

Exponential Savings in Agnostic Active Learning through Abstention

Nikita Puchkin

HSE University and Institute for Information Transmission Problems RAS, Moscow

NPUCHKIN@HSE.RU

Nikita Zhivotovskiy

ETH, Zürich

NIKITA.ZHIVOTOVSKII@MATH.ETHZ.CH

Editors: Mikhail Belkin and Samory Kpotufe

Abstract

We show that in pool-based active classification without assumptions on the underlying distribution, if the learner is given the power to abstain from some predictions by paying the price marginally smaller than the average loss $1/2$ of a random guess, exponential savings in the number of label requests are possible whenever they are possible in the corresponding realizable problem. We extend this result to provide a necessary and sufficient condition for exponential savings in pool-based active classification under the model misspecification.

Keywords: active learning, sample complexity, abstention, reject option, Chow’s risk, VC dimension, model selection aggregation, Massart’s noise

1. Introduction

Pool-based *active classification* can be seen as an extension of the classical PAC classification setup, where instead of learning from the labeled sample $(X_1, Y_1), \dots, (X_n, Y_n)$, one can adaptively request the labels from a large pool X_1, X_2, \dots of i.i.d. unlabeled instances round by round. Our hope is to request significantly fewer labels Y_i and get the same statistical guarantees as in *passive learning*. A textbook example is the one of learning the class \mathcal{F} of threshold classifiers in the realizable (noise-free) case where a binary search based algorithm can improve the sample complexity (the number of requested labels) from the passive sample complexity $O(\frac{1}{\varepsilon})$ to the exponentially better $O(\log \frac{1}{\varepsilon})$ sample complexity, where ε is the desired probability of error. For more general classes the realizable case sample complexity is understood quite well in the distribution dependent (Dasgupta, 2005) and the minimax (Hanneke and Yang, 2015) senses.

The improvements in active learning are less impressive once the problem is not realizable. Our starting point is the foundational work of Kääriäinen (2006) containing the following observations formulated (informally) as follows:

1. (*Arbitrary noise*) Active learning cannot bypass the classical agnostic passive learning bound on the sample complexity $\Omega(\frac{1}{\varepsilon^2})$ (Vapnik and Chervonenkis, 1974) in the noisy case¹. Indeed, if there is just one “heavy” instance X with the noisy label $Y \in \{0, 1\}$ such that $\Pr(Y = 1|X) = \frac{1}{2} \pm \varepsilon$, the sample complexity of any active learning algorithm can be reduced to the estimation of $\Pr(Y = 1|X)$. A well-known result (Anthony and Bartlett, 2009, Lemma 5.1) gives the lower bound $\Omega(\frac{1}{\varepsilon^2})$.

1. When only the dependence on ε is considered.

2. (*Misspecification*) No improvements in active learning over passive learning are possible in the misspecified case even if there is no noise in the labeling mechanism. That is, even when $Y = f^*(X)$ almost surely for some $f^* \notin \mathcal{F}$, one cannot bypass the passive learning lower bound $\Omega(\frac{1}{\varepsilon^2})$. To demonstrate this phenomenon a class \mathcal{F} consisting of only two specific functions is sufficient.
3. (*Bounded noise*) The sample complexity of active classification in the bounded noise case, that is, when each label of the true function is corrupted independently with probability strictly smaller than 0.5, is essentially the same as in the realizable case whenever the Bayes optimal rule is in the class. Therefore, in this case, exponential savings are also possible and well understood by now.

To avoid the aforementioned lower bounds of [Kääriäinen](#) and to show significant superiority of active learning, many authors are focusing on the favorable *noise assumptions*. These assumptions take their roots in passive learning and include the realizable or the bounded noise cases ([Massart and Nédélec, 2006](#)), Tsybakov’s noise ([Tsybakov, 2004](#)), and the Bernstein condition ([Bartlett and Mendelson, 2006](#)). Unfortunately, these assumptions are difficult to satisfy in practice, as they require that the Bayes optimal classifier is in the class. The Bernstein condition avoids this problem, but little is known about the cases where it holds in classification without assuming that the Bayes optimal rule is in the class (see ([Gelbart and El-Yaniv, 2019](#))). Based on his findings, [Kääriäinen](#) writes: “*The implication of this lower bound is that exponential savings should not be expected in realistic models of active learning, and thus the label complexity goals in active learning should be refined*”.

This paper aims to provide such a refinement. Our method will be as follows: instead of restricting the distribution of (X, Y) , we use the power to abstain from some predictions. To do so, we allow the learner to output a $\{0, 1, *\}$ -valued classifier, where $*$ corresponds to the reject option. For $p \in [0, \frac{1}{2}]$ and a $\{0, 1, *\}$ -valued classifier f we define the Chow’s risk ([Chow, 1970](#)) as

$$R^p(f) = \Pr(f(X) \neq Y \text{ and } f(X) \in \{0, 1\}) + \left(\frac{1}{2} - p\right) \Pr(f(X) = *), \quad (1)$$

which is the binary risk as long as we predict in $\{0, 1\}$, and the price of abstention is equal to $\frac{1}{2} - p$. The special case $p = 0$ corresponds to the situation where the price of abstention is the same as the average loss of a random guess. In what follows, we always think of p as a small parameter, so that the price of $\frac{1}{2} - p = 0.49$ for abstention will always suffice.

In the standard active learning setup, given a class \mathcal{F} of $\{0, 1\}$ -valued classifiers and using only a small number of label requests, we aim to construct a classifier f such that, with high probability,

$$R(\tilde{f}) - \inf_{f \in \mathcal{F}} R(f) \leq \varepsilon,$$

where the binary risk $R(f)$ is defined as $R(f) = \Pr(f(X) \neq Y)$. Our simplified aim is instead to construct for the same class \mathcal{F} a $\{0, 1, *\}$ -classifier \hat{f} such that, with high probability,

$$R^p(\hat{f}) - \inf_{f \in \mathcal{F}} R(f) \leq \varepsilon, \quad (2)$$

that is, together with a reject option, \hat{f} predicts almost as good as the best classifier in the class. This approach has a practical motivation: if the learner can identify the instances where the output

classifier predicts no better than a random guess, it is reasonable to use some external source of information (for example, expert advice) to make a prediction. It will be clear that our setup is essentially an active learning version of the *model selection aggregation* problem (Tsybakov, 2003); in the model selection aggregation one is allowed to output an *improper* (not necessarily in \mathcal{F}) classifier and use the “curvature” of the loss function to predict as good as the best classifier in \mathcal{F} . It has been recently shown in (Bousquet and Zhivotovskiy, 2019) that Chow’s risk (1) gives exactly the right amount of “curvature” to exploit the techniques in the model selection aggregation and improve the agnostic sample complexity in passive learning. Another closely related setup is the prediction of individual sequences with expert advice where the “curvature” of the loss expressed in terms of *mixability* gives significant improvements with no assumptions on the data generating mechanism (Vovk, 1990; Haussler et al., 1998) (see also (Cesa-Bianchi and Lugosi, 2006, Chapter 3)).

In what follows, we assume that the price of abstention is only marginally smaller than the average loss of a random guess; that is, for instance, $\frac{1}{2} - p = 0.49$. With such an option, we have to be conservative when to abstain: our algorithm should be adaptive to the bounded noise case and avoid the reject option in this situation. Indeed, if the distribution is such that the true labels are corrupted independently with a probability of only 0.25, any abstention with the price of 0.49 can worsen the situation. We show in Proposition 3.6 that our algorithm is adaptive to the bounded noise assumption and abstains rarely if this assumption holds; thus, we are always recovering the standard guarantees.

We are ready to make an informal statement of our main result. We use the notions of the *VC dimension* and the *disagreement coefficient*, both of which are standard in the active learning literature. Formal definitions are presented in Section 2.

Theorem 1.1 (A simplified statement) *Fix $\varepsilon, \delta \in (0, 1]$ and $p \in (0, \frac{1}{2}]$. There is an active learning algorithm such that for any distribution P of (X, Y) , after requesting*

$$n = O\left(\frac{d \theta(\varepsilon/p)}{p^2} \log^2\left(\frac{d}{p\varepsilon\delta}\right)\right), \quad (3)$$

*labels it produces a $\{0, 1, *\}$ -valued classifier \hat{f}_p satisfying, with probability at least $1 - \delta$,*

$$R^P(\hat{f}_p) - \inf_{f \in \mathcal{F}} R(f) \leq \varepsilon.$$

Here $\theta(\cdot)$ is the disagreement coefficient, d is the VC dimension of a $\{0, 1\}$ -valued class \mathcal{F} .

A formal statement of this result is Theorem 3.1. Observe that if the disagreement coefficient $\theta(\cdot)$ is bounded, which holds, for example, for threshold classifiers on the real line and homogeneous linear separators in \mathbb{R}^d under a uniform distribution (Hanneke, 2014), then exponential savings are always possible by the above result. Indeed, in this case, the sample complexity bound (3) scales as $O(\log^2(\frac{1}{\varepsilon}))$. Also, the definition of $\theta(\cdot)$ implies that the dependence on ε in (3) is never worse than $O(\frac{1}{\varepsilon} \log^2(\frac{1}{\varepsilon}))$. This is superior to the passive $\Theta(\frac{1}{\varepsilon^2})$ sample complexity. To be more specific, we illustrate our result by the following basic example.

Example 1.2 *For threshold classifiers on the real line, if the price of abstention is $1/2 - p = 0.49$, our result implies that for any distribution of the data, $O(\log^2(\frac{1}{\varepsilon}))$ label requests are sufficient to guarantee that $R^P(\hat{f}_p) - \inf_{f \in \mathcal{F}} R(f) \leq \varepsilon$. If either abstention is not allowed or $p = 0$, the number of label requests $\Theta(\frac{1}{\varepsilon^2})$ cannot be generally improved in the active learning setup.*

The reader can recall that a similar sample complexity bound holds in the bounded noise model of [Massart and Nédélec \(2006\)](#), that is, when the *Bayes optimal classifier* f_B^* belongs to \mathcal{F} and $|2\Pr(Y = 1|X) - 1| \geq h > 0$ almost surely (see, for example, [Hanneke and Yang, 2015](#), Section 7.1)). And at least on an intuitive level, when $f_B^* \in \mathcal{F}$, an option to abstain can be potentially used to eliminate the noise and reduce the problem to the bounded noise case. More importantly, our result is also robust to the model misspecification, that is, we allow $f_B^* \notin \mathcal{F}$. Indeed, according to [Kääriäinen \(2006\)](#), the model misspecification alone can result in the $\Omega(\frac{1}{\varepsilon^2})$ lower bound even if there is no noise in the labeling mechanism. Therefore, our estimator with a reject option avoids *both known reasons* of $\Omega(\frac{1}{\varepsilon^2})$ lower bounds.

Our second result is the minimax analysis of the standard active learning setup. We exploit our classifier with a reject option as an intermediate step. By the minimax analysis, we usually mean the sample complexity bounds valid for any marginal distribution (denoted by P_X) of the unlabeled data. In this setup, an aforementioned lower bound of [Kääriäinen \(2006\)](#) implies that one should restrict the noise of the problem to get exponential savings, that is, we assume

$$|2\Pr(Y = 1|X) - 1| \geq h \text{ almost surely for some } h > 0. \tag{4}$$

We answer the following question.

Assuming Massart’s noise (4) what is the characterization of \mathcal{F} allowing exponential savings in active learning?

Under a strong assumption that the Bayes optimal classifier f_B^* is in \mathcal{F} , this question has been answered in [\(Hanneke and Yang, 2015, Theorem 4\)](#): exponential savings are possible under (4) and $f_B^* \in \mathcal{F}$ if and only if the *star number* s is finite (defined in Section 2). We need to define the *diameter* of \mathcal{F} . It is the smallest integer D such that

$$\sup_{f,g \in \mathcal{F}} |\{x \in \mathcal{X} : f(x) \neq g(x)\}| \leq D,$$

where \mathcal{X} is our instance space (see Section 2).

Theorem 1.3 (An informal statement) *Exponential savings are possible in active classification for any distribution satisfying Massart’s noise assumption (4) (without assuming $f_B^* \in \mathcal{F}$) if and only if both the star number s (or respectively the disagreement coefficient $\theta(\cdot)$) if the dependence on P_X is allowed) and the diameter D are finite.*

A formal version of this result is Theorem 4.1. As we mentioned, one may show that Massart’s assumption (4) is also inevitable if one wants to have exponential savings for any marginal distribution P_X of the unlabeled data. We discuss this in more detail in Section 4.

The disagreement coefficient and the star number of [Hanneke \(2007\)](#); [Hanneke and Yang \(2015\)](#) play an important role in the active learning literature while the diameter of [Ben-David and Uner \(2014\)](#) is used in the analysis of passive learning with *deterministic labeling*, that is, when $Y = f_B^*(X)$ almost surely. It appears that both the star number and the diameter are infinite in many natural scenarios. Fortunately, our main result shows that active learning is useful if the disagreement coefficient $\theta(\cdot)$ is small and the price of abstention is separated from $1/2$.

1.1. Our contributions

- In Section 3, we present our main result as well as the performance bound for a passive algorithm called the *mid-point algorithm*.
- In Section 3.2, we show the adaptivity of our results to the bounded noise assumption. In particular, our algorithm abstains rarely when this assumption holds.
- In Section 4, we return to the standard active learning setup where the reject option is not available. We characterize the case where the Bayes optimal rule is not in the class, the Bernstein assumption is vacuous, but exponential savings are still possible.

1.2. Related work

We start with a concise literature overview followed by a more detailed comparison with some related recent results.

The most standard algorithm in the realizable case is referred to as the CAL algorithm (after the names of [Cohn, Atlas, and Ladner \(1994\)](#)). This algorithm can be shown to provide exponential savings in some cases. The analysis of the realizable case with the complexity measure depending on the marginal distribution of the unlabeled data is by [Dasgupta \(2005\)](#). In particular, [Dasgupta](#) generalizes various examples of exponential savings in realizable active learning. The fact that exponential savings are also possible in the bounded noise case for the threshold functions is attributed to [Burnashev and Zigangirov \(1974\)](#); their ideas were later developed in ([Korostelev, 1999](#); [Golubev and Levit, 2003](#); [Castro and Nowak, 2008](#)). The first general agnostic active learning algorithm is presented in ([Balcan, Beygelzimer, and Langford, 2009](#)) followed by a more refined analysis in ([Hanneke, 2007](#); [Dasgupta et al., 2008](#); [Beygelzimer et al., 2009](#); [Hsu, 2010](#); [Koltchinskii, 2010](#); [Hanneke, 2011](#); [Zhang and Chaudhuri, 2014](#); [Hanneke and Yang, 2015](#)) and other works. Most of the known upper bounds are based on the *disagreement coefficient* introduced by [Hanneke \(2007\)](#) to analyze the performance of active learning algorithms; essentially the same quantity also appeared in the analysis of ratio-type empirical processes ([Alexander, 1987](#)) and was later reintroduced to the passive learning literature by [Giné and Koltchinskii \(2006\)](#). We refer to the survey ([Hanneke, 2014](#)) for a detailed exposition of these results.

The risk (1) was analyzed in the seminal work of [Chow \(1970\)](#). The statistical analysis in the context of passive learning was first provided in ([Herbei and Wegkamp, 2006](#); [Bartlett and Wegkamp, 2008](#)). The authors consider the reject option as an action available not only to the learner but also to the classifiers in the base class. For a more extensive survey and some related results, we refer to ([Freund et al., 2004](#); [El-Yaniv and Wiener, 2010](#); [Cortes et al., 2016](#); [Yan et al., 2016](#)), and ([Gelbhart and El-Yaniv, 2019](#)). Recent works ([Bousquet and Zhivotovskiy, 2019](#); [Neu and Zhivotovskiy, 2020](#)) show that if the learner is given an option to abstain, and the risk of Chow (1) is used, then the so-called *fast rates* are possible without additional assumptions in passive and online classification. There is also a line of research devoted to active learning in the non-parametric setup not covered in this paper ([Castro and Nowak, 2008](#); [Koltchinskii, 2010](#); [Minsker, 2012](#); [Locatelli et al., 2017, 2018](#)). An extension of our results to the non-parametric setup is the natural direction of future work.

In the context of active learning, Chow’s risk has been recently analyzed in ([Shekhar et al., 2020](#)). The authors consider a non-parametric classification problem and make some margin-type assumptions, which is different from our setup. Chow’s risk is also connected to surrogate losses

appearing in the context of active learning in (Hanneke and Yang, 2019), where the main purpose of using these losses is to simplify the computational problems associated with minimizing the binary loss. In particular, their statistical results are always not better than for the binary loss, which is again different from our findings.

Relations to (Bousquet and Zhivotovskiy, 2019). The model we are considering has been recently considered in the context of passive learning. Bousquet and Zhivotovskiy show that the passive learning sample complexity $\Theta(\frac{1}{\varepsilon^2})$ can be improved to $O(\frac{1}{p\varepsilon} \log \frac{1}{\varepsilon})$ whenever (2) is used.

We use an improved version of their argument as a subroutine and provide a simplified analysis for it. It appears that in the context of active learning just being able to provide a $O(\frac{1}{p\varepsilon} \log \frac{1}{\varepsilon})$ sample complexity is not sufficient. We instead directly exploit a phenomenon first observed by Audibert (2008), which in our case can be described as follows: for some realizations of the labels the difference $R^p(\hat{f}) - R(f^*)$ can be negative and (2) immediately follows. One of our key technical observations is that there is a way to detect this event using only the learning sample.

Relations to selective classification. In selective classification, a learner aims to provide a pair of $\{0, 1\}$ -valued functions (\hat{f}, \hat{g}) called a selective classifier. The classifier $\hat{f}(x)$ must be pointwise competitive, that is, $\hat{f}(x) = f^*(x)$, where $f^* = \arg \min_{f \in \mathcal{F}} R(f)$ if and only if $\hat{g}(x) = 1$. If $\hat{g}(x) = 0$ the selective classifier abstains. The learner is interested in minimizing the rejection rate $\Pr(\hat{g}(X) = 0)$. The pointwise competitiveness requirement is quite restrictive and is not always desirable, especially for those instances where f^* predicts differently from the Bayes optimal classifier f_B^* . To the best of our knowledge, only the realizable case (El-Yaniv and Wiener, 2010, 2012) and the case of a small $R(f^*)$ (Gelbhart and El-Yaniv, 2019) were considered so far. In this paper, we are pursuing a less ambitious goal and allow our classifier to make a small portion of mistakes when it does not abstain. As a result, our improvements are somewhat more substantial.

Relations to the minimax analysis of Hanneke and Yang (2015). The work of Hanneke and Yang (2015) provides an almost complete picture of the sample complexity bounds in the minimax sense under the bounded noise and the Bernstein assumption. It remains open if these savings are possible in other cases. We make one step forward and show that these improvements can be obtained for some classes even if the problem is misspecified, that is, the Bayes optimal rule is not in the class, and the Bernstein condition is vacuous.

Relations to confidence-rated predictors. Zhang and Chaudhuri (2014) analyze the disagreement based active learning algorithms via the confidence-rated predictor: instead of requesting a label of a specific point in the disagreement set, their algorithm can randomly abstain from doing so. Similar techniques were used in (Balcan et al., 2007; Balcan and Long, 2013) for linear separators. This approach leads to some improvements in the sample complexity bounds under various low noise assumptions. However, our analysis uses an option to abstain only when classifying some of the problematic instances.

2. Notation and Setup

We introduce some notation and basic definitions that will be used throughout the text. The symbol $\mathbb{1}[A]$ denotes an indicator function of the event A . The notation $f \lesssim g$ or $g \gtrsim f$ means that for some universal constant $c > 0$ we have $f \leq cg$. To avoid the problems with the logarithmic function we assume that $\log x$ means $\max\{\log x, 1\}$. Throughout the paper we also use the standard

$O(\cdot), \Omega(\cdot), \Theta(\cdot)$ notation. We set $a \wedge b = \min\{a, b\}$ and $a \vee b = \max\{a, b\}$. For $s \geq 1$ we define the $L_s(P)$ norm as $\|g\|_{L_s} = (\mathbb{E}|g(Z)|^s)^{\frac{1}{s}}$, where the expectation is taken with respect to some measure P always clear from the context. The $L_s(P)$ diameter of \mathcal{F} is

$$\mathcal{D}(\mathcal{F}, L_s) = \sup_{f, g \in \mathcal{F}} \|f - g\|_{L_s}.$$

We define the instance space \mathcal{X} and the label space $\mathcal{Y} = \{0, 1\}$. We assume that the set $\mathcal{X} \times \mathcal{Y}$ is equipped with some σ -algebra and a probability measure $P = P_{X, Y}$ on measurable subsets is defined. We also assume that we are given a set of classifiers \mathcal{F} mapping \mathcal{X} to \mathcal{Y} . In *passive learning* we observe $S_n = \{(X_1, Y_1), \dots, (X_n, Y_n)\}$ sampled according to P . In the *pool-based active learning*, we define an active learning algorithm as an algorithm taking as input a budget $n \in \mathbb{N}$, and proceeding as follows. The algorithm initially uses an unlabeled infinite data sequence X_1, X_2, \dots distributed according to P_X . The algorithm may select an index i_1 and request the label Y_{i_1} . In this case we observe the value of Y_{i_1} , sampled according to the conditional distribution $Y|X_{i_1}$; then based on both the unlabeled sequence and Y_{i_1} , it may select another index i_2 and request to observe Y_{i_2} . This continues for at most n rounds. Finally the algorithm outputs a classifier f .

Given $S_n = \{(X_1, Y_1), \dots, (X_n, Y_n)\}$ let P_{S_n} denote the expectation with respect to the empirical measure induced by this sample. We sometimes write $P_n f$ instead of $P_{S_n} f(X)$ and $P f$ instead of $\mathbb{E}f(X)$. For a set $\{x_1, \dots, x_k\} \subseteq \mathcal{X}$ and a class of $\{0, 1\}$ -valued functions \mathcal{F} , we denote the restriction of \mathcal{F} on $\{x_1, \dots, x_k\}$ by $\mathcal{F}_{\{x_1, \dots, x_k\}} = \{(f(x_1), \dots, f(x_k)) : f \in \mathcal{F}\}$. The value of the growth function $\mathcal{S}_{\mathcal{F}}(k)$ is defined as the largest cardinality of $\mathcal{F}_{\{x_1, \dots, x_k\}}$ among all $x_1, \dots, x_k \in \mathcal{X}$. The VC dimension of \mathcal{F} is the largest integer d such that $\mathcal{S}_{\mathcal{F}}(d) = 2^d$ (Vapnik and Chervonenkis, 1968). For any set \mathcal{F} of classifiers let the disagreement set of \mathcal{F} be defined as

$$\text{DIS}(\mathcal{F}) = \{x \in \mathcal{X} : \text{there are } f, g \in \mathcal{F} \text{ such that } f(x) \neq g(x)\}.$$

As above, define the prediction risk as $R(f) = \Pr(f(X) \neq Y)$ and the Chow's risk R^p risk is given by (1). The Bayes optimal rule f_B^* and the best classifier in the class f^* are given by

$$f_B^*(x) = \mathbb{1}[\Pr(Y = 1|X = x) \geq 1/2] \quad \text{and} \quad f^* = \arg \min_{f \in \mathcal{F}} R(f).$$

The largest $h \geq 0$ such that almost surely

$$|2 \Pr(Y = 1|X) - 1| \geq h \tag{5}$$

is called the Massart's margin parameter (Massart and Nédélec, 2006). Let

$$R_{S_n}(f) = \frac{1}{n} \sum_{(X_i, Y_i) \in S_n} \mathbb{1}[f(X_i) \neq Y_i]$$

denote the empirical risk with respect to S_n . We sometimes write $R_n(f)$ instead of $R_{S_n}(f)$ when the sample is clear from the context. Any minimizer of the empirical risk $R_{S_n}(f)$ in \mathcal{F} is called ERM. For a $\{0, 1, *\}$ -valued classifier g we define the empirical Chow's risk as

$$R_{S_n}^p(g) = \frac{1}{n} \sum_{(X_i, Y_i) \in S_n} \mathbb{1}[g(X_i) \neq Y_i \text{ and } g(X_i) \in \{0, 1\}] + \frac{1/2 - p}{n} \sum_{(X_i, Y_i) \in S_n} \mathbb{1}[g(X_i) = *].$$

Fix $\varepsilon \geq 0$. The disagreement coefficient $\theta(\cdot)$ of [Hanneke \(2007\)](#) is defined as

$$\theta(\varepsilon) = \sup_{g \in \mathcal{F}, \varepsilon_0 \geq \varepsilon} \frac{P_X(\text{DIS}(\{g \in \mathcal{F} : \|f - g\|_{L_1} \leq \varepsilon_0\}))}{\varepsilon_0} \vee 1.$$

Finally, the star number of [Hanneke and Yang \(2015\)](#) is the largest integer s such that there exist $f_0, f_1, \dots, f_s \in \mathcal{F}$ and $x_1, \dots, x_s \in \mathcal{X}$ such that for all $i \in \{1, \dots, s\}$,

$$\text{DIS}(\{f_0, f_i\}) \cap \{x_1, \dots, x_s\} = \{x_i\}.$$

3. Active learning with abstention

We present our main result, which is slightly sharper than our simplified statement in [Section 1](#).

Theorem 3.1 *Fix $\varepsilon, \delta \in (0, 1], p \in [0, \frac{1}{2}]$. Assume that the VC dimension of \mathcal{F} is equal to d . There is an active learning algorithm (namely, [Algorithm 3.2](#)) such that after requesting at most*

$$n = O\left(\frac{\theta(\varepsilon/p)}{p^2} \left(d \log^2\left(\frac{d}{p\varepsilon}\right) + \log^2\left(\frac{1}{\delta}\right)\right)\right).$$

labels, it returns, with probability at least $1 - \delta$, a classifier \hat{f}_p satisfying

$$R^p(\hat{f}_p) - R(f^*) \leq \varepsilon.$$

The proof of this result is included in [Appendix B](#). We are ready to present the algorithm achieving these guarantees.

Algorithm 3.2

- Let $V_0 = \mathcal{F}$ and

$$\alpha^2(n, \delta) = \frac{4}{n} \left(3d \log \frac{e(2n \vee d)}{d} + \log \frac{56}{\delta} \right), \quad (6)$$

and set $J = \min \{k \in \mathbb{N} : 148\alpha^2(2^{k-1}, \delta/(k+1)^2)/p \leq \varepsilon\}$.

- **for** j from 1 to J **do**

1. Sample $n_j = 2^{j-1}$ fresh i.i.d. instances $X_{2^{j-1}}, \dots, X_{2^j-1}$ from P_X and denote them by $Q_j = \{X_{2^{j-1}}, \dots, X_{2^j-1}\}$.
2. Define $D_j = \text{DIS}(V_{j-1}) \cap Q_j$.
3. Request labels for all instances in D_j .
4. Set $S_j = \bigcup_{X_m \in D_j} \{(X_m, Y_m)\}$.
5. Compute (any) ERM $\hat{f}_j \in \arg \min_{f \in V_{j-1}} R_{S_j}(f)$.
6. For $n_j = 2^{j-1}$ and $\delta_j = \delta/(j+1)^2$, update

$$V_j = \left\{ f \in V_{j-1} : \frac{|S_j|}{n_j} \left(R_{S_j}(f) - R_{S_j}(\hat{f}_j) \right) \leq 2\alpha^2(n_j, \delta_j) + 2\alpha(n_j, \delta_j) \sqrt{P_{Q_j}|f - \hat{f}_j|} \right\}.$$

7. **if** $\mathcal{D}(V_j, L_2(P_{Q_j})) > 49\alpha(n_j, \delta_j)/p$ **or** $j = J$.

- Consider the class $\widehat{\mathcal{G}}_j = \left\{ \frac{f + \widehat{f}_j}{2} : f \in V_j \right\}$ of $\{0, 1, 1/2\}$ -valued functions and convert it into $\{0, 1, *\}$ -valued class U_j by replacing $1/2$ with $*$.
- Define the mid-point classifier as (also Algorithm 3.3 below)

$$\widehat{f}_p \in \arg \min_{f \in U_j} R_{S_j}^p(f).$$

- **return** \widehat{f}_p .

• **end for**

Let us discuss the mechanism behind this algorithm. At iteration j , our strategy maintains a set V_j of candidate classifiers and requests the labels of instances that belong to the disagreement set of V_j . Then we update the set V_j by removing all classifiers making a large number of mistakes on the requested labels. This part of our algorithm is standard and corresponds to the principle standing behind all disagreement-based algorithms. Our first modification is that at each iteration we also compute the empirical diameter $\mathcal{D}(V_j, L_2(P_{Q_j}))$ and compare it with the threshold value $49\alpha(n_j, \delta_j)/p$. It follows that a large value of $\mathcal{D}(V_j, L_2(P_{Q_j}))$ indicates that the current iteration is too “noisy” and the reject option can help. Otherwise, if $\mathcal{D}(V_j, L_2(P_{Q_j}))$ is small, we proceed with the standard active learning strategy described above. Observe that $\mathcal{D}(V_j, L_2(P_{Q_j})) \leq 1$, but $49\alpha(n_j, \delta_j)/p$ is always greater than 1 for small values of j so that we never return \widehat{f}_p too early. Our second modification is that \widehat{f}_p is built using a two-step aggregation procedure described in detail in Section 3.1.

In Section 3.2 we show that Algorithm 3.2 is adaptive to the favorable noise assumptions: under the bounded noise assumption the event $\mathcal{D}(V_j, L_2(P_{Q_j})) \geq 49\alpha(n_j, \delta_j)/p$ almost never happens, $\Pr(\widehat{f}_p(X) = *)$ is small, and our algorithm mimics the behavior of the standard active learning strategy such as, for example, the one of [Dasgupta, Hsu, and Monteleoni \(2008\)](#).

3.1. Strategy of the proof and the mid-point algorithm

In this section, we introduce the *mid-point algorithm* used in Algorithm 3.2. This algorithm is a simplified version of the aggregation procedure in ([Bousquet and Zhivotovskiy, 2019](#)), inspired in turn by several key aggregation algorithms for the squared loss ([Audibert, 2008](#); [Lecué and Mendelson, 2009](#); [Mendelson, 2019](#)).

Algorithm 3.3 (Mid-point Algorithm)

- Given the labeled sample S_n and the class \mathcal{F} and the confidence δ and the abstention margin $p \in (0, \frac{1}{2}]$.
- Find (any) ERM $\widehat{g} \in \arg \min_{f \in \mathcal{F}} R_{S_n}(f)$.

- Let $\alpha(n, \delta)$ be as in (6) and define

$$V = \left\{ f \in \mathcal{F} : R_{S_n}(f) - R_{S_n}(\widehat{g}) \leq 2\alpha^2(n, \delta) + 2\alpha(n, \delta)\sqrt{P_n|f - \widehat{g}|} \right\},$$

- Consider the (random) set $\left\{ \frac{f + \widehat{g}}{2} : f \in V \right\}$ of $\{0, 1, 1/2\}$ -valued functions and convert it into $\{0, 1, *\}$ -valued set $\widehat{\mathcal{G}}$ by replacing $1/2$ with $*$.
- Define the mid-point classifier as $\widetilde{f}_p \in \arg \min_{f \in \widehat{\mathcal{G}}} R_{S_n}^p(f)$.
- **return** \widetilde{f}_p .

We are ready to provide a data-dependent bound for this algorithm.

Theorem 3.4 Fix $p \in (0, \frac{1}{2}]$, $\delta \in (0, 1)$. Assume that the VC dimension of \mathcal{F} is equal to d . In the notation of Algorithm 3.3, we have that, with probability at least $1 - \delta$, $f^* \in V$ and

$$R^p(\widetilde{f}_p) - R(f^*) \leq 8\alpha^2(n, \delta) + 12\alpha(n, \delta)\mathcal{D}(V, L_2(P_n)) - \frac{p}{4}\mathcal{D}^2(V, L_2(P_n)), \quad (7)$$

where $\alpha(n, \delta)$ is given by (6). In particular, on this event whenever $\mathcal{D}(V, L_2(P_n)) \geq 49\alpha(n, \delta)/p$, we have

$$R^p(\widetilde{f}_p) < R(f^*). \quad (8)$$

For the rest of this section, we discuss how this passive learning result is used in the proof of Theorem 3.1. The second part of the statement of Theorem 3.4 is one of our main technical insights. This result means that for any labeled sample of size m we may compute the data-dependent value $\mathcal{D}(V, L_2(P_m))$ and if it is larger than $49\alpha(m, \delta)/p$, we conclude that \widetilde{f}_p outperforms f^* . This is a favorable scenario in our context. Our second observation is that if $\mathcal{D}(V, L_2(P_m))$ is smaller than $49\alpha(m, \delta)/p$, then one may show that the region of disagreement of V is small. Indeed, by the definition of $\theta(\cdot)$ and the uniform convergence, one can show that

$$P(\text{DIS}(V)) \leq \theta(\mathcal{D}^2(V, L_2(P)))\mathcal{D}^2(V, L_2(P)) \approx \theta(\mathcal{D}^2(V, L_2(P_m)))\mathcal{D}^2(V, L_2(P_m)).$$

From this moment on, we use a standard active learning analysis following closely the well-proven techniques of Dasgupta, Hsu, and Monteleoni (2008) (see also (Hsu, 2010; Zhang and Chaudhuri, 2014)). In some sense, our analysis reveals a dichotomy: at each iteration of Algorithm 3.2, we have that either the noise of the problem is so high that even the negative excess risk is possible through the reject option, or the problem is as good as if the bounded noise assumption holds. Moreover, both situations can be empirically detected. As we pointed out, the negativity of the excess risk (regret) for improper learners in passive (online) learning as in (8) is not well understood. Among the few works exploring this is the paper of Audibert (2008) where the negativity of the excess risk is used to explain why the so-called progressive mixture rules are deviation suboptimal. More recently, Mourtada, Vaškevičius, and Zhivotovskiy (2021) used the negativity of the excess risk to observe the same suboptimality for truncated linear least squares. The analysis of this paper reveals that the negativity of the excess risk is helpful in active learning.

Remark 3.5 Maximizing (7) with respect to $\mathcal{D}(V, L_2(P_n))$, we have, with probability at least $1 - \delta$,

$$R^p(\tilde{f}_p) - R(f^*) \lesssim \frac{d \log(n/d) + \log(1/\delta)}{np}.$$

This bound is achieved in (Bousquet and Zhivotovskiy, 2019, Theorem 2.1) by an algorithm requiring an additional sample splitting step.

The proof of Theorem 3.4 blends the ideas of Audibert (2007); Mendelson (2019); Bousquet and Zhivotovskiy (2019). We defer these derivations to Appendix A.

3.2. Adaptation to the bounded noise assumption

Assume that Massart’s noise condition (5) holds with $h > 0$ and that the Bayes rule f_B^* belongs to the class \mathcal{F} . Under these conditions, there are active learning algorithms (see, for example, Zhang and Chaudhuri (2014); Hanneke and Yang (2015)) showing exponential savings in active learning whenever the disagreement coefficient is bounded. We show that Algorithm 3.2 adapts to the bounded noise condition in the sense that if $p \leq h/4$, it also provides exponential savings and outputs a classifier \hat{f}_p such that $\Pr(\hat{f}_p(X) = *)$ is small.

It is known (see, for example, (Herbei and Wegkamp, 2006, Equation (5))) that the optimal $\{0, 1, *\}$ -valued classifier with respect to Chow’s risk (1) is given by

$$f_p^*(x) = \begin{cases} f_B^*(x), & \text{if } |2 \Pr(Y = 1|x) - 1| \geq 2p, \\ *, & \text{otherwise.} \end{cases}$$

We see that if Massart’s noise condition holds, the Bayes rule f_B^* minimizes the risk (1) for all $p \leq h/2$. Therefore, if $p \leq h/2$ and $f_B^* \in \mathcal{F}$, the excess risk $R^p(\tilde{f}_j) - R(f_B^*)$ cannot be negative. Algorithm 3.2 is constructed in such a way that if it terminates before the J -th iteration, then the excess risk $R^p(\tilde{f}_j) - R(f^*)$ is negative (see the details of the proof of Theorem 3.1). This yields that Algorithm 3.2 finishes after J iterations in our case. We also have the following result.

Proposition 3.6 *In the notation of Theorem 3.1, assume that the noise condition (5) with the parameter $h > 0$ holds and that $f_B^* \in \mathcal{F}$. We have for $p \in [0, h/4]$ that the output classifier \hat{f}_p of Algorithm 3.2, with the number of label requests as in Theorem 3.1, satisfies, with probability at least $1 - \delta$,*

$$\Pr(\hat{f}_p(X) = *) \leq 4\varepsilon/h.$$

Moreover, on the same event, if all $*$ -s are replaced by random guessing, which corresponds to the risk R^0 , we also have

$$R^0(\hat{f}_p) - R(f^*) \leq 2\varepsilon.$$

With minor efforts, a similar result is achievable if for a chosen value of p , instead of assuming that $f_B^* \in \mathcal{F}$ and (5) holds, we have that the Bernstein assumption holds. That is, for any $f \in \mathcal{F}$,

$$h \Pr(f(X) \neq f^*(X)) \leq R(f) - R(f^*).$$

We omit these derivations in favor of a more transparent Proposition 3.6. The proof of Proposition 3.6 reveals that in the case where $f_B^* \in \mathcal{F}$, our passive Algorithm 3.3 abstains most of the times only on the instances where the Bayes optimal classifier does not satisfy the bounded noise assumption. This observation can be useful in the context of selective classification described above.

Proposition 3.7 *Assume that $f_B^* \in \mathcal{F}$. Then, the classifier \tilde{f}_p of passive Algorithm 3.3 trained on S_n satisfies, with probability at least $1 - \delta$,*

$$\Pr\left(\tilde{f}_p(X) = * \text{ and } |2\Pr(Y = 1|X) - 1| \geq 4p\right) \leq \frac{592}{np^2} \left(3d \log \frac{e(2n \vee d)}{d} + \log \frac{56}{\delta}\right).$$

4. Exponential savings under the model misspecification

In this section, we return to the setting where abstention is not allowed. The result of (Hanneke and Yang, 2015, Theorem 4) implies that if Massart’s noise assumption (5) holds, $\varepsilon \in (0, h/24)$, $\delta \in [0, 1/24]$, and the Bayes optimal rule f_B^* belongs to \mathcal{F} , then at least

$$\Omega\left(\frac{1}{h^2} \left((1-h) \min\left\{s, \frac{h}{\varepsilon}\right\} \log \frac{1}{\delta} + d\right)\right), \quad (9)$$

label requests are needed to construct \tilde{f} satisfying $R(\tilde{f}) - R(f^*) \leq \varepsilon$, with probability at least $1 - \delta$, for some distribution of the unlabeled data P_X . In particular, this result implies that the condition $s < \infty$ is necessary for exponential savings in the number of label requests in this setup. Further, the aforementioned lower bound in (Kääriäinen, 2006, Section 4.1) shows that the bounded noise assumption (5) is necessary on the set where at least two functions in \mathcal{F} disagree. Otherwise, one can easily choose P_X , so that the passive lower bound $\Omega\left(\frac{1}{\varepsilon^2}\right)$ holds.

However, it is not immediately clear if the usual assumption $f_B^* \in \mathcal{F}$ is also needed when (5) holds for some $h > 0$. The lower bound in (Kääriäinen, 2006, Theorem 3) exploits a specific situation where $f_B^* \notin \mathcal{F}$, and there exist $f, g \in \mathcal{F}$ that disagree on a set of *infinite size* and this leads to the agnostic lower bound $\Omega\left(\frac{1}{\varepsilon^2}\right)$. To avoid this obstacle in passive learning with deterministic labeling, Ben-David and Urner (2014)² introduced the notion of the *diameter* of \mathcal{F} . Recall that it is the smallest integer D such that

$$\sup_{f, g \in \mathcal{F}} |\{x \in \mathcal{X} : f(x) \neq g(x)\}| \leq D.$$

Similar to the star number, this complexity measure is infinite for many natural classes. However, it is still relevant, as many existing lower bounds in classification use the classes with a finite diameter (Massart and Nédélec, 2006; Audibert, 2009). Our second main result shows that if D is finite, one can avoid the model misspecification problem in active learning under Massart’s noise.

Theorem 4.1 *Assume that the diameter of \mathcal{F} is equal to D and the VC dimension of \mathcal{F} is equal to d . Fix $\varepsilon, \delta \in (0, 1]$. If Massart’s noise condition (5) is satisfied and $h > 0$ is known, then there is an active learning algorithm (namely, Algorithm 4.2) such that after requesting at most*

$$n = O\left(\frac{d\theta(\varepsilon/h)}{h^2} \log^2\left(\frac{d}{\varepsilon h \delta}\right) + \frac{D}{h^2} \log\left(\frac{D}{\delta}\right)\right)$$

labels, it returns a classifier \hat{f} satisfying, with probability at least $1 - \delta$,

$$R(\hat{f}) - R(f^*) \leq \varepsilon.$$

2. Their analysis is extended to the bounded noise case in (Bousquet and Zhivotovskiy, 2019).

This result implies that if both the diameter D and the disagreement coefficient $\theta(\cdot)$ are bounded, then exponential savings are possible (without assuming $f_B^* \in \mathcal{F}$) if Massart’s noise condition is satisfied with $h > 0$. Further, by (Hanneke and Yang, 2015, Theorem 10) we have $\theta(\varepsilon/h) \leq s$ for any distribution P_X of the unlabeled data. Combining this result with Theorem 4.1, (Kääriäinen, 2006, Theorem 3) and (9), we see that exponential savings are possible for all distributions satisfying Massart’s noise condition (5) with $h > 0$ if and only if both the diameter D and the star number s are finite. As we mentioned, both assumptions are quite restrictive and are not likely to be simultaneously satisfied for any non-trivial class of interest. One may also argue as in (Hanneke, 2014, Section 5.2) that to get exponential savings, instead of requiring that s is finite, it is necessary to assume that the disagreement coefficient $\theta(\cdot)$ is bounded. However, this necessary assumption leads to some technical complications and we prefer to avoid this statement for clarity.

The proof of Theorem 4.1 goes as follows. First, we fix $p = 1/2 - h/2$ and use Algorithm 3.2 to construct \hat{f}_p . By Theorem 3.1 we have, with probability at least $1 - \delta/3$,

$$R^p(\hat{f}_p) - R(f^*) \leq \varepsilon.$$

Observe that by the construction of Algorithm 3.2, and since we abstain only on the disagreement set of two classifiers, it holds that $|\{x \in \mathcal{X} : \hat{f}_p(x) = *\}| \leq D$. Therefore, if we specify the labels on these at most D instances, we obtain a $\{0, 1\}$ -valued classifier \hat{f} . Since Massart’s noise condition holds, a simple repeated-querying algorithm, similar to the one used in (Kääriäinen, 2006, Theorem 1), allows us to estimate the Bayes optimal rule f_B^* on the finite set $\{x \in \mathcal{X} : \hat{f}_p(x) = *\}$. The value $p = 1/2 - h/2$ is chosen to guarantee that

$$\Pr\left(f_B^*(X) \neq Y \text{ and } \hat{f}_p(X) = *\right) \leq (1/2 - p) \Pr\left(\hat{f}_p(X) = *\right),$$

implying

$$R(\hat{f}) - R(f^*) \leq R^p(\hat{f}_p) - R(f^*).$$

The formal description of the algorithm of Theorem 4.1 is as follows.

Algorithm 4.2

1. Fix $p = h/2$. Run Algorithm 3.2 with the number of label requests sufficient to output \hat{f}_p satisfying, with probability at least $1 - \delta/3$,

$$R^p(\hat{f}_p) - R(f^*) \leq \varepsilon/2.$$

2. Set $\mathcal{X}_{\hat{f}_p} = \{x \in \mathcal{X} : \hat{f}_p(x) = *\}$.
3. Sample $\frac{28D \log(6D/\delta)}{3h^2\varepsilon}$ fresh i.i.d. instances from P_X and denote them by Q .
4. Define $\tilde{f}_D : \mathcal{X}_{\hat{f}_p} \rightarrow \{0, 1\}$ as follows: for each $x \in \mathcal{X}_{\hat{f}_p}$ request the labels of all, but no more than the first $\frac{2 \log(6D/\delta)}{h^2}$ appearances of x in the sample Q . Set $\tilde{f}_D(x)$ to be equal to the majority vote of the labels of x obtained this way with ties broken arbitrarily.

5. Set $\hat{f}(x) = \begin{cases} \tilde{f}_D(x), & \text{if } x \in \mathcal{X}_{\hat{f}_p}, \\ \hat{f}_p(x), & \text{otherwise.} \end{cases}$
6. **return** \hat{f} .

The full proof of Theorem 4.1 appears in Appendix D. Finally, we remark that the algorithm of Theorem 4.1 is improper. That is, \hat{f} is not necessarily in \mathcal{F} . Ben-David and Uner (2014, Theorem 13) show that in passive learning with deterministic labeling, it is necessary to use improper learning algorithms to obtain the optimal sample complexity. A natural question is to understand if it is also the case in the context of Theorem 4.1.

Acknowledgments

The article was prepared within the framework of the HSE University Basic Research Program. Nikita Puchkin was supported in part by Young Russian Mathematics award. Nikita Zhivotovskiy is funded in part by ETH Foundations of Data Science (ETH-FDS).

References

- Kenneth S Alexander. Rates of growth and sample moduli for weighted empirical processes indexed by sets. *Probability Theory and Related Fields*, 75(3):379–423, 1987.
- Martin Anthony and Peter L Bartlett. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, 2009.
- Jean-Yves Audibert. Proof of the optimality of the empirical star algorithm. *Technical note*, 2007.
- Jean-Yves Audibert. Progressive mixture rules are deviation suboptimal. In *Advances in Neural Information Processing Systems*, pages 41–48, 2008.
- Jean-Yves Audibert. Fast learning rates in statistical inference through aggregation. *The Annals of Statistics*, 37(4):1591–1646, 2009.
- Maria-Florina Balcan and Phil Long. Active and passive learning of linear separators under log-concave distributions. In *Conference on Learning Theory*, pages 288–316, 2013.
- Maria-Florina Balcan, Andrei Broder, and Tong Zhang. Margin based active learning. In *International Conference on Computational Learning Theory*, pages 35–50. Springer, 2007.
- Maria-Florina Balcan, Alina Beygelzimer, and John Langford. Agnostic active learning. *Journal of Computer and System Sciences*, 75(1):78–89, 2009.
- Peter L Bartlett and Shahar Mendelson. Empirical minimization. *Probability Theory and Related Fields*, 135(3):311–334, 2006.
- Peter L Bartlett and Marten H Wegkamp. Classification with a reject option using a hinge loss. *Journal of Machine Learning Research*, 9(Aug):1823–1840, 2008.

- Shai Ben-David and Ruth Uerner. The sample complexity of agnostic learning under deterministic labels. In *Conference on Learning Theory*, pages 527–542, 2014.
- Alina Beygelzimer, Sanjoy Dasgupta, and John Langford. Importance weighted active learning. In *Proceedings of the 26th Annual International Conference on Machine Learning (ICML)*, pages 49–56, 2009.
- Stéphane Boucheron, Olivier Bousquet, and Gábor Lugosi. Theory of classification: A survey of some recent advances. *ESAIM: Probability and Statistics*, 9:323–375, 2005.
- Olivier Bousquet and Nikita Zhivotovskiy. Fast classification rates without standard margin assumptions. *Information and Inference: A Journal of the IMA (forthcoming)*, 2019.
- Marat Burnashev and Kamil’ Zigangirov. An interval estimation problem for controlled observations. *Problemy Peredachi Informatsii*, 10(3):51–61, 1974.
- Rui M Castro and Robert D Nowak. Minimax bounds for active learning. *IEEE Transactions on Information Theory*, 54(5):2339–2353, 2008.
- Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge university press, 2006.
- C Chow. On optimum recognition error and reject tradeoff. *IEEE Transactions on information theory*, 16(1):41–46, 1970.
- David Cohn, Les Atlas, and Richard Ladner. Improving generalization with active learning. *Machine learning*, 15(2):201–221, 1994.
- Corinna Cortes, Giulia DeSalvo, and Mehryar Mohri. Learning with rejection. In *International Conference on Algorithmic Learning Theory*, pages 67–82. Springer, 2016.
- Sanjoy Dasgupta. Coarse sample complexity bounds for active learning. *Advances in neural information processing systems*, 18:235–242, 2005.
- Sanjoy Dasgupta, Daniel J Hsu, and Claire Monteleoni. A general agnostic active learning algorithm. In *Advances in neural information processing systems*, pages 353–360, 2008.
- Ran El-Yaniv and Yair Wiener. On the foundations of noise-free selective classification. *Journal of Machine Learning Research*, 11(May):1605–1641, 2010.
- Ran El-Yaniv and Yair Wiener. Active learning via perfect selective classification. *Journal of Machine Learning Research*, 13(1):255–279, 2012.
- Yoav Freund, Yishay Mansour, and Robert E Schapire. Generalization bounds for averaged classifiers. *The Annals of Statistics*, 32(4):1698–1722, 2004.
- Roei Gelbhart and Ran El-Yaniv. The relationship between agnostic selective classification, active learning and the disagreement coefficient. *Journal of Machine Learning Research*, 20:33–1, 2019.

- Evarist Giné and Vladimir Koltchinskii. Concentration inequalities and asymptotic results for ratio type empirical processes. *The Annals of Probability*, 34(3):1143–1216, 2006.
- G Golubev and B Levit. Sequential recovery of analytic periodic edges in binary image models. *Mathematical Methods of Statistics*, 12(1):95–115, 2003.
- Steve Hanneke. A bound on the label complexity of agnostic active learning. In *Proceedings of the 24th international conference on Machine learning*, pages 353–360, 2007.
- Steve Hanneke. Rates of convergence in active learning. *The Annals of Statistics*, 39(1):333–361, 2011.
- Steve Hanneke. Theory of disagreement-based active learning. *Foundations and Trends in Machine Learning*, 7(2–3):131–309, 2014.
- Steve Hanneke and Liu Yang. Minimax analysis of active learning. *Journal of Machine Learning Research*, 16(1):3487–3602, 2015.
- Steve Hanneke and Liu Yang. Surrogate losses in passive and active learning. *Electronic Journal of Statistics*, 13(2):4646–4708, 2019.
- David Haussler, Jyrki Kivinen, and Manfred K Warmuth. Sequential prediction of individual sequences under general loss functions. *IEEE Transactions on Information Theory*, 44(5):1906–1925, 1998.
- Radu Herbei and Marten H Wegkamp. Classification with reject option. *Canadian Journal of Statistics*, 34(4):709–721, 2006.
- Daniel Hsu. Algorithms for Active Learning. *PhD Thesis. University of California, San Diego*, 2010.
- Matti Kääriäinen. Active learning in the non-realizable case. In *International Conference on Algorithmic Learning Theory*, pages 63–77. Springer, 2006.
- Vladimir Koltchinskii. Rademacher complexities and bounding the excess risk in active learning. *Journal of Machine Learning Research*, 11:2457–2485, 2010.
- Alexander Korostelev. On minimax rates of convergence in image models under sequential design. *Statistics & Probability Letters*, 43(4):369–375, 1999.
- Guillaume Lecué and Shahar Mendelson. Aggregation via empirical risk minimization. *Probability Theory and Related Fields*, 145(3–4):591–613, 2009.
- Andrea Locatelli, Alexandra Carpentier, and Samory Kpotufe. Adaptivity to noise parameters in nonparametric active learning. In *Conference on Learning Theory*, pages 1383–1416, 2017.
- Andrea Locatelli, Alexandra Carpentier, and Samory Kpotufe. An adaptive strategy for active learning with smooth decision boundary. In *Algorithmic Learning Theory*, pages 547–571, 2018.
- Pascal Massart and Élodie Nédélec. Risk bounds for statistical learning. *The Annals of Statistics*, 34(5):2326–2366, 2006.

- Shahar Mendelson. An unrestricted learning procedure. *Journal of ACM*, 66(6), 2019.
- Stanislav Minsker. Plug-in approach to active learning. *Journal of Machine Learning Research*, 13(1):67–90, 2012.
- Jaouad Mourtada, Tomas Vaškevičius, and Nikita Zhivotovskiy. Distribution-free robust linear regression. *arXiv preprint arXiv:2102.12919*, 2021.
- Gergely Neu and Nikita Zhivotovskiy. Fast rates for online prediction with abstention. In *Conference on Learning Theory*, pages 3030–3048, 2020.
- Norbert Sauer. On the density of families of sets. *Journal of Combinatorial Theory, Series A*, 13(1):145 – 147, 1972.
- Shubhanshu Shekhar, Mohammad Ghavamzadeh, and Tara Javidi. Active learning for classification with abstention. In *2020 IEEE International Symposium on Information Theory (ISIT)*, pages 2801–2806. IEEE, 2020.
- Alexander B Tsybakov. Optimal aggregation of classifiers in statistical learning. *The Annals of Statistics*, 32(1):135–166, 2004.
- Alexandre B Tsybakov. Optimal rates of aggregation. In *Learning Theory and Kernel Machines*, pages 303–313. Springer, 2003.
- Vladimir Vapnik and Alexey Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Proceedings of the USSR Academy of Sciences*, 1968.
- Vladimir Vapnik and Alexey Chervonenkis. *Theory of Pattern Recognition*. Nauka. Moscow., 1974.
- Volodimir Vovk. Aggregating strategies. *Proceedings of the Third Annual Workshop on Computational Learning Theory*, pages 371–386, 1990.
- Songbai Yan, Kamalika Chaudhuri, and Tara Javidi. Active learning from imperfect labelers. In *Advances in Neural Information Processing Systems*, pages 2128–2136, 2016.
- Chicheng Zhang and Kamalika Chaudhuri. Beyond disagreement-based agnostic active learning. In *Advances in Neural Information Processing Systems*, pages 442–450, 2014.
- Nikita Zhivotovskiy and Steve Hanneke. Localization of VC classes: Beyond local Rademacher complexities. *Theoretical Computer Science*, 742:27–49, 2018.

Appendix A. Proof of Theorem 3.4

Throughout the proof, we identify $*$ with $1/2$ and convert $\{0, 1, *\}$ -valued functions to $\{0, 1, 1/2\}$ -valued ones by replacing $*$ with $1/2$ and vice versa. We refer to (Bousquet and Zhivotovskiy, 2019), where the connections between Chow’s risk, strong convexity and the model selection aggregation are presented. In contrast, we provide a short and direct proof with explicit constants. As a result, our algorithm is simpler (see Remark 3.5) and the proof is based only on the tools available in (Vapnik and Chervonenkis, 1974). In the notation of Theorem 3.4 and Algorithm 3.3, we need the following auxiliary lemma.

Lemma A.1 *With probability at least $1 - 4\delta/7$, for all $f \in V$, we have*

$$R^p(\tilde{f}_p) - R(f) \leq 4\alpha^2(n, \delta) + 8\alpha(n, \delta)\mathcal{D}(V, L_2(P_n)) - pP_n(f - \hat{g})^2,$$

where $\alpha(n, \delta)$ is given by (6).

Proof Lemma E.3 implies that, with probability at least $1 - 4\delta/7$, for any $f \in V$, it holds that

$$R^p(\tilde{f}_p) - R^p(f) \leq R_n^p(\tilde{f}_p) - R_n^p(f) + 4\alpha^2(n, \delta) + 8\alpha(n, \delta)\sqrt{P_n(\tilde{f}_p - f)^2}.$$

By the definition of \tilde{f}_p , there exists $f_0 \in V$ such that $\tilde{f}_p = (f_0 + \hat{g})/2$. Then, by the convexity of the seminorm $L_2(P_n)$,

$$\sqrt{P_n(\tilde{f}_p - f)^2} \leq \left(\sqrt{P_n(f_0 - f)^2} + \sqrt{P_n(\hat{g} - f)^2} \right) / 2 \leq \mathcal{D}(V, L_2(P_n)).$$

This implies

$$R^p(\tilde{f}_p) - R^p(f) \leq R_n^p(\tilde{f}_p) - R_n^p(f) + 4\alpha^2(n, \delta) + 8\alpha(n, \delta)\mathcal{D}(V, L_2(P_n)).$$

Define the loss function corresponding to the risk (1) as

$$\ell^p(y, f(x)) = \mathbb{1}[y \neq f(x) \text{ and } f(x) \neq 1/2] + \left(\frac{1}{2} - p \right) \mathbb{1}[f(x) = 1/2]. \quad (10)$$

A direct calculation shows that for any $\{0, 1\}$ -valued functions f, g and any $y \in \{0, 1\}$ it holds that for all $x \in \mathcal{X}$,

$$\ell^p\left(y, \frac{f(x) + g(x)}{2}\right) = \frac{1}{2}\ell^p(y, f(x)) + \frac{1}{2}\ell^p(y, g(x)) - p(f(x) - g(x))^2.$$

Observe that $R_n^p(f) = P_n\ell^p(Y, f(X))$. By the definition of \tilde{f}_p and the empirical risk minimizer \hat{g} , we have for any $f \in V$,

$$R_n^p(\tilde{f}_p) \leq R_n^p\left(\frac{f + \hat{g}}{2}\right) = \frac{1}{2}R_n^p(f) + \frac{1}{2}R_n^p(\hat{g}) - pP_n(f - \hat{g})^2 \leq R_n^p(f) - pP_n(f - \hat{g})^2,$$

and the claim of the lemma follows. ■

Proof of Theorem 3.4 To prove (7), take $h \in V$, such that $P_n(h - \hat{g})^2 \geq \mathcal{D}^2(V, L_2(P_n))/4$. The existence of such an h follows from the definition of V . Due to Lemma E.2, there is an event E such that $\Pr(E) \geq 1 - 3\delta/7$ and f^* belongs to V on E . Furthermore, on this event, it holds that

$$\begin{aligned} R(h) - R(f^*) &\leq R_n(h) - R_n(f^*) + 2\alpha^2(n, \delta) + 2\alpha(n, \delta)\sqrt{P_n(h - f^*)^2} \\ &\leq R_n(h) - R_n(f^*) + 2\alpha^2(n, \delta) + 2\alpha(n, \delta)\mathcal{D}(V, L_2(P_n)). \end{aligned}$$

Since $R_n(f^*) \geq R_n(\hat{g})$, we have

$$\begin{aligned} R_n(h) - R_n(f^*) &\leq R_n(h) - R_n(\hat{g}) \\ &\leq 2\alpha^2(n, \delta) + 2\alpha(n, \delta)\sqrt{P_n(h - \hat{g})^2} \\ &\leq 2\alpha^2(n, \delta) + 2\alpha(n, \delta)\mathcal{D}(V, L_2(P_n)), \end{aligned}$$

where the second and third inequalities hold since $h, \hat{g} \in V$. This yields

$$R(h) - R(f^*) \leq 4\alpha^2(n, \delta) + 4\alpha(n, \delta)\mathcal{D}(V, L_2(P_n)).$$

Applying Lemma A.1 and the union bound, we have, with probability at least $1 - \delta$,

$$\begin{aligned} R^p(\tilde{f}_p) - R(f^*) &= R^p(\tilde{f}_p) - R(h) + R(h) - R(f^*) \\ &= R^p(\tilde{f}_p) - R^p(h) + R(h) - R(f^*) \\ &\leq 8\alpha^2(n, \delta) + 12\alpha(n, \delta)\mathcal{D}(V, L_2(P_n)) - pP_n(h - \hat{g})^2 \\ &\leq 8\alpha^2(n, \delta) + 12\alpha(n, \delta)\mathcal{D}(V, L_2(P_n)) - \frac{p}{4}\mathcal{D}^2(V, L_2(P_n)). \end{aligned}$$

Hence, the proof of (7) is finished. To prove (8), we consider the largest root of the equation

$$\frac{p}{4}x^2 - 12\alpha(n, \delta)x - 8\alpha^2(n, \delta) = 0$$

and show that it is smaller than $49\alpha(n, \delta)/p$. Indeed, taking into account that $p \leq 1/2$, we obtain

$$\frac{2}{p} \left(12\alpha(n, \delta) + \sqrt{144\alpha^2(n, \delta) + 8p\alpha^2(n, \delta)} \right) \leq \frac{2\alpha(n, \delta)}{p} \left(12 + \sqrt{148} \right) < \frac{49\alpha(n, \delta)}{p}.$$

Thus, if $\mathcal{D}(V, L_2(P_n)) \geq 49\alpha(n, \delta)/p$, then $R^p(\tilde{f}_p) - R(f^*) < 0$. ■

Appendix B. Proof of Theorem 3.1

As we mentioned, our proof follows the standard arguments with several technical modifications needed to incorporate the result of Theorem 3.4. For the ease of exposure, we split the proof into several steps.

Step 0. Note that, for any $f, g \in V_{j-1}$, it holds that

$$|S_j| (R_{S_j}(f) - R_{S_j}(g)) = n_j (R_{Q_j}(f) - R_{Q_j}(g)),$$

because $f(x) = g(x)$ for all $x \in Q_j \setminus D_j$.

Step 1. Let E_1 be an event such that $\Pr(E_1) \geq 1 - \delta_1$, f^* belongs to V_1 on E_1 , and, moreover,

$$R^p(\tilde{f}_1) - R(f^*) \leq 8\alpha^2(n_1, \delta_1) + 12\alpha(n_1, \delta_1)\mathcal{D}(V_1, L_2(P_n)) - \frac{p}{4}\mathcal{D}^2(V_1, L_2(P_n)),$$

where \tilde{f}_1 is the output of Algorithm 3.3 applied to the class $V_0 = \mathcal{F}$ with the confidence δ_1 . The existence of such E_1 is guaranteed by Theorem 3.4. Given an integer $j \geq 2$, define an event E_j as follows. Let E_j be such that $f^* \in V_j$ on E_j , and, on the same event, it holds that

$$R^p(\tilde{f}_j) - R(f^*) \leq 8\alpha^2(n_j, \delta_j) + 12\alpha(n_j, \delta_j)\mathcal{D}(V_j, L_2(P_n)) - \frac{p}{4}\mathcal{D}^2(V_j, L_2(P_n)),$$

where \tilde{f}_j is the output of Algorithm 3.3 applied to the class V_{j-1} with the confidence δ_j . Theorem 3.4 implies that $\Pr(E_j | E_1, \dots, E_{j-1}) \geq 1 - \delta_j$. Note that, by the definition of E_1, \dots, E_J ,

$$\begin{aligned} \Pr\left(\bigcap_{j=1}^J E_j\right) &= \Pr(E_J | E_1, \dots, E_{J-1}) \Pr(E_{J-1} | E_1, \dots, E_{J-2}) \dots \Pr(E_2 | E_1) \Pr(E_1) \\ &\geq \prod_{j=1}^J (1 - \delta_j) \geq 1 - \sum_{j=1}^J \delta_j = 1 - \sum_{j=1}^J \frac{\delta}{(1+j)^2} \geq 1 - \sum_{j=1}^{\infty} \frac{\delta}{(1+j)^2} \geq 1 - \frac{2\delta}{3}. \end{aligned}$$

In particular, this yields that, with probability at least $1 - 2\delta/3$, $f^* \in V_j$ for any $j \in \{1, \dots, J\}$.

Step 2. Consider an event $E_{\cap} = \bigcap_{j=1}^J E_j$, $\Pr(E_{\cap}) \geq 1 - 2\delta/3$, where E_1, \dots, E_J were introduced in the previous step. We prove that on this event

$$R^p(\hat{f}_p) - R(f^*) \leq \varepsilon.$$

Theorem 3.4 implies that, for any $j \in \{1, \dots, J\}$, one has either $\mathcal{D}(V_j, L_2(P_n)) < 49\alpha(n_j, \delta_j)/p$ or $R^p(\tilde{f}_j) - R(f^*) < 0$ on E_{\cap} . Hence, if the procedure terminates ahead of time, we have $R^p(\hat{f}_p) - R(f^*) < 0$, with probability at least $1 - 2\delta/3$. In this case the proof is complete. Otherwise, we have $\mathcal{D}(V_j, L_2(P_{Q_j})) < 49\alpha(n_j, \delta_j)/p$ for all $j \in \{1, \dots, J\}$, with probability at least $1 - 2\delta/3$. Then, on the final iteration, we obtain

$$R^p(\tilde{f}_J) - R(f^*) \leq 8\alpha^2(n, \delta) + 12\alpha(n, \delta)\mathcal{D}(V_J, L_2(P_n)) - \frac{p}{4}\mathcal{D}^2(V_J, L_2(P_n)).$$

Maximizing the right-hand side over $\mathcal{D}(V_J, L_2(P_n))$ and taking into account that $p \leq 1/2$, we get

$$R^p(\tilde{f}_J) - R(f^*) \leq 8\alpha^2(n_J, \delta_J) + \frac{144\alpha^2(n_J, \delta_J)}{p} \leq \frac{148\alpha^2(n_J, \delta_J)}{p},$$

with probability at least $1 - 2\delta/3$. Recall that $n_J = 2^{J-1}$ and $\delta_J = \frac{\delta}{(J+1)^2}$. Since J satisfies the condition

$$\frac{148\alpha^2(n_J, \delta_J)}{p} \leq \varepsilon,$$

we have $R^p(\hat{f}_p) - R(f^*) = R^p(\tilde{f}_J) - R(f^*) \leq \varepsilon$, with probability at least $1 - 2\delta/3$.

Step 3. The total number of labels requested by Algorithm 3.2 is equal to

$$\sum_{j=1}^{T \wedge J} \sum_{i=1}^{n_j} \mathbb{1}(X_i \in \text{DIS}(V_{j-1})),$$

where T is the iteration when the procedure terminates.

Previously, we proved that either $\mathcal{D}(V_{j-1}, L_2(P_n)) < 49\alpha(n_{j-1}, \delta_{j-1})/p$ for all $j \leq J$ or we terminate at the moment $j < J$ and $R^p(\tilde{f}_j) - R(f^*) < 0 < \varepsilon$, with probability at least $1 - 2\delta/3$. Therefore, we may assume in the analysis that $\mathcal{D}(V_{j-1}, L_2(P_n)) < 49\alpha(n_{j-1}, \delta_{j-1})/p$

for all $j \leq T$. Applying Lemma E.2, we have, with probability at least $1 - 3\delta_j/7$,

$$\begin{aligned} \sup_{f,g \in V_j} P|f - g| &\leq \sup_{f,g \in V_j} \left(P_{Q_j}|f - g| + \alpha(n_j, \delta_j) \sqrt{P_{Q_j}|f - g| + \alpha^2(n_j, \delta_j)} \right) \\ &\leq \mathcal{D}^2(V_j, L_2(P_{Q_j})) + \alpha(n_j, \delta_j) \mathcal{D}(V_j, L_2(P_{Q_j})) + \alpha^2(n_j, \delta_j) \\ &\leq \frac{50^2 \alpha^2(n_j, \delta_j)}{p^2}. \end{aligned}$$

Consequently, on this event, $\text{DIS}(V_{j-1}) \subseteq \text{DIS}(\mathcal{B}(f^*, \xi_{j-1}))$, where we defined

$$\xi_{j-1} = \frac{50^2 \alpha^2(n_{j-1}, \delta_{j-1})}{p^2},$$

and

$$\mathcal{B}(f^*, r) = \{f \in \mathcal{F} : P|f - f^*| \leq r\}.$$

This yields $P_X(\text{DIS}(V_{j-1})) \leq \theta(\xi_{j-1})\xi_{j-1}$. Due to Bernstein's inequality and since for $a, b \geq 0$, $\sqrt{2ab} \leq a/2 + b$, (conditionally on V_{j-1}) it holds that, with probability at least $1 - \delta_j/56$,

$$\begin{aligned} \sum_{i=1}^{n_j} \mathbb{1}(X_i \in \text{DIS}(V_{j-1})) &\leq n_j \theta(\xi_{j-1}) \xi_{j-1} + \sqrt{2n_j \theta(\xi_{j-1}) \xi_{j-1} \log(56/\delta_j)} + 2 \log(56/\delta_j) \\ &\leq \frac{3}{2} n_j \theta(\xi_{j-1}) \xi_{j-1} + 3 \log(56/\delta_j) \\ &= 3n_{j-1} \theta(\xi_{j-1}) \xi_{j-1} + 3 \log(56/\delta_j) \\ &\leq \frac{4 \cdot 50^2 \theta(\xi_{j-1})}{p^2} (9d + 3dj + 2 \log(1+j) + \log(56/\delta)) \\ &\quad + 6 \log(1+j) + 3 \log(56/\delta) \\ &\leq \frac{4 \cdot 50^2 \theta(\xi_{j-1})}{p^2} (9d + (3d+3)j + 2 \log(56/\delta)). \end{aligned}$$

By the union bound, with probability at least

$$\Pr(E_\cap) - \sum_{j=1}^J \frac{3\delta_j}{7} - \sum_{j=1}^J \frac{\delta_j}{56} \geq 1 - \frac{2\delta}{3} - \frac{2\delta}{3} \left(\frac{3}{7} + \frac{1}{56} \right) \geq 1 - \delta,$$

the total number of requested labels is not greater than

$$\sum_{j=1}^{T \wedge J} \frac{4 \cdot 50^2 \theta(\xi_{j-1})}{p^2} (9d + (3d+2)j + \log(56/\delta)) + 6 \log(1+j) + 3 \log(56/\delta),$$

and, using the fact that $\theta(\cdot)$ is a decreasing function, we obtain

$$\begin{aligned}
 & \sum_{j=1}^{T \wedge J} \frac{4 \cdot 50^2 \theta(\xi_{j-1})}{p^2} (9d + (3d + 3)j + 2 \log(56/\delta)) \\
 & \leq \sum_{j=1}^{T \wedge J} \frac{4 \cdot 50^2 \theta(\xi_{j-1})}{p^2} (9d + (3d + 3)j + 2 \log(56/\delta)) \\
 & \leq \frac{36 \cdot 50^2 \theta(\xi_{J-1}) d J}{p^2} + \frac{2 \cdot 50^2 \theta(\xi_{J-1}) (3d + 3) J (J + 1)}{p^2} + \frac{8 \cdot 50^2 \theta(\xi_{J-1}) J \log(56/\delta)}{p^2} \\
 & \lesssim \frac{\theta(\varepsilon/p)}{p^2} \left(d \log^2 \left(\frac{d}{p\varepsilon} \right) + \log \left(\frac{d}{p\varepsilon} \right) \log \left(\frac{1}{\delta} \right) + \log \left(\frac{1}{\delta} \right) \log \log \left(\frac{1}{\delta} \right) \right). \tag{11}
 \end{aligned}$$

To prove the last inequality, we took into account that, by the definition of J ,

$$\xi_{J-1} = 50^2 \alpha^2 (n_{J-1}, \delta_{J-1}) / p^2 \geq (50^2 \varepsilon) / (148p) > \varepsilon / p.$$

Moreover, it is easy to see that

$$n_{J-1} \lesssim \frac{d \log(d/\varepsilon)}{p\varepsilon} + \log \frac{1}{\delta_{J-1}} \lesssim \max \left\{ \frac{d \log(d/\varepsilon)}{p\varepsilon}, \log \frac{1}{\delta_{J-1}} \right\},$$

which yields $J \lesssim \log \left(\frac{d}{p\varepsilon} \right) + \log \log(1/\delta)$. Therefore, the upper bound (11) follows. For the sake of presentation, in our statement we use a simple relaxation of (11).

Appendix C. Proof of Proposition 3.6 and Proposition 3.7

Proof of Proposition 3.6 Let $\eta(x) = \Pr(Y = 1 | X = x)$. Theorem 3.1 yields that $R^p(\widehat{f}_p) - R(f_B^*) \leq \varepsilon$, with probability at least $1 - \delta$. By (Boucheron et al., 2005, an equality on page 341), we have

$$\begin{aligned}
 \varepsilon & \geq R^p(\widehat{f}_p) - R(f_B^*) = \mathbb{E} |2\eta(X) - 1| \mathbb{1} \left[\widehat{f}_p(X) \neq f_B^*(X) \text{ and } \widehat{f}_p(X) \neq * \right] \\
 & \quad + \mathbb{E} \left(\frac{1}{2} - p - \min\{\eta(X), 1 - \eta(X)\} \right) \mathbb{1} \left[\widehat{f}_p(X) = * \right] \\
 & \geq 0 + \left(\frac{h}{2} - p \right) \Pr \left(\widehat{f}_p(X) = * \right) \\
 & \geq \frac{h}{4} \Pr \left(\widehat{f}_p(X) = * \right),
 \end{aligned}$$

where we used the noise condition (5). To finish the proof, note that on the same event we have

$$R^0(\widehat{f}_p) - R(f^*) = R^p(\widehat{f}_p) - R(f^*) + p \Pr(\widehat{f}_p(X) = *) \leq \varepsilon + \frac{4p\varepsilon}{h} \leq 2\varepsilon.$$

The claim follows. ■

Proof of Proposition 3.7 As above, let $\eta(x) = \Pr(Y = 1 | X = x)$. We show a slightly more general result, that is, for any $u > 0$, with probability at least $1 - \delta$,

$$\Pr \left(\widetilde{f}_p(X) = * \text{ and } |2\eta(X) - 1| \geq 2(p + u) \right) \leq \frac{592}{np u} \left(3d \log \frac{e(2n \vee d)}{d} + \log \frac{56}{\delta} \right).$$

First, maximizing (7) with respect to $\mathcal{D}(V, L_2(P_n))$, we obtain that

$$R^p(\tilde{f}_p) - R(f_B^*) \leq \frac{148\alpha^2(n, \delta)}{p} = \frac{592}{np} \left(3d \log \frac{e(2n \vee d)}{d} + \log \frac{56}{\delta} \right),$$

with probability at least $1 - \delta$. On the other hand, similarly to the proof of Proposition 3.6, we have

$$\begin{aligned} & R^p(\hat{f}_p) - R(f_B^*) \\ &= \mathbb{E} |2\eta(X) - 1| \mathbb{1} \left[\tilde{f}_p(X) \neq f_B^*(X) \text{ and } \tilde{f}_p(X) \neq * \right] \\ &\quad + \mathbb{E} \left(\frac{1}{2} - p - \min\{\eta(X), 1 - \eta(X)\} \right) \mathbb{1} \left[\tilde{f}_p(X) = * \right] \\ &\geq 0 + \mathbb{E} \left(\frac{1}{2} - p - \min\{\eta(X), 1 - \eta(X)\} \right) \mathbb{1} \left[\tilde{f}_p(X) = * \text{ and } |2\eta(X) - 1| \geq 2(p + u) \right] \\ &\geq u \Pr \left(\tilde{f}_p(X) = * \text{ and } |2\eta(X) - 1| \geq 2(p + u) \right). \end{aligned}$$

Combining these two bounds, we obtain

$$\Pr \left(\tilde{f}_p(X) = * \text{ and } |2\eta(X) - 1| \geq 2(p + u) \right) \leq \frac{592}{np u} \left(3d \log \frac{e(2n \vee d)}{d} + \log \frac{56}{\delta} \right).$$

The claim follows by choosing $u = p$. ■

Appendix D. Proof of Theorem 4.1

Step 1. First, we prove that conditionally on the observations required to construct \hat{f}_p , with probability at least $1 - \delta/3$, for any $x \in \mathcal{X}_{\hat{f}_p}$ it holds that either

$$f_B^*(x) = \tilde{f}_D(x), \text{ or } \Pr(\{x\}) < \frac{\varepsilon}{2D}. \quad (12)$$

Assume that there is at least one $x \in \mathcal{X}_{\hat{f}_p}$ such that $\Pr(\{x\}) \geq \frac{\varepsilon}{2D}$, since otherwise we already have (12). For some integer m , we estimate the number of unlabeled instances sufficient to observe each $x \in \mathcal{X}_{\hat{f}_p}$ having $\Pr(\{x\}) \geq \frac{\varepsilon}{2D}$ at least m times. Fix any such x and apply the Bernstein inequality to Bernoulli random variables $\eta_j = \mathbb{1}[x \text{ occurred on the } j\text{-th trial}]$, $1 \leq j \leq N$. Then the probability that a point $x \in \mathcal{X}_{\hat{f}_p}$ with probability mass $\Pr(\{x\}) \geq \frac{\varepsilon}{2D}$ was observed less than m times after $N \geq m \Pr(\{x\})$ trials does not exceed

$$\exp \left(- \frac{N(\Pr(\{x\}) - m/N)^2}{2\Pr(\{x\})(1 - \Pr(\{x\})) + 2(\Pr(\{x\}) - m/N)/3} \right) \leq \exp \left(- \frac{3N\varepsilon - 6mD}{16D} \right).$$

Since the diameter is finite we have by the construction of Algorithm 3.2 that $|\mathcal{X}_{\hat{f}_p}| \leq D$. Thus, by the union bound, the probability that there exists $x \in \mathcal{X}_{\hat{f}_p}$, $\Pr(\{x\}) \geq \frac{\varepsilon}{2D}$ such that it occurred less than m times after N trials is not greater than $D \exp \left(- \frac{3N\varepsilon - 6mD}{16D} \right)$. Thus, we need at most $D(6m + 16 \log(3D/\delta)) / (3\varepsilon)$ unlabeled instances to satisfy this, with probability at least $1 - \delta/3$.

Recall that $f_B^*(x) = \mathbb{1}[\Pr(Y = 1|X = x) \geq 1/2]$. By our assumption, we have for any x , $|2\Pr(Y = 1|X = x) - 1| \geq h$. By Hoeffding's inequality and the union bound, the probability that there is $x \in \mathcal{X}_{\hat{f}_p}$ such that the majority vote $f_D(x)$ is not equal to $f_B^*(x)$, is bounded by

$$2D \exp(-2m(h/2)^2) \leq \delta/3,$$

whenever the number of label requests m for each instance satisfies $m \geq \frac{2\log(6D/\delta)}{h^2}$. This proves (12). Since $|2\Pr(Y = 1|X) - 1| \geq h$ almost surely, we have

$$\Pr(f_B^*(X) \neq Y|X = x) = \Pr(Y = 1|X = x) \wedge (1 - \Pr(Y = 1|X = x)) \leq (1 - h)/2.$$

This implies that on the event where (12) holds, since $|\mathcal{X}_{\hat{f}_p}| \leq D$, we also have

$$\Pr(\tilde{f}_D(X) \neq Y \text{ and } X \in \mathcal{X}_{\hat{f}_p}) \leq \left((1 - h)P_X(\mathcal{X}_{\hat{f}_p}) + \varepsilon \right) / 2. \quad (13)$$

Step 2. By the union bound, (13) and the first step of our algorithm we have, with probability at least $1 - \delta$,

$$\begin{aligned} R(\hat{f}) &= \Pr(\hat{f}_p(X) \neq Y \text{ and } X \notin \mathcal{X}_{\hat{f}_p}) + \Pr(\tilde{f}_D(X) \neq Y \text{ and } X \in \mathcal{X}_{\hat{f}_p}) \\ &\leq \Pr(\hat{f}_p(X) \neq Y \text{ and } X \notin \mathcal{X}_{\hat{f}_p}) + \left((1 - h)P_X(\mathcal{X}_{\hat{f}_p}) + \varepsilon \right) / 2 \\ &= R^p(\hat{f}_p) + \varepsilon/2 \leq R(f^*) + \varepsilon. \end{aligned}$$

Thus, the desired risk bound follows.

Step 3. It is only left to estimate the number of label requests. Using the sample complexity bound of Theorem 3.1 and that for each $x \in \mathcal{X}_{\hat{f}_p}$ we request at most $\left\lceil \frac{2\log(6D/\delta)}{h^2} \right\rceil$ labels, we have that the total number of label requests is

$$n = O\left(\frac{d\theta(\varepsilon/h)}{h^2} \log^2\left(\frac{d}{\varepsilon h \delta} \right) + \frac{D}{h^2} \log\left(\frac{D}{\delta} \right) \right).$$

The claim follows. ■

Appendix E. Auxiliary results

The next result is due to [Vapnik and Chervonenkis \(1974, Theorem 12.2\)](#) presented in the form of [Boucheron et al. \(2005, Theorem 5.1\)](#).

Lemma E.1 *Let \mathcal{F} be a class of $\{0, 1\}$ -valued functions and for $\delta \in (0, 1)$, introduce*

$$\sigma^2(n, \delta) = \frac{4}{n} \left(\log \mathcal{S}_{\mathcal{F}}(2n) + \log \frac{8}{\delta} \right).$$

Then, with probability at least $1 - \delta$, for all $f \in \mathcal{F}$, it holds that

$$P_n f - P f \leq \min \left\{ \sigma^2(n, \delta) + \sigma(n, \delta) \sqrt{P f}, \sigma(n, \delta) \sqrt{P_n f} \right\}$$

and

$$P f - P_n f \leq \min \left\{ \sigma^2(n, \delta) + \sigma(n, \delta) \sqrt{P_n f}, \sigma(n, \delta) \sqrt{P f} \right\}.$$

Lemma E.2 *Let \mathcal{F} be a class of $\{0, 1\}$ -valued functions with VC dimension d . Fix $\delta \in (0, 1)$. Let*

$$\beta^2(n, \delta) = \frac{4}{n} \left(2d \log \frac{e(2n \vee d)}{d} + \log \frac{24}{\delta} \right).$$

Then, the following inequalities hold simultaneously, with probability at least $1 - \delta$ for all $f, g \in \mathcal{F}$:

$$|R(f) - R(g) - R_n(f) + R_n(g)| \leq 2\beta^2(n, \delta) + 2\beta(n, \delta) \sqrt{P_n|f - g| \wedge P|f - g|}, \quad (14)$$

and

$$|P|f - g| - P_n|f - g|| \leq \beta^2(n, \delta) + \beta(n, \delta) \sqrt{P_n|f - g| \wedge P|f - g|}. \quad (15)$$

The result of Lemma E.2 is a simple corollary of Lemma E.1. Similar bounds are used in the proofs in (Dasgupta et al., 2008; Hsu, 2010; Zhang and Chaudhuri, 2014; Bousquet and Zhivotovskiy, 2019). We provide the proof of Lemma E.2 for the sake of completeness. We remark that in our analysis, the logarithmic factors can be improved. In particular, the results in (Giné and Koltchinskii, 2006; Zhivotovskiy and Hanneke, 2018) allow refining the logarithmic factors in Lemma E.2. Moreover, the techniques in (Hanneke and Yang, 2015, Theorem 4) give a better joint dependence on the VC dimension d and the star number s in a similar context. An adaptation of these techniques is a natural direction of future research.

Proof of Lemma E.2 By Sauer's lemma (Sauer, 1972, Theorem 1), we have

$$\log S_{\mathcal{F}}(2n) \leq d \log \frac{e(2n \vee d)}{d}.$$

Then, the inequality (15) follows from Lemma E.1 applied to the class $\mathcal{F}\Delta\mathcal{F} = \{|f - g| : f, g \in \mathcal{F}\}$ with the confidence $\delta/3$ and the fact that

$$\log S_{\mathcal{F}\Delta\mathcal{F}}(2n) \leq 2 \log S_{\mathcal{F}}(2n).$$

To prove (14), we rewrite $\mathbb{1}[f(x) \neq y] - \mathbb{1}[g(x) \neq y]$ in the form

$$\mathbb{1}[f(x) \neq y] - \mathbb{1}[g(x) \neq y] = \mathbb{1}[f(x) \neq y \text{ and } g(x) = y] - \mathbb{1}[g(x) \neq y \text{ and } f(x) = y]$$

and apply Lemma E.1 to the classes $\mathcal{F}_1 = \{(x, y) \mapsto \mathbb{1}[f(x) \neq y \text{ and } g(x) = y] : f, g \in \mathcal{F}\}$ and $\mathcal{F}_2 = \{(x, y) \mapsto \mathbb{1}[f(x) = y \text{ and } g(x) \neq y] : f, g \in \mathcal{F}\}$. It is easy to see that

$$\log S_{\mathcal{F}_1}(2n) = \log S_{\mathcal{F}_2}(2n) \leq 2 \log S_{\mathcal{F}}(2n).$$

Then, with probability at least $1 - \delta/3$, it holds that

$$\begin{aligned} & |P\mathbb{1}[f(X) \neq Y \text{ and } g(X) = Y] - P_n\mathbb{1}[f(X) \neq Y \text{ and } g(X) = Y]| \\ & \leq \beta^2(n, \delta) + \beta(n, \delta) \sqrt{P\mathbb{1}[f(X) \neq Y \text{ and } g(X) = Y] \wedge P_n\mathbb{1}[f(X) \neq Y \text{ and } g(X) = Y]} \\ & \leq \beta^2(n, \delta) + \beta(n, \delta) \sqrt{P(f - g)^2 \wedge P_n(f - g)^2}, \end{aligned}$$

and, with the same probability,

$$\begin{aligned} & |P\mathbb{1}[g(X) \neq Y \text{ and } f(X) = Y] - P_n\mathbb{1}[g(X) \neq Y \text{ and } f(X) = Y]| \\ & \leq \beta^2(n, \delta) + \beta(n, \delta) \sqrt{P\mathbb{1}[g(X) \neq Y \text{ and } f(X) = Y] \wedge P_n\mathbb{1}[g(X) \neq Y \text{ and } f(X) = Y]} \\ & \leq \beta^2(n, \delta) + \beta(n, \delta) \sqrt{P(f - g)^2 \wedge P_n(f - g)^2}. \end{aligned}$$

Then, with probability at least $1 - 2\delta/3$, we have

$$|R(f) - R(g) - R_n(f) + R_n(g)| \leq 2\beta^2(n, \delta) + 2\beta(n, \delta)\sqrt{P(f-g)^2 \wedge P_n(f-g)^2}.$$

Finally, the union bound concludes the proof. \blacksquare

Our next result provides a similar uniform bound for Chow's risk.

Lemma E.3 *Let \mathcal{F} be a class of $\{0, 1\}$ -valued functions with VC-dimension d and let*

$$\mathcal{G} = \frac{\mathcal{F} + \mathcal{F}}{2} = \left\{ \frac{f_1 + f_2}{2} : f_1, f_2 \in \mathcal{F} \right\}.$$

Denote

$$\gamma^2(n, \delta) = \frac{4}{n} \left(3d \log \frac{e(2n \vee d)}{d} + \log \frac{32}{\delta} \right).$$

Then, with probability at least $1 - \delta$, for all $f \in \mathcal{F}$, $g \in \mathcal{G}$, it holds that

$$|R^p(f) - R^p(g) - R_n^p(f) + R_n^p(g)| \leq 4\gamma^2(n, \delta) + 8\gamma(n, \delta)\sqrt{P_n(f-g)^2 \wedge P(f-g)^2}.$$

Proof Recall the definition (10) of the ℓ^p loss. We have $R^p(f) = P\ell^p(Y, f(X))$ and $R_n^p(f) = P_n\ell^p(Y, f(X))$. For any $f \in \mathcal{F}$, $g \in \mathcal{G}$, and for any $y \in \{0, 1\}$, $x \in \mathcal{X}$, it holds that

$$\begin{aligned} & \ell^p(y, g(x)) - \ell^p(y, f(x)) \\ &= \mathbb{1}[g(x) = y \text{ and } f(x) \neq y] - \mathbb{1}[g(x) = 1 - y \text{ and } f(x) = y] \\ & \quad + \left(\frac{1}{2} - p\right) \mathbb{1}\left[g(x) = \frac{1}{2} \text{ and } f(x) = y\right] - \left(\frac{1}{2} + p\right) \mathbb{1}\left[g(x) = \frac{1}{2} \text{ and } f(x) \neq y\right]. \end{aligned}$$

Apply Lemma E.1 to each term in the right hand side. First, Sauer-Shelah lemma yields

$$\log S_{\mathcal{G}}(2n) \leq 2 \log S_{\mathcal{F}}(2n) \leq 2d \log \frac{e(2n \vee d)}{d}.$$

Define the classes

$$\begin{aligned} \mathcal{G}_1 &= \{(x, y) \mapsto \mathbb{1}[g(x) = y \text{ and } f(x) \neq y] : g \in \mathcal{G}, f \in \mathcal{F}\}, \\ \mathcal{G}_2 &= \{(x, y) \mapsto \mathbb{1}[g(x) = 1 - y \text{ and } f(x) = y] : g \in \mathcal{G}, f \in \mathcal{F}\}, \\ \mathcal{G}_3 &= \{(x, y) \mapsto \mathbb{1}[g(x) = 1/2 \text{ and } f(x) = y] : g \in \mathcal{G}, f \in \mathcal{F}\}, \\ \mathcal{G}_4 &= \{(x, y) \mapsto \mathbb{1}[g(x) = 1/2 \text{ and } f(x) \neq y] : g \in \mathcal{G}, f \in \mathcal{F}\}. \end{aligned}$$

It holds that

$$\max_{1 \leq k \leq 4} \log S_{\mathcal{G}_k}(2n) \leq \log S_{\mathcal{G}}(2n) + \log S_{\mathcal{F}}(2n) \leq 3d \log \frac{e(2n \vee d)}{d}.$$

Using $\mathbb{1}[g(x) = y \text{ and } f(x) \neq y] \leq \mathbb{1}[g(x) \neq f(x)] \leq 4(f(x) - g(x))^2$, we deduce that, with probability at least $1 - \delta/4$, it holds that

$$\begin{aligned} & |P\mathbb{1}[g(X) = Y \text{ and } f(X) \neq Y] - P_n\mathbb{1}[g(X) = Y \text{ and } f(X) \neq Y]| \\ & \leq \gamma^2(n, \delta) + 2\gamma(n, \delta)\sqrt{P_n(f(X) - g(X))^2 \wedge P(f(X) - g(X))^2}. \end{aligned}$$

Similarly, we have, with probability at least $1 - \delta/4$,

$$\begin{aligned} & |P\mathbf{1}[g(x) = 1 - y \text{ and } f(x) = y] - P_n\mathbf{1}[g(x) = 1 - y \text{ and } f(x) = y]| \\ & \leq \gamma^2(n, \delta) + 2\gamma(n, \delta)\sqrt{P_n(f(X) - g(X))^2 \wedge P(f(X) - g(X))^2}. \end{aligned}$$

Finally, using the fact that

$$\mathbf{1}[g(x) = 1/2 \text{ and } f(x) = y] \leq 4(f(x) - g(x))^2$$

and

$$\mathbf{1}[g(x) = 1/2 \text{ and } f(x) \neq y] \leq 4(f(x) - g(x))^2,$$

we have, with probability at least $1 - \delta/4$,

$$\begin{aligned} & \left| P\mathbf{1}\left(g(x) = \frac{1}{2} \text{ and } f(x) = y\right) - P_n\mathbf{1}\left(g(x) = \frac{1}{2} \text{ and } f(x) = y\right) \right| \\ & \leq \gamma^2(n, \delta) + 2\gamma(n, \delta)\sqrt{P_n(f(X) - g(X))^2 \wedge P(f(X) - g(X))^2} \end{aligned}$$

and

$$\begin{aligned} & \left| P\mathbf{1}\left(g(x) = \frac{1}{2} \text{ and } f(x) \neq y\right) - P_n\mathbf{1}\left(g(x) = \frac{1}{2} \text{ and } f(x) \neq y\right) \right| \\ & \leq \gamma^2(n, \delta) + 2\gamma(n, \delta)\sqrt{P_n(f(X) - g(X))^2 \wedge P(f(X) - g(X))^2}. \end{aligned}$$

Hence, by the union bound, with probability at least $1 - \delta$, we have

$$|R^p(f) - R^p(g) - R_n^p(f) + R_n^p(g)| \leq 4\gamma^2(n, \delta) + 8\gamma(n, \delta)\sqrt{P_n(f - g)^2 \wedge P(f - g)^2}.$$

The proof is complete. ■