Symmetry:** $d_G(x, x') = d_G(x', x)$ for all $x, x' \in V$, since G is undirected.
- **Triangle inequality:** $d_G(x, x') \leq d_G(x, z) + d_G(z, x')$ for all $x, x', z \in V$.

Our results in Section 5 extend to the more general improvement function $\Delta(x) = \{x' \in V \mid d_G(x, x') \leq r\}$ by applying our arguments to G^r , the r th power of G .

Definition D.2 (Dominating Set). Let $G = (V, E)$ be an undirected graph, where V is the set of vertices and $E \subseteq V \times V$ is the set of edges. A subset of vertices $S \subseteq V$ is called a dominating set if every vertex in V is either in S or adjacent to at least one vertex in S . Formally, S is a dominating set if:

$$\forall x \in V, \quad x \in S \text{ or } \exists x' \in S \text{ such that } (x, x') \in E.$$

D.2. Proof of Theorem 5.1

Proof. Given a sample S labeled by f^* , let $S^+ = \{x \in S \mid f^*(x) = +1\}$ denote the set of positive points in S . To achieve the claimed upper bound on the sample complexity of learning with zero-error, the learner outputs h_S with $h_S(x) = \mathbb{I}\{x \in S^+\}$, that is the classifier which positively classifies exactly the points in S^+ . We will now show that with a sample of size $m = O\left(\frac{n(\log n + \log \frac{1}{\delta})}{d_{\min}^+ + 1}\right)$, the proposed h_S achieves zero generalization error with probability at least $1 - \delta$.

Let $V^+ = \{x \in V \mid f^*(x) = +1\}$ denote the set of vertices in G^+ . We say that $x \in V^+$ is *covered* by the sample S if $x \in S$ or there exists $x' \in S$ such that $x' \in V^+$ and $(x, x') \in E$. Note that if every $x \in V^+$ is covered by S , then $\text{LOSS}_{\mathcal{D}}(h_S, f^*) = 0$ (formally established in Theorem 5.2). It is therefore sufficient to bound determine the sample size needed to guarantee that every positive vertex is covered with high probability.

Let $x \in V^+$ be a vertex in the positive subgraph G^+ . For x to be covered, it must either be included in the sample S , or have at least one of its neighbors $x' \in \Delta(x)$ included in S . The probability of sampling x directly in one draw is $\frac{1}{n}$. The probability of sampling any of its neighbors is proportional to its degree in G^+ . Thus, the total probability of covering x in one draw is:

$$p_{\text{cover}}(x) = \frac{1}{n} \cdot (1 + d(x)),$$

where $d(x)$ is the degree of x in G^+ . Since $d(x) \geq d_{\min}^+$, we have:

$$p_{\text{cover}}(x) \geq \frac{d_{\min}^+ + 1}{n}.$$

To ensure the desired coverage holds with probability at least $1 - \delta$, we analyze the failure probability for a single vertex. The probability that a given vertex $x \in V^+$ is not covered after m samples is

$$\mathbb{P}[x \text{ is not covered}] \leq \left(1 - \frac{d_{\min}^+ + 1}{n}\right)^m.$$

To ensure that this holds for all $|V^+| \leq n$ vertices, we apply the union bound

$$\mathbb{P}[\exists x \in V^+ \text{ not covered}] \leq n \left(1 - \frac{d_{\min}^+ + 1}{n}\right)^m.$$

Therefore, the sample size m required to ensure that the probability of the above bad event is at most δ is given by

$$m = O\left(\frac{n(\log n + \log \frac{1}{\delta})}{d_{\min}^+ + 1}\right).$$

To establish the lower bound, consider a graph G on n vertices consisting of k disjoint cliques, each of the same size $\frac{n}{k} = d_{\min}^+ + 1$, see Figure 3. Now note that if our sample S does not contain any node from any one of the cliques (say C), then zero-error is not possible. This is because, one of two cases occur. If the learner's hypothesis h predicts any point in C as positive then we can select an f^* that predicts C entirely as negative while being consistent with S , causing the population loss to be at least $\frac{1}{k}$. On the other hand, if the learned h predicts all points in C as negative, then we can set f^* to label C entirely as positive, again incurring a population loss of at least $\frac{1}{k}$.

Our goal therefore is to determine a lower bound on the number of points required to ensure that every clique has at least one of its vertices included in the training sample S , which ensures that for every positive vertex, either the vertex itself or one of its neighbors is included. Using the standard coupon collector analysis, the number of trials needed to collect $k = \frac{n}{d_{\min}^+ + 1}$ coupons is $\Omega(k \log k)$ with high constant probability.

□




Figure 3. The graph G used to establish our lower bound on the zero-error sample complexity. The graph consists of k components, each of size $\frac{n}{k}$.

D.3. Enabling Improvement Whenever It Helps

Theorem D.3. Let $G = (V, E)$ be an undirected graph with $n = |V|$ vertices, and let $f^* : V \rightarrow \{0, +1\}$ denote the ground truth labeling function. Define: $V^+ = \{x \in V \mid f^*(x) = +1\}$, the set of positive vertices. Define $V^- = \{x \in V \mid f^*(x) = 0\}$ the set of negative vertices, and $N = \{x \in V^- \mid \exists x' \in V^+ \text{ such that } (x, x') \in E\}$ denote the set of negative vertices that have a positive neighbor. Let $d_{\min}^N = \min_{x \in N} |\{x' \in V^+ \mid (x, x') \in E\}|$ denote the minimum number of positive neighbors of vertices in N . Assume that the data distribution \mathcal{D} is uniform on V . For any $\delta > 0$, and training sample $S \stackrel{i.i.d.}{\sim} \mathcal{D}^m$ of size $m = O\left(\frac{n(\log n + \log \frac{1}{\delta})}{d_{\min}^N}\right)$, there exists a learner that outputs a hypothesis h such that $\text{LOSS}_{\mathcal{D}}^E(h, f^*) = 0$, with probability at least $1 - \delta$ over the draw of S . Moreover, there exists a graph G for which any learner that always outputs h with $\text{LOSS}_{\mathcal{D}}^E(h, f^*) = 0$ for any \mathcal{D}, f^* must see at least $\Omega\left(\frac{n}{d_{\min}^N} \log \frac{n}{d_{\min}^N}\right)$ labeled points in the training sample, with high constant probability.

Proof. The proof of the upper bound is technically similar to the proof of Theorem 5.1. Essentially, to ensure that $\text{LOSS}^E = 0$, we need to cover all the vertices in N by some vertices in V^+ . The probability that any fixed node in N is covered by a random sample can be lower-bounded in terms of d_{\min}^N as

$$p_{\text{cover}}(x) \geq \frac{d_{\min}^N}{n}.$$

Using the same argument as in Theorem 5.1, we obtain an upper bound of $m = O\left(\frac{n(\log n + \log \frac{1}{\delta})}{d_{\min}^N}\right)$ on the sample complexity of the classifier h which outputs exactly the positively labeled points in its sample as positive to guarantee that $\text{LOSS}_{\mathcal{D}}^E(h, f^*) = 0$ with probability at least $1 - \delta$.

To establish the lower bound, consider a graph $G = (V, E)$ with two types of nodes, i.e. $V = V_1 \cup V_2$, $|V_1| = k$, $|V_2| = n - k$. f^* labels all nodes in V_1 as negative. Each node $x_i \in V_1$ has $d_{\min}^N = \frac{n-k}{k}$ neighboring nodes Δ_i in V_2 and the sets of these neighbors are pairwise disjoint. Now, suppose our training sample S does not contain any point in Δ_i for some $i \in [k]$. If the learned hypothesis h predicts any point Δ_i as positive, we have $h(x_i) \neq \{x_i\}$ but if f^* labels all points in Δ_i negative, $f^*(x_i) = \{x_i\}$ and we incur loss corresponding to x_i . Similarly, if h labels all points in Δ_i as negative then $h(x_i) = \{x_i\}$ but we can label Δ_i consistent with S such that $f^*(x_i) \neq \{x_i\}$.




Figure 4. The graph G used to establish our lower bound in Theorem D.3.

Therefore it is sufficient to determine a lower bound on the number of points required to ensure that every Δ_i has at least one of its vertices included in the training sample S . Using the standard coupon collector analysis, the number of trials needed to collect $k = \frac{n}{d_{\min}^N + 1}$ coupons is $\Omega(k \log k)$ with high constant probability. \square

Theorem D.4. Let $G = (V, E)$ be an undirected graph with $n = |V|$ vertices, and let $f^* : V \rightarrow \{0, +1\}$ denote the ground truth labeling function. Define $V^+ = \{x \in V \mid f^*(x) = +1\}$, the set of positive vertices and d_{\min}^+ denote the minimum degree of a vertex in V^+ in the subgraph of G induced by V^+ . Define $V^- = \{x \in V \mid f^*(x) = 0\}$ the set of negative vertices, and $N = \{x \in V^- \mid \exists x' \in V^+ \text{ such that } (x, x') \in E\}$ denote the set of negative vertices that have a positive neighbor. Let $d_{\min}^N = \min_{x \in N} |\{x' \in V^+ \mid (x, x') \in E\}|$ denote the minimum number of positive neighbors of vertices in N . Assume that the data distribution \mathcal{D} is uniform on V . For any $\delta > 0$, and training sample $S \stackrel{i.i.d.}{\sim} \mathcal{D}^m$ of size $m = O\left(\frac{n(\log n + \log \frac{1}{\delta})}{\min\{d_{\min}^N, d_{\min}^+\}}\right)$, there exists a learner that outputs a hypothesis h such that $\text{LOSS}_{\mathcal{D}}(h, f^*) = 0$ and $\text{LOSS}_{\mathcal{D}}^E(h, f^*) = 0$, with probability at least $1 - \delta$ over the draw of S . Moreover, there exists a graph G for which any learner that always outputs h with $\text{LOSS}_{\mathcal{D}}(h, f^*) = \text{LOSS}_{\mathcal{D}}^E(h, f^*) = 0$ for any \mathcal{D}, f^* must see at least $\Omega\left(\max\left\{\frac{n}{d_{\min}^+} \log \frac{n}{d_{\min}^+}, \frac{n}{d_{\min}^N} \log \frac{n}{d_{\min}^N}\right\}\right)$ labeled points in the training sample, with high constant probability.

Proof. See Theorems 5.1 and D.3. □

D.4. Proof of Theorem 5.2: Teaching a Risk-averse Student

Proof. Since S^+ is a dominating set of $G^+ = (V^+, E^+)$, for any $x \in V^+$, either $x \in S^+$ or there exists $x' \in S^+$ such that $(x, x') \in E^+$. In the first case, $h_{S^+}(x) = 1$ and therefore $\Delta_{h_{S^+}}(x) = \{x\}$. Thus, $h_{S^+}(x) = 1 = f^*(x)$ implies that $\text{LOSS}(x; h_{S^+}, f^*) = 0$ in this case. In the second case, $x \notin S^+$, but there exists a neighbor $\tilde{x} \in \Delta(x)$ such that $\tilde{x} \in S^+$ by the definition of S^+ . Thus, for any point $x' \in \Delta_{h_{S^+}}(x) \subseteq S^+$, we have that $h_{S^+}(x') = 1 = f^*(x')$, ensuring that $\text{LOSS}(x; h_{S^+}, f^*) = 0$ in this case as well.

For $x \in V \setminus V^+$, if x has no positive neighbors in S^+ , $\Delta_{h_{S^+}}(x) = \{x\}$ because there is no neighboring vertex $x' \in \Delta(x)$ that would induce a reaction. Thus, $h_{S^+}(x) = 0 = f^*(x)$ in this case, implying the loss $\text{LOSS}(x; h_{S^+}, f^*)$ on x is zero.

Finally, if $x \in V \setminus V^+$ has positive neighbors contained in the dominating set, i.e., $\Delta(x) \cap S^+ \neq \emptyset$, then $h_{S^+}(x') = +1$. The reaction set $\Delta_{h_{S^+}}(x)$ ensures that x moves to one of these neighbors. Specifically, the reaction set allows x to improve and move to a neighboring vertex $x' \in \Delta(x) \cap S^+$ such that $f^*(x') = +1$. Thus, $h(x') = f^*(x') = +1$ for any $x' \in \Delta_{h_{S^+}}(x)$ implying $\text{LOSS}(x; h_{S^+}, f^*) = 0$. □

E. Evaluation: Supplementary Details

This section includes supplementary details on the datasets and classifiers used, how improvement is done and results of the empirical evaluations.

E.1. Datasets

We utilize three real-world datasets: the Adult Income dataset from UCI and the Open University Learning Analytics Dataset (OULAD) and Law School datasets sourced from Le Quy et al. (Le Quy et al., 2022b). The preprocessing steps for all the datasets, similar to those described in Le Quy et al. (Le Quy et al., 2022b), include removing missing data and applying label encoding to categorical variables. In addition to the real-world datasets, we generate an 8-dimensional synthetic dataset with increased separability ($class_sep = 4$) using the *make_classification* function from Scikit-learn. We clean the dataset by removing duplicates and outliers, with Z-scores applied with thresholds (0.9 for class 0 and 0.8 for class 1). The cleaned synthetic dataset is then balanced using SMOTE (Chawla et al., 2002) to ensure class balance.

Statistical details of the datasets, including test/train sizes and number of features, are in Table 1. We examine the structural variations within datasets to gain deeper insights into how the characteristics influence the impact of improvements on error drop rates. Figure 5 highlights the target distribution across training datasets: the Adult dataset has a higher proportion of negative examples, whereas the OULAD and Law School datasets have a higher percentage of positive examples. The synthetic dataset, by contrast, is balanced. Figure 6 and 7 illustrate dataset separability properties, showing that the synthetic dataset (k -NN error: 0.1016) and the Law School dataset (k -NN error: 0.1010) have the highest separability. However, as Figure 7 shows, the Adult and synthetic datasets exhibit the lowest false positive (FP) outlier rates.

Dataset	Target variable	d	Train/Test	Improvable features
Adult	$\{1(> 50K), 0(\leq 50K)\}$	14	21113/9049	{“hours-per-week, capital-gain, capital-loss, fnlwgt, educational-num, workclass, education, occupation”}
OULAD	$\{1(\text{pass}), 0(\text{fail})\}$	11	15093/6469	{“code_module, code_presentation, imd_band, highest_education, num_of_prev_attempts, studied_credits”}
Law school	$\{1(\text{pass}), 0(\text{fail})\}$	11	14558/6240	{“decile1b, decile3, lsat, ugpa, zfygpa, zgpa, fulltime, fam_inc, tier”}
Synthetic	$\{1(\text{positive}), 0(\text{negative})\}$	8	1561/669	{all features are used}

Table 1. Details of the tabular datasets, both synthetic and real, used in the experiments.




Figure 5. Target variable distributions of synthetic and real-world train datasets used in the experiments for the ((a)) Adult, ((b)) OULAD, ((c)) Law School, and ((d)) Synthetic datasets.




Figure 6. Inspection of clusteredness and class separation using Davies–Bouldin index (Davies & Bouldin, 1979) and Silhouettes scores (Rousseeuw, 1987; Shahapure & Nicholas, 2020).

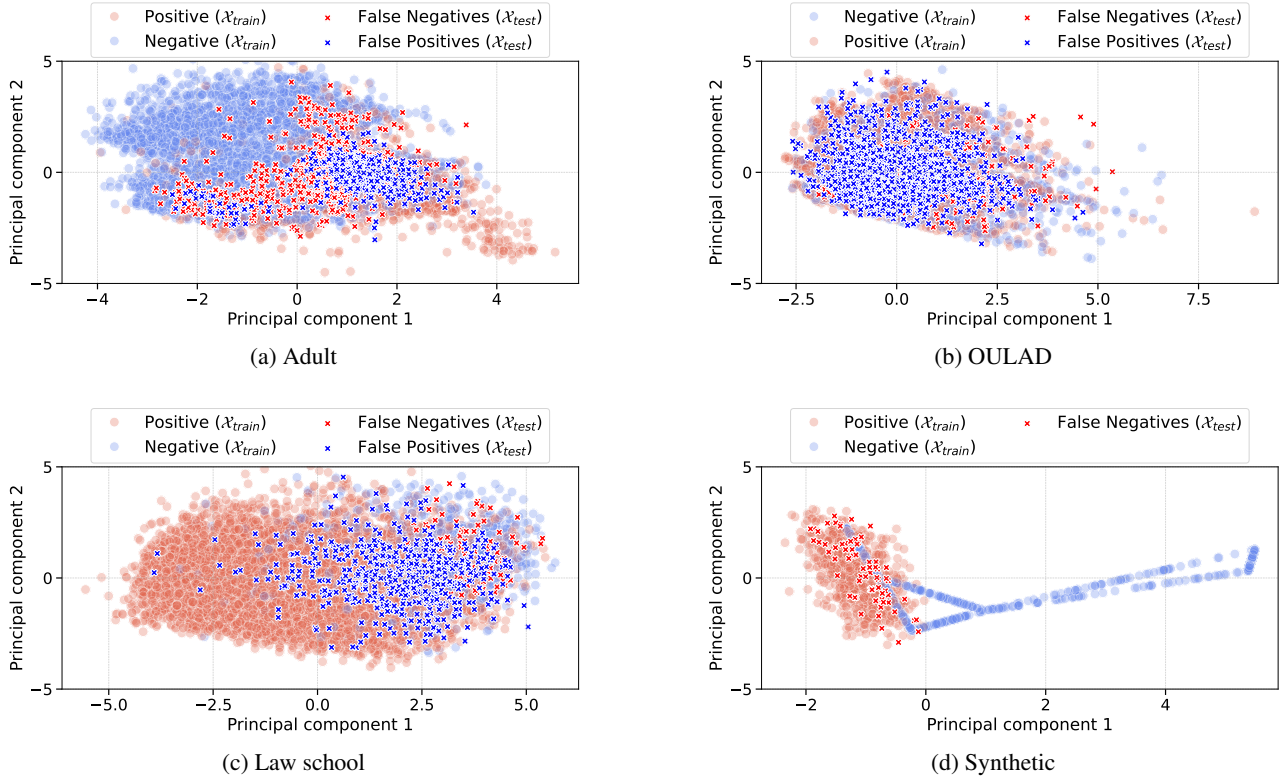


Figure 7. Scatter plots of the two principal components of the training data and of the k -NN misclassification on test data for the ((a)) Adult (k -NN error: 0.1670, FNR: 0.4610, FPR: 0.0678), ((b)) OULAD (k -NN error: 0.3291, FNR: 0.1195, FPR: 0.7645), ((c)) Law School (k -NN error: 0.1010, FNR: 0.0180, FPR: 0.7875), and ((d)) Synthetic (k -NN error: 0.1016, FNR: 0.1960, FPR: 0.000) datasets.

Dataset	DTC1	DTC2	RFC1	RFC2	XGB
Adult	0.999967 \pm 0.000049	0.999967 \pm 0.000049	0.999934 \pm 0.000079	0.999967 \pm 0.000049	0.999967 \pm 0.000049
Law	1.000000 \pm 0.000000	1.000000 \pm 0.000000	1.000000 \pm 0.000000	1.000000 \pm 0.000000	0.999952 \pm 0.000071
OULAD	1.000000 \pm 0.000000	1.000000 \pm 0.000000	1.000000 \pm 0.000000	1.000000 \pm 0.000000	1.000000 \pm 0.000000
Synthetic	1.000000 \pm 0.000000	1.000000 \pm 0.000000	1.000000 \pm 0.000000	1.000000 \pm 0.000000	1.000000 \pm 0.000000

 Table 2. Accuracy score of the f^* models when trained and tested on \mathcal{S}_T across different datasets.

Dataset	DTC1 (LOO)	DTC2 (LOO)	RFC1 (LOO)	RFC2 (LOO)	XGB (LOO)
Adult	0.8084 \pm 0.0044	0.8053 \pm 0.0045	0.8541 \pm 0.0040	0.8545 \pm 0.0040	—
Law	0.8507 \pm 0.0048	0.8428 \pm 0.0049	0.8962 \pm 0.0041	0.8972 \pm 0.0041	0.8894 \pm 0.0043
OULAD	0.5941 \pm 0.0066	0.5952 \pm 0.0066	0.6684 \pm 0.00663	0.6689 \pm 0.0063	—
Synthetic	0.9955 \pm 0.0028	0.9951 \pm 0.0029	0.9969 \pm 0.0023	0.9973 \pm 0.0021	0.9955 \pm 0.0028

 Table 3. Average leave one out (LOO) score of the 5, f^* models on \mathcal{S}_T across different datasets.

E.2. Classifiers

In all experiments we set the random seed to 42 to ensure reproducibility and consistency across all runs. All experiments were conducted on a laptop computer with the following hardware specifications: 2.6-GHz 6-Core Intel Core i7 processor, 16 GB of 2400-MHz DDR4 RAM, and an Intel UHD Graphics 630 graphics card with 1536 MB of memory. Below are supplementary details about the classification models used. Below are supplementary details about the classification models used.

E.2.1. THE f^* MODEL

The function f^* served as the ground truth labeler, assessing whether the agent’s modifications led to a successful improvement. We evaluated five standard machine learning binary classification models, each achieving near 100% accuracy when trained and tested on \mathcal{S}_T (see Table 2). These models include two decision tree classifiers (DTC1 and DTC2), two random forest classifiers (RFC1 and RFC2), and a gradient boosting classifier (XGB). Descriptions of these models are provided below.

1. **Model f_1^* (DTC1):** A decision tree classifier with the following hyperparameters: criterion = “entropy”, min_samples_split = 2, min_samples_leaf = 1, and random_state = 42.
2. **Model f_2^* (DTC2):** A decision tree classifier with the following hyperparameters: criterion = “gini”, min_samples_split = 2, min_samples_leaf = 1, and random_state = 42.
3. **Model f_3^* (RFC1):** A random forest classifier with default settings and random_state = 42.
4. **Model f_4^* (RFC2):** A random forest classifier with the following hyperparameters: n_estimators = 500, min_samples_split = 2, min_samples_leaf = 1, max_features = “sqrt”, bootstrap = True, oob_score = True, and random_state = 42.
5. **Model f_5^* (XGB):** A gradient boosting classifier with the following hyperparameters: n_estimators = 500, max_depth = 50, learning_rate = 0.088, min_child_weight = 2, subsample=0.9, gamma = 0.088, and random_state = 42.

We define the ground truth labeler f^* either as a singular near-100% accuracy model (see Table 2) or as an agreement among multiple near-100% accuracy models.

The multi-defined f^* model. Although the five f^* models described above achieve nearly 100% accuracy when trained and tested on \mathcal{S}_T , we assessed their generalization using the leave-one-out (LOO) validation score. The observed differences in LOO validation scores (Table 3), despite similar and high accuracy scores (Table 2), highlight potential generalization gaps. To account for this, we employ a multi-defined f^* model to validate the experimental results.

For a given data point x , the five models: $f_1^*(x)$, $f_2^*(x)$, $f_3^*(x)$, $f_4^*(x)$ and $f_5^*(x)$ each make a prediction for the label of the data point. Based on these predictions, we define a boolean agreement mask $M(x)$ that checks whether all four models agree on the prediction:

$$M(x) = \mathbb{1}(f_1^*(x) = f_2^*(x) = f_3^*(x) = f_4^*(x) = f_5^*(x))$$

where $\mathbb{1}(\cdot)$ is the indicator function that outputs 1 if all four models agree, and 0 otherwise. Using this agreement mask, we define the ground truth labeling function $f^*(x)$ as follows:

$$f^*(x) = \begin{cases} f_1^*(x), & \text{if } M(x) = 1 \text{ (i.e., full agreement)} \\ 0, & \text{otherwise} \end{cases} \quad (8)$$

The singularly-defined f^* model. Alternatively, we define the labeling function $f^*(x)$ using a single near-100% accuracy model trained and tested on \mathcal{S}_T , selected from the set $\{f_1^*(x), f_2^*(x), f_3^*(x), f_4^*(x), f_5^*(x)\}$. Unless otherwise stated, all experimental results were obtained using the DTC2 model ($f_2^*(x)$) as the designated singularly-defined $f^*(x)$ function.

E.2.2. THE DECISION-MAKER’S MODEL (h)

We trained two-layer neural networks, denoted as h functions, using PyTorch with Adam optimizer with a learning rate of 0.001 and a batch size of 64. These h functions generate decisions for the test set agents. In cases where the test agent receives a negative classification, they can, if within budget, improve their feature values to get the desired classification from the h function. Table 4 summarizes the performance metrics of the f^* and h model functions, demonstrating their varied performance across the datasets.

Since the empirical setup evaluates the impact of improvement on h ’s error drop rates, we vary the loss functions we train the model h function with. We use the standard binary cross entropy loss (BCE) and the risk-averse weighted-BCE (wBCE) loss functions defined in Equation 5. In particular, because only negatively classified test-set agents improve, improvement (x'), if successful (that is, $f^*(x') = h(x') = 1$) reduces the false negative rate and turns true negatives into true positives. On the other hand, when unsuccessful (that is, $f^*(x') = 0$ and $h(x') = 1$), it increases the false positive rate by turning true negatives into false positives and false negatives into false positives.

The model trained with the weighted-BCE loss corresponds to a more risk-averse classifier that penalizes the false positive (FP) errors more heavily than the false negative (FN) one, creating a more compact positive agreement region that ensures more successful improvements. We prioritize minimizing FPs by ensuring the false positive to false negative weight ratio $\frac{w_{FP}}{w_{FN}} > 1$ is high, for example, $\frac{4.4}{0.001} = 4400$ for the adult dataset. Another form of risk-averse classification we consider is only classifying an agent as positive *iff* the probability of being positive is high. That is to say, we use the standard threshold 0.5 for the standard classifier and a higher threshold 0.9 for a more risk-averse classifier.

Dataset	Model (kind)	Accuracy	Precision	Recall	F1 Score
Adult	2-layer neural network ($h(x)$)	0.841087	0.699811	0.647677	0.672736
Law	2-layer neural network ($h(x)$)	0.901282	0.911691	0.984731	0.946805
OULAD	2-layer neural network ($h(x)$)	0.678003	0.698609	0.919853	0.794109
Synthetic	2-layer neural network ($h(x)$)	0.994021	1.000000	0.988473	0.994203

Table 4. Average performance of the standard models functions h across different datasets’ test sets $\mathcal{S}_{\text{test}}$ when test-set agent cannot improve that is, $r = 0$.

E.3. Agents Improvement

Given the feature vector of a negatively classified test-set agent, x_{orig} and its negative label $h(x_{\text{orig}})$, the loss function \mathcal{L} (BCE or wBCE), improvement budget r , step size α , number of iterations T and set of indices of improvement features S , compute the agent's improvement features. For each dataset, we predefine the improvable features that the agents can change in order to get a desirable (positive) model outcome (see Table 1). We vary the improvement budget in the empirical setup to so as to assess the impact of improvement on the the error drop rates.

Below are the steps of the improvement algorithm we used to compute each agent's improvement features.

Initialization:

$$x'_{(0)} = x_{\text{orig}}$$

Iterative updates: For $t = 0, 1, \dots, T - 1$:

1. Compute the gradient of the loss \mathcal{L} with respect to the agent's updates $x'_{(t)}$:

$$\mathbf{g}_{(t)} = \nabla_{x'_{(t)}} \mathcal{L}(h(x'_{(t)}), h(x_{\text{orig}}))$$

2. Update the improvement features by taking a step in the direction of the sign of the gradient:

$$\rho_{(t)}[i] = \begin{cases} \alpha \cdot \text{sign}(\mathbf{g}_{(t)}[i]), & \text{if } i \in S \\ 0, & \text{otherwise} \end{cases}, \quad \forall i \in [d]$$

$$x'_{(t+1)} = x'_{(t)} + \rho_{(t)}$$

3. Project the updated improvement features back onto the r -ball around the original features x_{orig} :

$$x'_{(t+1)} = (x_{\text{orig}} + \text{clip}_{[-r, r]}(x'_{(t+1)} - x_{\text{orig}}))$$

Improvement vector: After T iterations, the final agent's improvement is given by:

$$x' = x'_{(T)}$$

E.4. Evaluation Results: Supplementary Details

Following the key insights mentioned in the main paper in Section 6, below are the detailed observations from the experimental evaluation and the supplementary figures of the experimental results.

Effect of dataset characteristics. For all datasets we consider, Figure 2 shows that as the improvement budget increases, error rates drop significantly, particularly when agents improve in response to a risk-averse model trained using the $\mathcal{L}_{\text{wBCE}}$ loss function. Dataset characteristics notably influence performance. For instance, as shown in Figure 6, the Law school and Synthetic datasets exhibit the highest separability and require relatively less risk aversion to achieve substantial and close to zero error reductions. Among all datasets, the Synthetic dataset shows a sharper error decline, reaching zero as the improvement budget increases (refer to Figure 2). Furthermore, as depicted in Figure 7, the Adult and Synthetic datasets demonstrate the lowest k -NN test-set false positive rates (FPRs) of 0.0678 and 0.000, respectively, compared to the OULAD and Law School datasets. As seen in Figure 2, for both datasets, the error drops close to 0 as r increases.

Effect of risk-aversion and improvement budget. Figures 10a, 10c, 10e and 10g show that when agents improve to a risk-averse model trained using $\mathcal{L}_{\text{wBCE}}$ with $\frac{w_{\text{FP}}}{w_{\text{FN}}} > 1$, the error rate decreases rapidly, particularly as the improvement budget increases. Notably, the higher the false positive to false negative weight ratio $\frac{w_{\text{FP}}}{w_{\text{FN}}}$, the faster the reduction in error (see Figure 8). In contrast, models trained with the standard \mathcal{L}_{BCE} loss function exhibit a slower error reduction rate, almost looking like a line, under the same conditions as the effects of improvement are minimal and cancel each other out (see Figure 9).

Furthermore, the false negative rate (FNR) decreases as the improvement budget r grows when the agents respond (improve) to an $\mathcal{L}_{\text{wBCE}}$ -trained model (see Figures 9b, 9d, and 9f). This is because almost 100% of the false negatively classified agents improve to become true positives as shown in Figures 9a, 9c, and 9e. On the other hand, on datasets where the weighted-BCE loss function effectively removed false positives (close to 0 false positive rate (FPR)) before agents' improvement, remains very low, in some cases close to 0 (e.g., in Figure 9f), since no or few agents become false positives after improvement (see Figure 9e). However, when the weighted-BCE loss function wasn't as effective, the false positive rate increases as the improvement budget $r < 2.0$ grows (see Figure 9d).

Although we observe similar trends when the threshold for classifying an agent as positive increases to 0.9 (instead of 0.5), the error is slightly higher and the reduction becomes slower (Figure 10). Additionally, while Figures 2 and 8 demonstrate diminishing returns for $r > 2.0$, a different trend emerges in Figures 10b, 10d, 10f and 10h where we classify agents positive with high probability (0.9).

Effect of choice of f^* model. Although different f^* models achieved $\sim 100\%$ accuracy on a given dataset while exhibiting varied LOO accuracy (cf. Tables 2 and 3), the error drop rate showed consistent patterns when different f^* models are used. As shown in Figure 11, although some f^* models, such as RFC2 (f_4^*) (cf. Figure 11c), had higher performance gains than others, in all cases, the error drops rapidly as improvement budget increases and agents respond (improve) to wBCE-trained models.

Additionally, performance gains observed with evaluation of successful improvement using the multi-defined f^* (cf. Figure 12) were quite similar in trend and gains to those when a singularly-defined model function $f^* = f_2^*$ was used (cf. Figure 10, column one ((a), (c), (e) and (g))).

Our results indicate that while the f^* models achieve $\sim 100\%$ accuracy when trained and tested on the unsplit dataset S_T , they often overfit, as shown by the LOO scores (cf. Table 3). Nevertheless, they yield comparable performance gains when assessing agents' improvement (cf. Figures 12 and 11).

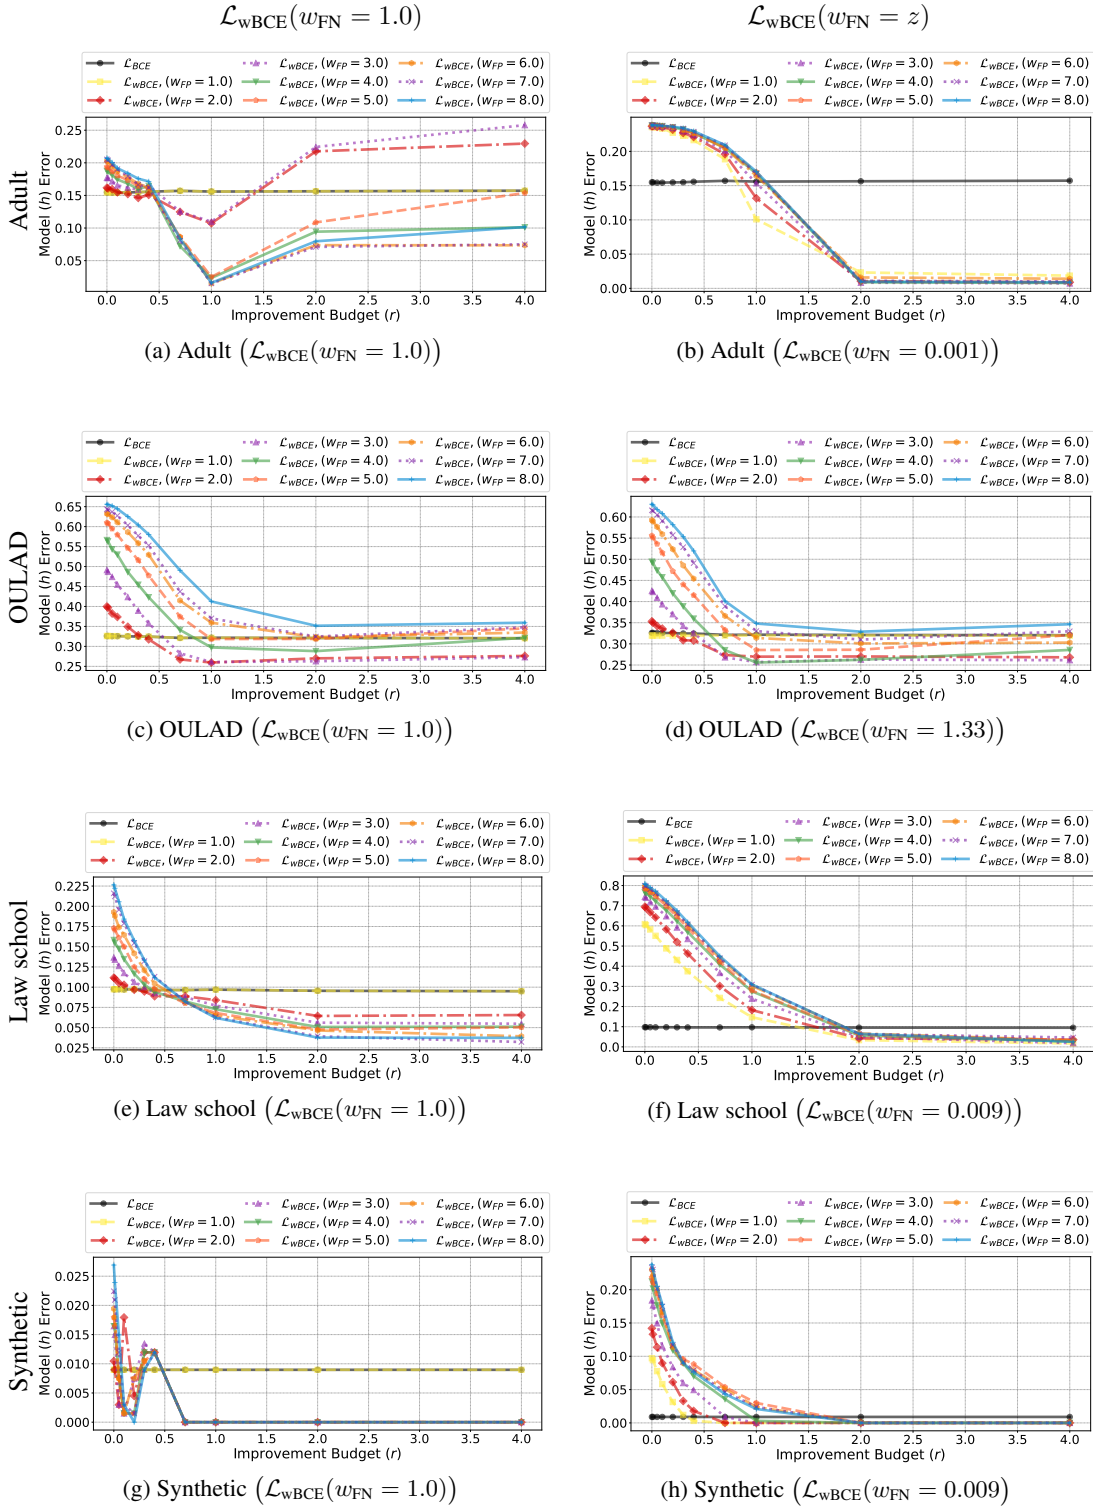





Figure 8. Comparison of the error drop rate when agents improve to the risk-averse models trained with \mathcal{L}_{WBCE} where $w_{FN} = 1.0$ ((a), (c), (e), and (g)) and where w_{FN} is optimized and false positive weight is varied $w_{FP} = \{i\}_{i=1}^8$ across different datasets (Adult, OULAD, Law school and Synthetic). The standard model (black line) trained with \mathcal{L}_{BCE} loss function is such that $w_{FP} = w_{FN} = 1$, and in all cases an agent is classified as positive if the probability of being positive is above 0.5.




(a) Adult (Movement of agents from TN/FN to TP/FP)




(b) Adult (FNR/FPR before and after agents' improvement)




(c) OULAD (Movement of agents from TN/FN to TP/FP)



(d) OULAD (FNR/FPR before and after agents' improvement)



(e) Synthetic (Movement of agents from TN/FN to TP/FP)



(f) Synthetic (FNR/FPR before and after agents' improvement)

Figure 9. ((a), (c), and (e)) The percentage of negatively classified agents (TN and FN) that transition to TP and FP after responding to the classifier ($h(x)$). ((b), (d), and (f)) The FPR and FNR before and after agents move. On the adult dataset, the wBCE-trained model used $w_{FP} = 0.001$ and $w_{FN} = 4.4$, while on the OULAD dataset, it used $w_{FP} = 1.33$ and $w_{FN} = 2$, and on the synthetic dataset, $w_{FP} = 0.009$ and $w_{FN} = 1.0$. In all cases, BCE-trained models used $w_{FP} = w_{FN} = 1$, and in all cases an agent is classified as positive if the probability of being positive is above 0.5.

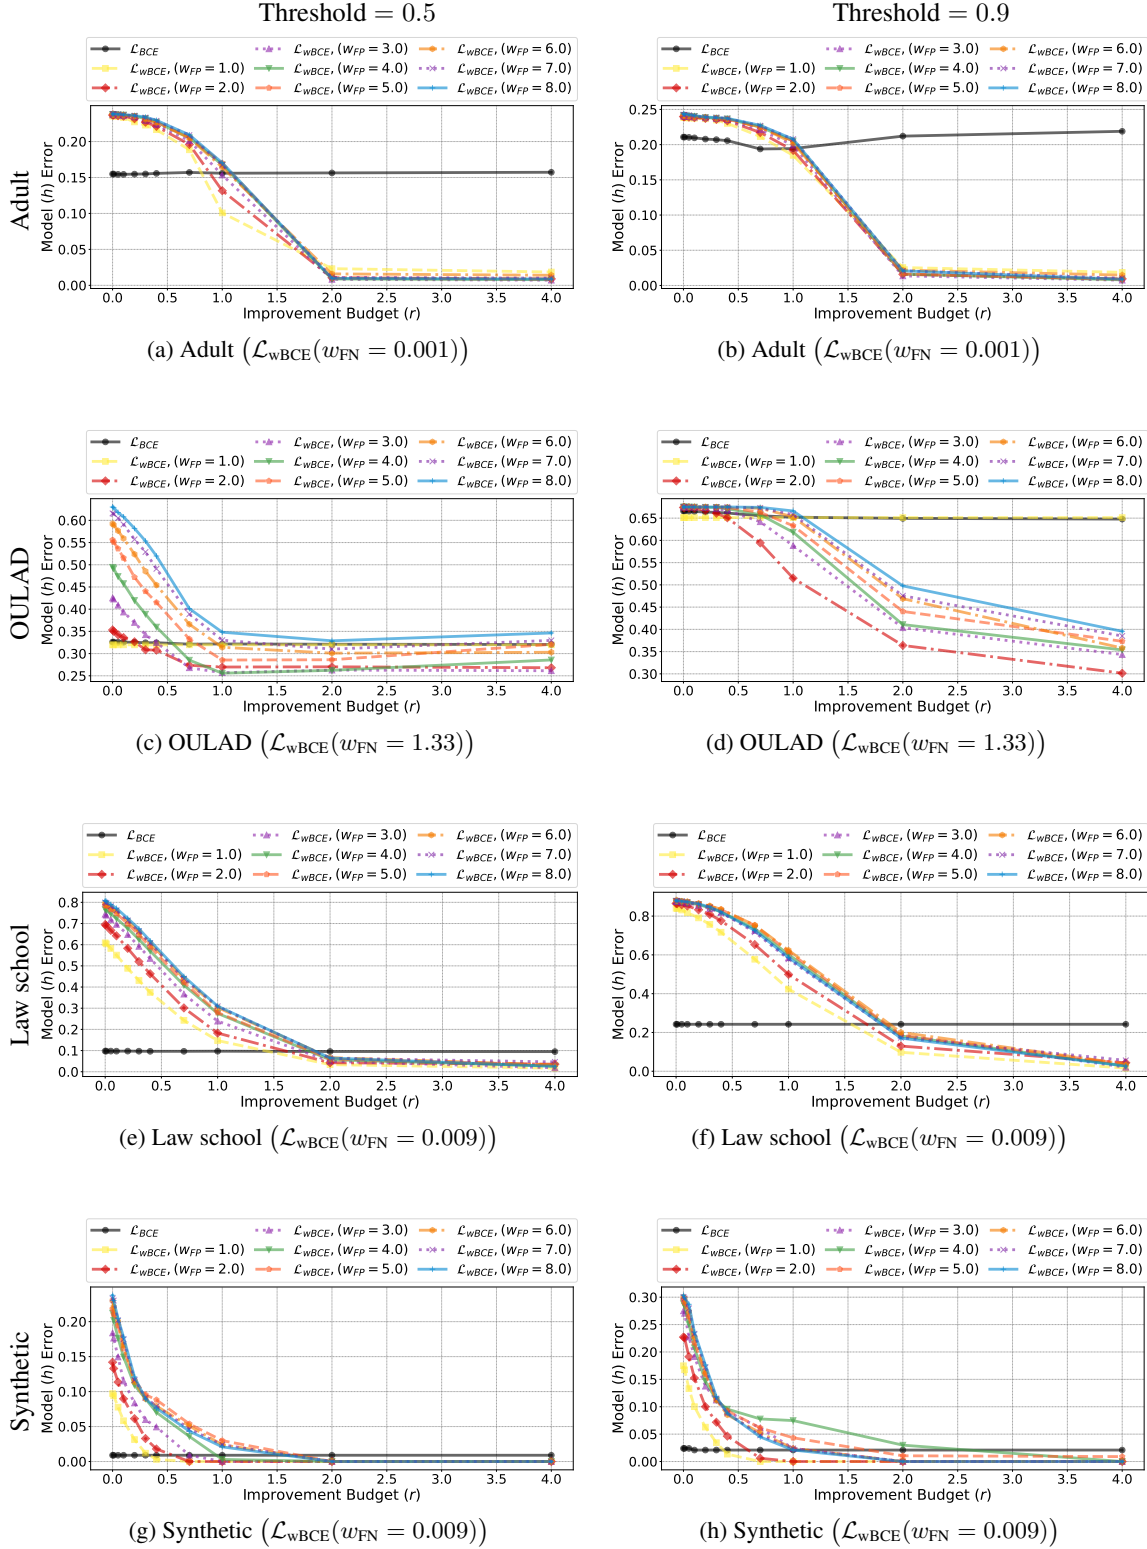


Figure 10. Comparison of the error drop rate when agents improve to the risk-averse ($\mathcal{L}_{wBCE}, \frac{w_{FP}}{w_{FN}} > 1, w_{FP} = \{i\}_{i=1}^8$) and standard ($\mathcal{L}_{BCE}, w_{FP} = w_{FN} = 1$) models across four datasets (Adult, OULAD, Law school, and Synthetic). **Column one** ((a), (c), (e) and (g)) considers models where an agent is classified as positive if the probability of being positive is above 0.5 and **column two** ((b), (d), (f) and (h)) considers models where a higher threshold is used 0.9. Increasing the improvement budget and classifier risk-aversion (high $\frac{w_{FP}}{w_{FN}}$) leads to a sharper error drop rate, and loss-based risk aversion is more effective than the threshold-based risk-aversion.

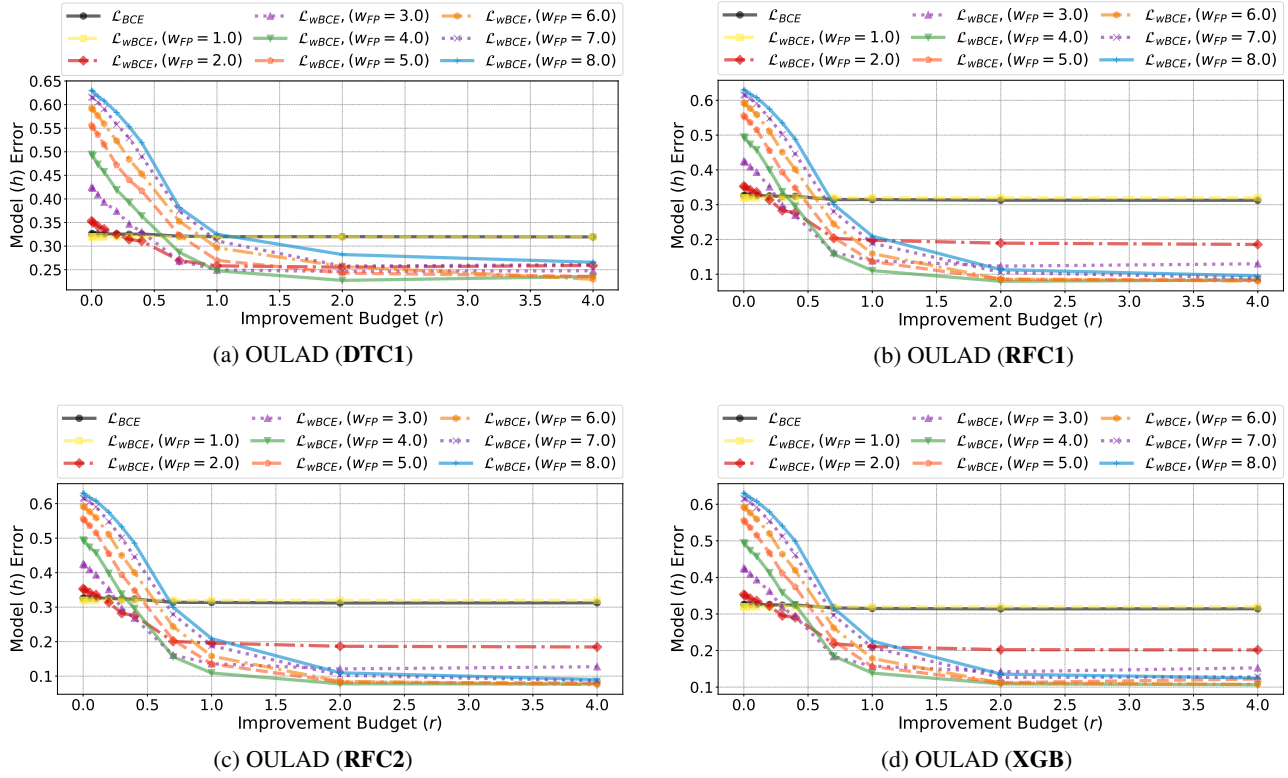


Figure 11. Risk-averse (\mathcal{L}_{WBCE} with $w_{FP} = \{i\}_{i=1}^8, w_{FN} = 1.33$) and standard ($\mathcal{L}_{BCE}, w_{FP} = w_{FN} = 1$) trained model function (h) error drop rates versus improvement budget (r) on the Adult dataset where different **singularly-defined** f^* are used to verify successful-ness of improvement: ((a)) with f_1^* , ((b)) with f_3^* , ((c)) with f_4^* , and ((d)) with f_5^* . For DTC2, f_2^* see Figure 10c. In all cases, the threshold for classifying an agent as positive is 0.5.

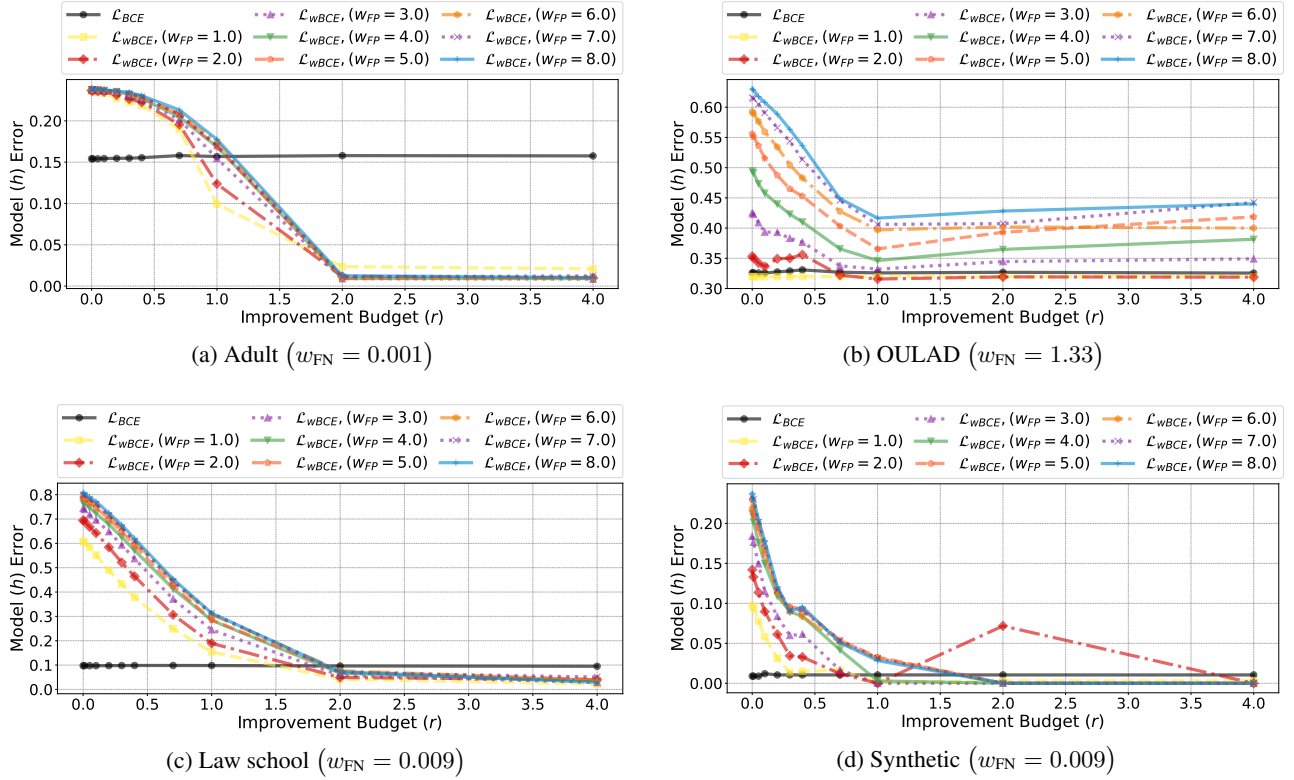


Figure 12. Risk-averse (\mathcal{L}_{WBCE} with $w_{FP} = \{i\}_{i=1}^8$) and standard (\mathcal{L}_{BCE}) trained model function (h) error versus improvement budget (r) across four datasets (Adult, OULAD, Law school, and Synthetic). For all cases, the threshold for classifying an agent as positive is 0.5 and a **multi-defined** f^* model is used to verify successful-ness of improvement. Increasing the improvement budget and classifier risk-aversion (high $\frac{w_{FP}}{w_{FN}}$) leads to a faster error drop rate.