
Active Learning with Neural Networks: Insights from Nonparametric Statistics

Yinglun Zhu

Department of Computer Sciences
University of Wisconsin-Madison
Madison, WI 53706
yinglun@cs.wisc.edu

Robert Nowak

Department of Electrical and Computer Engineering
University of Wisconsin-Madison
Madison, WI 53706
rdnowak@wisc.edu

Abstract

Deep neural networks have great representation power, but typically require large numbers of training examples. This motivates deep active learning methods that can significantly reduce the amount of labeled training data. Empirical successes of deep active learning have been recently reported in the literature, however, rigorous label complexity guarantees of deep active learning have remained elusive. This constitutes a significant gap between theory and practice. This paper tackles this gap by providing the first near-optimal label complexity guarantees for deep active learning. The key insight is to study deep active learning from the nonparametric classification perspective. Under standard low noise conditions, we show that active learning with neural networks can provably achieve the minimax label complexity, up to disagreement coefficient and other logarithmic terms. When equipped with an abstention option, we further develop an efficient deep active learning algorithm that achieves $\text{polylog}(\frac{1}{\epsilon})$ label complexity, without any low noise assumptions. We also provide extensions of our results beyond the commonly studied Sobolev/Hölder spaces and develop label complexity guarantees for learning in Radon BV^2 spaces, which have recently been proposed as natural function spaces associated with neural networks.

1 Introduction

We study active learning with neural network hypothesis classes, sometimes known as *deep active learning*. Active learning agent proceeds by selecting the most informative data points to label: The goal of active learning is to achieve the same accuracy achievable by passive learning, but with much fewer label queries (Settles, 2009; Hanneke, 2014). When the hypothesis class is a set of neural networks, the learner further benefits from the representation power of deep neural networks, which has driven the successes of passive learning in the past decade (Krizhevsky et al., 2012; LeCun et al., 2015). With these added benefits, deep active learning has become a popular research area, with empirical successes observed in many recent papers (Sener and Savarese, 2018; Ash et al., 2019; Citovsky et al., 2021; Ash et al., 2021; Kothawade et al., 2021; Emam et al., 2021; Ren et al., 2021). However, due to the difficulty of analyzing a set of neural networks, rigorous label complexity guarantees for deep active learning have remained largely elusive.

To the best of our knowledge, there are only two papers (Karzand and Nowak, 2020; Wang et al., 2021) that have made the attempts at theoretically quantifying active learning gains with neural networks. While insightful views are provided, these two works have their own limitations. The guarantees provided in Karzand and Nowak (2020) only work in the $1d$ case where data points are uniformly sampled from $[0, 1]$ and labeled by a well-separated piece-wise constant function in a noise-free way (i.e., without any labeling noise). Wang et al. (2021) study deep active learning by

linearizing the neural network at its random initialization and then analyzing it as a linear function; moreover, as the authors agree, their error bounds and label complexity guarantees can in fact be *vacuous* in certain cases. Thus, it’s fair to say that up to now researchers have not identified cases where deep active learning are provably near minimax optimal (or even with provably non-vacuous guarantees), which constitutes a significant gap between theory and practice.

In this paper, we bridge this gap by providing the first near-optimal label complexity guarantees for deep active learning. We obtain insights from the nonparametric setting where the conditional probability (of taking a positive label) is assumed to be a smooth function (Tsybakov, 2004; Audibert and Tsybakov, 2007). Previous nonparametric active learning algorithms proceed by partitioning the action space into exponentially many sub-regions (e.g., partitioning the unit cube $[0, 1]^d$ into ε^{-d} sub-cubes each with volume ε^d), and then conducting local mean (or some higher-order statistics) estimation within each sub-region (Castro and Nowak, 2008; Minsker, 2012; Locatelli et al., 2017, 2018; Shekhar et al., 2021; Kpotufe et al., 2021). We show that, with an appropriately chosen set of neural networks that *globally* approximates the smooth regression function, one can in fact recover the minimax label complexity for active learning, up to disagreement coefficient (Hanneke, 2007, 2014) and other logarithmic factors. Our results are established by (i) identifying the “right tools” to study neural networks (ranging from approximation results (Yarotsky, 2017, 2018) to complexity measure of neural networks (Bartlett et al., 2019)), and (ii) developing novel extensions of agnostic active learning algorithms (Balcan et al., 2006; Hanneke, 2007, 2014) to work with a set of neural networks.

While matching the minimax label complexity in nonparametric active learning is existing, such minimax results scale as $\Theta(\text{poly}(\frac{1}{\varepsilon}))$ (Castro and Nowak, 2008; Locatelli et al., 2017) and do not resemble what is practically observed in deep active learning: A fairly accurate neural network classifier can be obtained by training with only a few labeled data points. Inspired by recent results in *parametric* active learning with abstention (Puchkin and Zhivotovskiy, 2021; Zhu and Nowak, 2022), we develop an oracle-efficient algorithm showing that deep active learning provably achieves $\text{polylog}(\frac{1}{\varepsilon})$ label complexity when equipped with an abstention option (Chow, 1970). Our algorithm not only achieves an exponential saving in label complexity (*without any low noise assumptions*), but is also highly practical: In real-world scenarios such as medical imaging, it makes more sense for the classifier to abstain from making prediction on hard examples (e.g., those that are close to the boundary), and ask medical experts to make the judgments.

1.1 Problem setting

Let \mathcal{X} denote the instance space and \mathcal{Y} denote the label space. We focus on the binary classification problem where $\mathcal{Y} := \{+1, -1\}$. The joint distribution over $\mathcal{X} \times \mathcal{Y}$ is denoted as $\mathcal{D}_{\mathcal{X}\mathcal{Y}}$. We use $\mathcal{D}_{\mathcal{X}}$ to denote the marginal distribution over the instance space \mathcal{X} , and use $\mathcal{D}_{\mathcal{Y}|x}$ to denote the conditional distribution of \mathcal{Y} with respect to any $x \in \mathcal{X}$. We consider the standard active learning setup where $x \sim \mathcal{D}_{\mathcal{X}}$ but its label $y \sim \mathcal{D}_{\mathcal{Y}|x}$ is only observed after issuing a label query. We define $\eta(x) := \mathbb{P}_{y \sim \mathcal{D}_{\mathcal{Y}|x}}(y = +1)$ as the conditional probability of taking a positive label. The Bayes optimal classifier h^* can thus be expressed as $h^*(x) := \text{sign}(2\eta(x) - 1)$. For any classifier $h : \mathcal{X} \rightarrow \mathcal{Y}$, its (standard) error is calculated as $\text{err}(h) := \mathbb{P}_{(x,y) \sim \mathcal{D}_{\mathcal{X}\mathcal{Y}}}(h(x) \neq y)$; and its (standard) excess error is defined as $\text{excess}(h) := \text{err}(h) - \text{err}(h^*)$. Our goal is to learn an accurate classifier with a small number of label querying.

The nonparametric setting. We consider the nonparametric setting where the conditional probability η is characterized by a smooth function. Fix any $\alpha \in \mathbb{N}_+$, the *Sobolev norm* of a function $f : \mathcal{X} \rightarrow \mathbb{R}$ is defined as $\|f\|_{\mathcal{W}^{\alpha,\infty}} := \max_{|\bar{\alpha}| \leq \alpha} \text{ess sup}_{x \in \mathcal{X}} |D^{\bar{\alpha}} f(x)|$, where $\alpha = (\alpha_1, \dots, \alpha_d)$, $|\alpha| = \sum_{i=1}^d \alpha_i$ and $D^{\alpha} f$ denotes the standard α -th weak derivative of f . The unit ball in the Sobolev space is defined as $\mathcal{W}_1^{\alpha,\infty}(\mathcal{X}) := \{f : \|f\|_{\mathcal{W}^{\alpha,\infty}} \leq 1\}$. Following the convention of nonparametric active learning (Castro and Nowak, 2008; Minsker, 2012; Locatelli et al., 2017, 2018; Shekhar et al., 2021; Kpotufe et al., 2021), we assume $\mathcal{X} = [0, 1]^d$ and $\eta \in \mathcal{W}_1^{\alpha,\infty}(\mathcal{X})$ (except in Section 4).

Neural Networks. We consider *feedforward neural networks* with Rectified Linear Unit (ReLU) activation function, which is defined as $\text{ReLU}(x) := \max\{x, 0\}$. Each neural network $f_{\text{dnn}} : \mathcal{X} \rightarrow \mathbb{R}$ consists of several input units (which corresponds to the covariates of $x \in \mathcal{X}$), one output unit (which corresponds to the prediction in \mathbb{R}), and multiple hidden computational units. Each hidden

computational unit takes inputs $\{\bar{x}_i\}_{i=1}^N$ (which are outputs from previous layers) and perform the computation $\text{ReLU}(\sum_{i=1}^N w_i \bar{x}_i + b)$ with *adjustable* parameters $\{w_i\}_{i=1}^N$ and b ; the output unit performs the same operation, but without the ReLU nonlinearity. We use W to denote the total number of parameters of a neural network, and L to denote the depth of the neural network.

1.2 Contributions and paper organization

Neural networks are known to be universal approximators (Cybenko, 1989; Hornik, 1991). In this paper, we argue that, in both passive and active regimes, the universal approximability makes neural networks “universal classifiers” for classification problems: With an appropriately chosen set of neural networks, one can recover known minimax rates (up to disagreement coefficients in the active setting) in the rich nonparametric regimes.¹ We provide informal statements of our main results in the sequel, with detailed statements and associated definitions/algorithms deferred to later sections.

In Section 2, we analyze the label complexity of deep active learning under the standard Tsybakov noise condition with smoothness parameter $\beta \geq 0$ (Tsybakov, 2004). Let \mathcal{H}_{dnn} be an appropriately chosen set of neural network classifiers and denote $\theta_{\mathcal{H}_{\text{dnn}}}(\varepsilon)$ as the disagreement coefficient (Hanneke, 2007, 2014) at level ε . We develop the following label complexity guarantees for deep active learning.

Theorem 1 (Informal). *There exists an algorithm that returns a neural network classifier $\hat{h} \in \mathcal{H}_{\text{dnn}}$ with excess error $\tilde{O}(\varepsilon)$ after querying $\tilde{O}(\theta_{\mathcal{H}_{\text{dnn}}}(\varepsilon^{\frac{\beta}{1+\beta}}) \cdot \varepsilon^{-\frac{d+2\alpha}{\alpha+\alpha\beta}})$ labels.*

The label complexity presented in Theorem 1 matches the active learning lower bound $\Omega(\varepsilon^{-\frac{d+2\alpha}{\alpha+\alpha\beta}})$ (Locatelli et al., 2017) up to the dependence on the disagreement coefficient (and other logarithmic factors). Since $\theta_{\mathcal{H}_{\text{dnn}}}(\varepsilon) \leq \varepsilon^{-1}$ by definition, the label complexity presented in Theorem 1 is never worse than the passive learning rates $\tilde{\Theta}(\varepsilon^{-\frac{d+2\alpha+\alpha\beta}{\alpha+\alpha\beta}})$ (Audibert and Tsybakov, 2007). We also discover conditions under which the disagreement coefficient with respect to a set of neural network classifiers can be properly bounded, i.e., $\theta_{\mathcal{H}_{\text{dnn}}}(\varepsilon) = o(\varepsilon^{-1})$ (implying strict improvement over passive learning) and $\theta_{\mathcal{H}_{\text{dnn}}}(\varepsilon) = o(1)$ (implying matching active learning lower bound).

In Section 3, we develop label complexity guarantees for deep active learning when an additional abstention option is allowed (Chow, 1970; Puchkin and Zhivotovskiy, 2021; Zhu and Nowak, 2022). Suppose a cost (e.g. 0.49) that is marginally smaller than random guessing (which has expected cost 0.5) is incurred whenever the classifier abstains from making a predication, we develop the following label complexity guarantees for deep active learning.

Theorem 2 (Informal). *There exists an efficient algorithm that constructs a neural network classifier \hat{h}_{dnn} with Chow’s excess error $\tilde{O}(\varepsilon)$ after querying $\text{polylog}(\frac{1}{\varepsilon})$ labels.*

The above $\text{polylog}(\frac{1}{\varepsilon})$ label complexity bound is achieved *without any low noise assumptions*. Such exponential label savings theoretically justify the great empirical performances of deep active learning observed in practice (e.g., in Sener and Savarese (2018)): It suffices to label a few data points to achieve a high accuracy level. Moreover, apart from an initialization step, our algorithm (Algorithm 4) developed for Theorem 2 can be *efficiently* implemented in $\tilde{O}(\varepsilon^{-1})$ time, given a convex loss regression oracle over an appropriately chosen set of neural networks; in practice, the regression oracle can be approximated by running stochastic gradient descent.

Technical contributions. Besides identifying the “right tools” (ranging from approximation results (Yarotsky, 2017, 2018) to complexity analyses (Bartlett et al., 2019)) to analyze deep active learning, our theoretical guarantees are empowered by novel extensions of active learning algorithms *under neural network approximations*. In particular, we deal with approximation error in active learning under Tsybakov noise, and identify conditions that greatly relax the approximation requirement in the learning with abstention setup; we also analyze the disagreement coefficient, both classifier-based and value function-based, with a set of neural networks. These analyses together lead to our main results for deep active learning (e.g., Theorem 1 and Theorem 2). More generally, we establish a

¹As a byproduct, our results also provide a new perspective on nonparametric active learning through the lens of neural network approximations. Nonparametric active learning was previously tackled through space partitioning and local estimations over exponentially many sub-regions (Castro and Nowak, 2008; Minsker, 2012; Locatelli et al., 2017, 2018; Shekhar et al., 2021; Kpotufe et al., 2021).

bridge between approximation theory and active learning; we provide these general guarantees in [Appendix B](#) (under Tsybakov noise) and [Appendix D](#) (with the abstention option), which can be of independent interests. Benefited from these generic algorithms and guarantees, in [Section 4](#), we extend our results into learning smooth functions in the Radon BV^2 space ([Ongie et al., 2020](#); [Parhi and Nowak, 2021, 2022a,b](#); [Unser, 2022](#)), which is recently proposed as a natural space to analyze neural networks.

1.3 Related work

Active learning concerns about learning accurate classifiers without extensive human labeling. One of the earliest work of active learning dates back to the CAL algorithm proposed by [Cohn et al. \(1994\)](#), which set the cornerstone for *disagreement-based* active learning. Since then, a long line of work have been developed, either directly working with a set classifier ([Balcan et al., 2006](#); [Hanneke, 2007](#); [Dasgupta et al., 2007](#); [Beygelzimer et al., 2009, 2010](#); [Huang et al., 2015](#); [Cortes et al., 2019](#)) or work with a set of regression functions ([Krishnamurthy et al., 2017, 2019](#)). These work mainly focus on the parametric regime (e.g., learning with a set of linear classifiers), and their label complexities rely on the boundedness of the so-called disagreement coefficient ([Hanneke, 2007, 2014](#); [Friedman, 2009](#)). Active learning in the nonparametric regime has been analyzed in [Castro and Nowak \(2008\)](#); [Minsker \(2012\)](#); [Locatelli et al. \(2017, 2018\)](#); [Kpotufe et al. \(2021\)](#). These algorithms rely on partitioning of the input space $\mathcal{X} \subseteq [0, 1]^d$ into exponentially (in dimension) many small cubes, and then conduct local mean (or some higher-order statistics) estimation within each small cube.

It is well known that, in the worst case, active learning exhibits no label complexity gains over the passive counterpart ([Kääriäinen, 2006](#)). To bypass these worst-case scenarios, active learning has been popularly analyzed under the so-called Tsybakov low noise conditions ([Tsybakov, 2004](#)). Under Tsybakov noise conditions, active learning has been shown to be strictly superior than passive learning in terms of label complexity ([Castro and Nowak, 2008](#); [Locatelli et al., 2017](#)). Besides analyzing active learning under favorable low noise assumptions, more recently, researchers consider active learning with an abstention option and analyze its label complexity under Chow’s error ([Chow, 1970](#)). In particular, [Puchkin and Zhivotovskiy \(2021\)](#); [Zhu and Nowak \(2022\)](#) develop active learning algorithms with $\text{polylog}(\frac{1}{\epsilon})$ label complexity when analyzed under Chow’s excess error. [Shekhar et al. \(2021\)](#) study nonparametric active learning under a different notion of the Chow’s excess error, and propose algorithms with $\text{poly}(\frac{1}{\epsilon})$ label complexity; their algorithms follow similar procedures of those partition-based nonparametric active learning algorithms (e.g., [Minsker \(2012\)](#); [Locatelli et al. \(2017\)](#)).

Inspired by the success of deep learning in the passive regime, active learning with neural networks has been extensively explored in recent years ([Sener and Savarese, 2018](#); [Ash et al., 2019](#); [Citovsky et al., 2021](#); [Ash et al., 2021](#); [Kothawade et al., 2021](#); [Emam et al., 2021](#); [Ren et al., 2021](#)). Great empirical performances are observed in these papers, however, rigorous label complexity guarantees have largely remains elusive (except in [Karzand and Nowak \(2020\)](#); [Wang et al. \(2021\)](#), with limitations discussed before). We bridge the gap between practice and theory by providing the first near-optimal label complexity guarantees for deep active learning. Our results are built upon approximation results of deep neural networks ([Yarotsky, 2017, 2018](#); [Parhi and Nowak, 2022b](#)) and VC/pseudo dimension analyses of neural networks with given structures ([Bartlett et al., 2019](#)).

2 Label complexity of deep active learning

We analyze the label complexity of deep active learning in this section. We first introduce the Tsybakov noise condition in [Section 2.1](#), and then identify the “right tools” to analyze classification problems with neural network classifiers in [Section 2.2](#) (where we also provide passive learning guarantees). We establish our main active learning guarantees in [Section 2.3](#).

2.1 Tsybakov noise condition

It is well known that active learning exhibits no label complexity gains over the passive counterpart without additional low noise assumptions ([Kääriäinen, 2006](#)). We next introduce the Tsybakov low noise condition ([Tsybakov, 2004](#)), which has been extensively analyzed in active learning literature.

Definition 1 (Tsybakov noise). A distribution $\mathcal{D}_{\mathcal{X}\mathcal{Y}}$ satisfies the Tsybakov noise condition with parameter $\beta \geq 0$ and a universal constant $c \geq 1$ if, $\forall \tau > 0$,

$$\mathbb{P}_{x \sim \mathcal{D}_{\mathcal{X}}}(|\eta(x) - 1/2| \leq \tau) \leq c\tau^\beta.$$

The case with $\beta = 0$ corresponds to the general case *without* any low noise conditions, where no active learning algorithm can outperform the passive counterpart (Audibert and Tsybakov, 2007; Locatelli et al., 2017). We use $\mathcal{P}(\alpha, \beta)$ to denote the set of distributions satisfying: (i) the smoothness conditions introduced in Section 1.1 with parameter $\alpha > 0$; and (ii) the Tsybakov low noise condition (i.e., Definition 1) with parameter $\beta \geq 0$. We assume $\mathcal{D}_{\mathcal{X}\mathcal{Y}} \in \mathcal{P}(\alpha, \beta)$ in the rest of Section 2. As in Castro and Nowak (2008); Hanneke (2014), we assume the knowledge of noise/smoothness parameters.

2.2 Approximation and expressiveness of neural networks

Neural networks are known to be universal approximators (Cybenko, 1989; Hornik, 1991): For any continuous function $g : \mathcal{X} \rightarrow \mathbb{R}$ and any error tolerance $\kappa > 0$, there exists a large enough neural network f_{dnn} such that $\|f_{\text{dnn}} - g\|_\infty := \sup_{x \in \mathcal{X}} |f_{\text{dnn}}(x) - g(x)| \leq \kappa$. Recently, *non-asymptotic* approximation rates by ReLU neural networks have been developed for smooth functions in the Sobolev space, which we restate in the following.²

Theorem 3 (Yarotsky (2017)). Fix any $\kappa > 0$. For any $f^* = \eta \in \mathcal{W}_1^{\alpha, \infty}([0, 1]^d)$, there exists a neural network f_{dnn} with $W = O(\kappa^{-\frac{d}{\alpha}} \log \frac{1}{\kappa})$ total number of parameters arranged in $L = O(\log \frac{1}{\kappa})$ layers such that $\|f_{\text{dnn}} - f^*\|_\infty \leq \kappa$.

The architecture of the neural network f_{dnn} appearing in the above theorem only depends on the smooth function space $\mathcal{W}_1^{\alpha, \infty}([0, 1]^d)$, but otherwise is independent of the true regression function f^* ; also see Yarotsky (2017) for details. Let \mathcal{F}_{dnn} denote the set of neural network *regression functions* with the same architecture. We construct a set of neural network *classifiers* by thresholding the regression function at $\frac{1}{2}$, i.e., $\mathcal{H}_{\text{dnn}} := \{h_f := \text{sign}(2f(x) - 1) : f \in \mathcal{F}_{\text{dnn}}\}$. The next result concerns about the expressiveness of the neural network classifiers, in terms of a well-known complexity measure: the VC dimension (Vapnik and Chervonenkis, 1971).

Theorem 4 (Bartlett et al. (2019)). Let \mathcal{H}_{dnn} be a set of neural network classifiers of the same architecture and with W parameters arranged in L layers. We then have

$$\Omega(WL \log(W/L)) \leq \text{VCdim}(\mathcal{H}_{\text{dnn}}) \leq O(WL \log(W)).$$

With these tools, we can construct a set of neural network classifiers \mathcal{H}_{dnn} such that (i) the best in-class classifier $\tilde{h} \in \mathcal{H}_{\text{dnn}}$ has small excess error, and (ii) \mathcal{H}_{dnn} has a well-controlled VC dimension that is proportional to smooth/noise parameters. More specifically, we have the following proposition.

Proposition 1. Suppose $\mathcal{D}_{\mathcal{X}\mathcal{Y}} \in \mathcal{P}(\alpha, \beta)$. One can construct a set of neural network classifier \mathcal{H}_{dnn} such that the following two properties hold simultaneously:

$$\inf_{h \in \mathcal{H}_{\text{dnn}}} \text{err}(h) - \text{err}(h^*) = O(\varepsilon) \quad \text{and} \quad \text{VCdim}(\mathcal{H}_{\text{dnn}}) = \tilde{O}(\varepsilon^{-\frac{d}{\alpha(1+\beta)}}).$$

With the approximation results obtained above, to learn a classifier with $O(\varepsilon)$ excess error, one only needs to focus on a set of neural networks \mathcal{H}_{dnn} with a well-controlled VC dimension. As a warm-up, we first analyze the label complexity of such procedure in the passive regime (with fast rates).

Theorem 5. Suppose $\mathcal{D}_{\mathcal{X}\mathcal{Y}} \in \mathcal{P}(\alpha, \beta)$. Fix any $\varepsilon, \delta > 0$. Let \mathcal{H}_{dnn} be the set of neural network classifiers constructed in Proposition 1. With $n = \tilde{O}(\varepsilon^{-\frac{d+2\alpha+\alpha\beta}{\alpha(1+\beta)}})$ i.i.d. sampled points, with probability at least $1 - \delta$, the empirical risk minimizer $\hat{h} \in \mathcal{H}_{\text{dnn}}$ achieves excess error $O(\varepsilon)$.

The label complexity results obtained in Theorem 5 matches, up to logarithmic factors, the passive learning lower bound $\Omega(\varepsilon^{-\frac{d+2\alpha+\alpha\beta}{\alpha(1+\beta)}})$ established in Audibert and Tsybakov (2007), indicating that our proposed learning procedure *with a set of neural networks* is near minimax optimal.³

²As in Yarotsky (2017), we hide constants that are potentially α -dependent and d -dependent into the Big-Oh notation.

³Similar passive learning guarantees have been developed with different tools and analyses, e.g., see results in Kim et al. (2021).

2.3 Deep active learning and guarantees

The passive learning procedure presented in the previous section treats every data point equally, i.e., it requests the label of every data point. Active learning reduces the label complexity by only querying labels of data points that are “more important”. We present deep active learning results in this section. Our algorithm ([Algorithm 1](#)) is inspired by RobustCAL ([Balcan et al., 2006](#); [Hanneke, 2007, 2014](#)) and the seminal CAL algorithm ([Cohn et al., 1994](#)); we call our algorithm NeuralCAL to emphasize that it works with a set of neural networks.

For any accuracy level $\varepsilon > 0$, NeuralCAL first initialize a set of neural network classifiers $\mathcal{H}_0 := \mathcal{H}_{\text{dnn}}$ such that (i) the best in-class classifier $\check{h} := \arg \min_{h \in \mathcal{H}_{\text{dnn}}} \text{err}(h)$ has excess error at most $O(\varepsilon)$, and (ii) the VC dimension of \mathcal{H}_{dnn} is upper bounded by $\tilde{O}(\varepsilon^{-\frac{\alpha}{\alpha(1+\beta)}})$ (see [Section 2.2](#) for more details). NeuralCAL then runs in epochs of geometrically increasing lengths. At the beginning of epoch m , based on previously *labeled* data points, NeuralCAL updates a set of active classifier \mathcal{H}_m such that, with high probability, the best classifier \check{h} remains *uneliminated*. Within each epoch m , NeuralCAL only queries the label y of a data point x if it lies in the *region of disagreement* with respect to the current active set of classifier \mathcal{H}_m , i.e., $\text{DIS}(\mathcal{H}_m) := \{x \in \mathcal{X} : \exists h_1, h_2 \in \mathcal{H}_m \text{ s.t. } h_1(x) \neq h_2(x)\}$. NeuralCAL returns any classifier $\hat{h} \in \mathcal{H}_m$ that remains uneliminated after $M - 1$ epoch.

Algorithm 1 NeuralCAL

Input: Accuracy level $\varepsilon \in (0, 1)$, confidence level $\delta \in (0, 1)$.

- 1: Let \mathcal{H}_{dnn} be a set of neural networks classifiers constructed in [Proposition 1](#).
 - 2: Define $T := \varepsilon^{-\frac{2+\beta}{1+\beta}} \cdot \text{VCdim}(\mathcal{H}_{\text{dnn}})$, $M := \lceil \log_2 T \rceil$, $\tau_m := 2^m$ for $m \geq 1$ and $\tau_0 := 0$.
 - 3: Define $\rho_m := O\left(\frac{\text{VCdim}(\mathcal{H}_{\text{dnn}}) \cdot \log(\tau_{m-1}) \cdot \log(M/\delta)}{\tau_{m-1}}\right)^{\frac{1+\beta}{2+\beta}}$ for $m \geq 2$ and $\rho_1 := 1$.
 - 4: Define $\hat{R}_m(h) := \sum_{t=1}^{\tau_{m-1}} Q_t \mathbb{1}(h(x_t) \neq y_t)$ with the convention that $\sum_{t=1}^0 \dots = 0$.
 - 5: Initialize $\mathcal{H}_0 := \mathcal{H}_{\text{dnn}}$.
 - 6: **for** epoch $m = 1, 2, \dots, M$ **do**
 - 7: Update active set $\mathcal{H}_m := \left\{ h \in \mathcal{H}_{m-1} : \hat{R}_m(h) \leq \inf_{h \in \mathcal{H}_{m-1}} \hat{R}_m(h) + \tau_{m-1} \cdot \rho_m \right\}$
 - 8: **if** epoch $m = M$ **then**
 - 9: **Return** any classifier $\hat{h} \in \mathcal{H}_M$.
 - 10: **for** time $t = \tau_{m-1} + 1, \dots, \tau_m$ **do**
 - 11: Observe $x_t \sim \mathcal{D}_{\mathcal{X}}$. Set $Q_t := \mathbb{1}(x_t \in \text{DIS}(\mathcal{H}_m))$.
 - 12: **if** $Q_t = 1$ **then**
 - 13: Query the label y_t of x_t .
-

Since NeuralCAL only queries labels of data points lying in the region of disagreement, its label complexity should intuitively be related to how fast the region of disagreement shrinks. More formally, the rate of collapse of the (probability measure of) region of disagreement is captured by the (*classifier-based*) *disagreement coefficient* ([Hanneke, 2007, 2014](#)), which we introduce next.

Definition 2 (Classifier-based disagreement coefficient). *For any ε_0 and classifier $h \in \mathcal{H}$, the classifier-based disagreement coefficient of h is defined as*

$$\theta_{\mathcal{H}, h}(\varepsilon_0) := \sup_{\varepsilon > \varepsilon_0} \frac{\mathbb{P}_{x \sim \mathcal{D}_{\mathcal{X}}}(\text{DIS}(\mathcal{B}_{\mathcal{H}}(h, \varepsilon)))}{\varepsilon} \vee 1,$$

where $\mathcal{B}_{\mathcal{H}}(h, \varepsilon) := \{g \in \mathcal{H} : \mathbb{P}(x \in \mathcal{X} : g(x) \neq h(x)) \leq \varepsilon\}$. We also define $\theta_{\mathcal{H}}(\varepsilon_0) := \sup_{h \in \mathcal{H}} \theta_{\mathcal{H}, h}(\varepsilon_0)$.

The guarantees of NeuralCAL follows from a more general analysis of RobustCAL under approximation. In particular, to achieve fast rates (under Tsybakov noise), previous analysis of RobustCAL requires that the Bayes classifier is in the class (or a Bernstein condition for every $h \in \mathcal{H}$) ([Hanneke, 2014](#)). These requirements are stronger compared to what we have in the case with neural network approximations. Our analysis extends the understanding of RobustCAL under approximation. We defer such general analysis to [Appendix B](#), and present the following guarantees.

Theorem 6. *Suppose $\mathcal{D}_{\mathcal{X}Y} \in \mathcal{P}(\alpha, \beta)$. Fix any $\varepsilon, \delta > 0$. With probability at least $1 - \delta$, [Algorithm 1](#) returns a classifier $\hat{h} \in \mathcal{H}_{\text{dnn}}$ with excess error $\tilde{O}(\varepsilon)$ after querying $\tilde{O}(\theta_{\mathcal{H}_{\text{dnn}}}(\varepsilon^{\frac{\beta}{1+\beta}}) \cdot \varepsilon^{-\frac{d+2\alpha}{\alpha+\beta}})$ labels.*

We next discuss in detail the label complexity of deep active learning proved in [Theorem 6](#).

- Ignoring the dependence on disagreement coefficient, the label complexity appearing in [Theorem 6](#) matches, up to logarithmic factors, the lower bound $\Omega(\varepsilon^{-\frac{d+2\alpha}{\alpha+\beta}})$ for active learning ([Locatelli et al., 2017](#)). At the same time, the label complexity appearing in [Theorem 6](#) is *never worse* than the passive counterpart (i.e., $\tilde{\Theta}(\varepsilon^{-\frac{d+2\alpha+\beta}{\alpha(1+\beta)}})$) since $\theta_{\mathcal{H}_{\text{dnn}}}(\varepsilon^{\frac{\beta}{1+\beta}}) \leq \varepsilon^{-\frac{\beta}{1+\beta}}$.
- We also identify cases when $\theta_{\mathcal{H}_{\text{dnn}}}(\varepsilon^{\frac{\beta}{1+\beta}}) = o(\varepsilon^{-\frac{\beta}{1+\beta}})$, indicating *strict* improvement over passive learning (e.g., when $\mathcal{D}_{\mathcal{X}}$ is supported on countably many data points), and when $\theta_{\mathcal{H}_{\text{dnn}}}(\varepsilon^{\frac{\beta}{1+\beta}}) = O(1)$, indicating matching the minimax active lower bound (e.g., when $\mathcal{D}_{\mathcal{X}\mathcal{Y}}$ satisfies conditions such as *decomposability* defined in [Definition 4](#). See [Appendix C.2](#) for detailed discussion).⁴

Our algorithm and theorems lead to the following results, which could benefit both deep active learning and nonparametric learning communities.

- **Near minimax optimal label complexity for deep active learning.** While empirical successes of deep active learning have been observed, rigorous label complexity analysis remains elusive except for two attempts made in [Karzand and Nowak \(2020\)](#); [Wang et al. \(2021\)](#). The guarantees provided in [Karzand and Nowak \(2020\)](#) only work in very special cases (i.e., data uniformly sampled from $[0, 1]$ and labeled by well-separated piece-constant functions in a noise-free way). [Wang et al. \(2021\)](#) study deep active learning in the NTK regime by linearizing the neural network at its random initialization and analyzing it as a linear function; moreover, as the authors agree, their error bounds and label complexity guarantees are *vacuous* in certain cases. On the other hand, our guarantees are minimax optimal, up to disagreement coefficient and other logarithmic factors, which bridge the gap between theory and practice in deep active learning.
- **New perspective on nonparametric learning.** Nonparametric learning of smooth functions have been mainly approached by partitioning-based methods ([Tsybakov, 2004](#); [Audibert and Tsybakov, 2007](#); [Castro and Nowak, 2008](#); [Minsker, 2012](#); [Locatelli et al., 2017, 2018](#); [Kpotufe et al., 2021](#)): Partition the unit cube $[0, 1]^d$ into exponentially (in dimension) many sub-cubes and conduct local mean estimation within each sub-cube (which additionally requires a strictly stronger membership querying oracle). Our results show that, in both passive and active settings, one can learn *globally* with a set of neural networks and achieve near minimax optimal label complexities.

3 Deep active learning with abstention: Exponential speedups

While the theoretical guarantees provided in [Section 2](#) are near minimax optimal, the label complexity scales as $\text{poly}(\frac{1}{\varepsilon})$, which doesn't match the great empirical performance observed in deep active learning. In this section, we fill in this gap by leveraging the idea of abstention and provide a deep active learning algorithm that achieves exponential label savings. We introduce the concepts of abstention and Chow's excess error in [Section 3.1](#), and provide our label complexity guarantees in [Section 3.2](#).

3.1 Active learning without low noise conditions

The previous section analyzes active learning under Tsybakov noise, which has been extensively studied in the literature since [Castro and Nowak \(2008\)](#). More recently, promising results are observed in active learning under Chow's excess error, but otherwise *without any low noise assumption* ([Puchkin and Zhivotovskiy, 2021](#); [Zhu and Nowak, 2022](#)). We introduce this setting in the following.

Abstention and Chow's error ([Chow, 1970](#)). We consider classifier of the form $\hat{h} : \mathcal{X} \rightarrow \mathcal{Y} \cup \{\perp\}$ where \perp denotes the action of abstention. For any fixed $0 < \gamma < \frac{1}{2}$, the Chow's error is defined as

$$\text{err}_{\gamma}(\hat{h}) := \mathbb{P}_{(x,y) \sim \mathcal{D}_{\mathcal{X}\mathcal{Y}}}(\hat{h}(x) \neq y, \hat{h}(x) \neq \perp) + (1/2 - \gamma) \cdot \mathbb{P}_{(x,y) \sim \mathcal{D}_{\mathcal{X}\mathcal{Y}}}(\hat{h}(x) = \perp).$$

⁴We remark that disagreement coefficient is usually bounded/analyzed under additional assumptions on $\mathcal{D}_{\mathcal{X}\mathcal{Y}}$, even for simple cases with a set of linear classifiers ([Friedman, 2009](#); [Hanneke, 2014](#)). The label complexity guarantees of partition-based nonparametric active algorithms (e.g., [Castro and Nowak \(2008\)](#)) do not depend on the disagreement coefficient, but they are analyzed under stronger assumptions, e.g., they require the strictly stronger membership querying oracle. See [Wang \(2011\)](#) for a discussion. We left a comprehensive analysis of the disagreement coefficient with a set of neural network classifiers for future work.

The parameter γ can be chosen as a small constant, e.g., $\gamma = 0.01$, to avoid excessive abstention: The price of abstention is only marginally smaller than random guess (which incurs cost 0.5). The *Chow’s excess error* is then defined as $\text{excess}_\gamma(\hat{h}) := \text{err}_\gamma(\hat{h}) - \text{err}(h^*)$ (Puchkin and Zhivotovskiy, 2021).

At a high level, analyzing with Chow’s excess error allows slackness in predications of hard examples (e.g., data points whose $\eta(x)$ is close to $\frac{1}{2}$) by leveraging the power of abstention. Puchkin and Zhivotovskiy (2021); Zhu and Nowak (2022) show that $\text{polylog}(\frac{1}{\varepsilon})$ is always achievable in the *parametric* settings. We generalize their results to the *nonparametric* setting and analyze active learning with a set of neural networks.

3.2 Exponential speedups with abstention

In this section, we work with a set of neural network *regression functions* $\mathcal{F}_{\text{dnn}} : \mathcal{X} \rightarrow [0, 1]$ (that approximates η) and then *construct* classifiers $h : \mathcal{X} \rightarrow \mathcal{Y} \cup \{\perp\}$ with an *additional abstention action*. To work with a set of regression functions \mathcal{F}_{dnn} , we analyze its “complexity” from the lenses of *pseudo dimension* $\text{Pdim}(\mathcal{F}_{\text{dnn}})$ (Pollard, 1984; Haussler, 1989, 1995) and *value function disagreement coefficient* $\theta_{\mathcal{F}_{\text{dnn}}}^{\text{val}}(\iota)$ (for some $\iota > 0$) (Foster et al., 2020). We defer detailed definitions of these complexity measures to [Appendix D.1](#).

Algorithm 2 NeuralCAL++

Input: Accuracy level $\varepsilon \in (0, 1)$, confidence level $\delta \in (0, 1)$, abstention parameter $\gamma \in (0, 1/2)$.

1: Let \mathcal{F}_{dnn} be a set of neural network regression functions obtained by (i) applying [Theorem 3](#) with an appropriate approximation level κ (which satisfies $\frac{1}{\kappa} = \text{poly}(\frac{1}{\gamma}) \text{polylog}(\frac{1}{\varepsilon\gamma})$), and (ii) applying a preprocessing step on the set of neural networks obtained from step (i). See [Appendix E](#) for details.

2: Define $T := \frac{\theta_{\mathcal{F}_{\text{dnn}}}^{\text{val}}(\gamma/4) \cdot \text{Pdim}(\mathcal{F}_{\text{dnn}})}{\varepsilon\gamma}$, $M := \lceil \log_2 T \rceil$, and $C_\delta := O(\text{Pdim}(\mathcal{F}_{\text{dnn}}) \cdot \log(T/\delta))$.

3: Define $\tau_m := 2^m$ for $m \geq 1$, $\tau_0 := 0$, and $\beta_m := 3(M - m + 1)C_\delta$.

4: Define $\hat{R}_m(f) := \sum_{t=1}^{\tau_m-1} Q_t(f(x_t) - y_t)^2$ with the convention that $\sum_{t=1}^0 \dots = 0$.

5: **for** epoch $m = 1, 2, \dots, M$ **do**

6: Get $\hat{f}_m := \arg \min_{f \in \mathcal{F}_{\text{dnn}}} \sum_{t=1}^{\tau_m-1} Q_t(f(x_t) - y_t)^2$.

7: (Implicitly) Construct active set $\mathcal{F}_m := \{f \in \mathcal{F}_{\text{dnn}} : \hat{R}_m(f) \leq \hat{R}_m(\hat{f}_m) + \beta_m\}$.

8: Construct classifier $\hat{h}_m : \mathcal{X} \rightarrow \{+1, -1, \perp\}$ as

$$\hat{h}_m(x) := \begin{cases} \perp, & \text{if } [\text{lcb}(x; \mathcal{F}_m) - \frac{\gamma}{4}, \text{ucb}(x; \mathcal{F}_m) + \frac{\gamma}{4}] \subseteq [\frac{1}{2} - \gamma, \frac{1}{2} + \gamma]; \\ \text{sign}(2\hat{f}_m(x) - 1), & \text{o.w.} \end{cases}$$

and query function $g_m(x) := \mathbb{1}(\frac{1}{2} \in (\text{lcb}(x; \mathcal{F}_m) - \frac{\gamma}{4}, \text{ucb}(x; \mathcal{F}_m) + \frac{\gamma}{4})) \cdot \mathbb{1}(\hat{h}_m(x) \neq \perp)$.

9: **if** epoch $m = M$ **then**

10: **Return** classifier \hat{h}_M .

11: **for** time $t = \tau_{m-1} + 1, \dots, \tau_m$ **do**

12: Observe $x_t \sim \mathcal{D}_{\mathcal{X}}$. Set $Q_t := g_m(x_t)$.

13: **if** $Q_t = 1$ **then**

14: Query the label y_t of x_t .

We now present NeuralCAL++ ([Algorithm 2](#)), a deep active learning algorithm that leverages the power of abstention. NeuralCAL++ first initialize a set of set of neural network regression functions \mathcal{F}_{dnn} by applying a preprocessing step on top of the set of regression functions obtained from [Theorem 3](#) with a carefully chosen approximation level κ . The preprocessing step mainly contains two actions: (1) clipping $f_{\text{dnn}} : \mathcal{X} \rightarrow \mathbb{R}$ into $\check{f}_{\text{dnn}} : \mathcal{X} \rightarrow [0, 1]$ (since we obviously have $\eta(x) \in [0, 1]$); and (2) filtering out $f_{\text{dnn}} \in \mathcal{F}_{\text{dnn}}$ that are clearly not a good approximation of η . After initialization, NeuralCAL++ runs in epochs of geometrically increasing lengths. At the beginning of epoch $m \in [M]$, NeuralCAL++ (implicitly) constructs an active set of regression functions \mathcal{F}_m that are “close” to the true conditional probability η . For any $x \sim \mathcal{D}_{\mathcal{X}}$, NeuralCAL++ constructs a lower bound $\text{lcb}(x; \mathcal{F}_m) := \inf_{f \in \mathcal{F}_m} f(x)$ and an upper bound $\text{ucb}(x; \mathcal{F}_m) := \sup_{f \in \mathcal{F}_m} f(x)$ as a confidence range of $\eta(x)$ (based on \mathcal{F}_m). An empirical classifier with an abstention option

$\widehat{h}_m : \mathcal{X} \rightarrow \{+1, -1, \perp\}$ and a query function $g_m : \mathcal{X} \rightarrow \{0, 1\}$ are then constructed based on the confidence range (and the abstention parameter γ). For any time step t within epoch m , NeuralCAL++ queries the label of the observed data point x_t if and only if $Q_t := g_m(x_t) = 1$. NeuralCAL++ returns \widehat{h}_M as the learned classifier.

NeuralCAL++ is adapted from the algorithm developed in [Zhu and Nowak \(2022\)](#), but with novel extensions. In particular, the algorithm presented in [Zhu and Nowak \(2022\)](#) requires the existence of a $\bar{f} \in \mathcal{F}$ such that $\|\bar{f} - \eta\|_\infty \leq \varepsilon$ (to achieve ε Chow’s excess error), Such an approximation requirement directly leads to $\text{poly}(\frac{1}{\varepsilon})$ label complexity *in the nonparametric setting*, which is unacceptable. The initialization step of NeuralCAL++ (line 1) is carefully chosen to ensure that $\text{Pdim}(\mathcal{F}_{\text{dnn}}), \theta_{\mathcal{F}_{\text{dnn}}}^{\text{val}}(\frac{\gamma}{4}) = \text{poly}(\frac{1}{\gamma}) \cdot \text{polylog}(\frac{1}{\varepsilon})$; together with a sharper analysis of concentration results, these conditions help us derive the following deep active learning guarantees (also see [Appendix D](#) for a more general guarantee).

Theorem 7. *Fix any $\varepsilon, \delta, \gamma > 0$. With probability at least $1 - \delta$, [Algorithm 2](#) (with an appropriate initialization at line 1) returns a classifier \widehat{h} with Chow’s excess error $\widetilde{O}(\varepsilon)$ after querying $\text{poly}(\frac{1}{\gamma}) \cdot \text{polylog}(\frac{1}{\varepsilon\delta})$ labels.*

We discuss two important aspects of [Algorithm 2/Theorem 7](#) in the following, i.e., exponential savings and computational efficiency. We defer more detailed discussions to [Appendix F.1](#).

- **Exponential speedups.** [Theorem 7](#) shows that, equipped with an abstention option, deep active learning enjoys $\text{polylog}(\frac{1}{\varepsilon})$ label complexity. This provides theoretical justifications for great empirical results of deep active learning observed in practice. Moreover, [Algorithm 2](#) outputs a classifier that abstains *properly*, i.e., it abstains only if abstention is the optimal choice; such a property further implies $\text{polylog}(\frac{1}{\varepsilon})$ label complexity under *standard* excess error and Massart noise ([Massart and Nédélec, 2006](#)).
- **Computational efficiency.** Suppose one can efficiently implement a (weighted) square loss regression oracle over the *initialized* set of neural networks \mathcal{F}_{dnn} : Given any set \mathcal{S} of weighted examples $(w, x, y) \in \mathbb{R}_+ \times \mathcal{X} \times \mathcal{Y}$ as input, the regression oracle outputs $\widehat{f}_{\text{dnn}} := \arg \min_{f \in \mathcal{F}_{\text{dnn}}} \sum_{(w, x, y) \in \mathcal{S}} w(f(x) - y)^2$.⁵ [Algorithm 2](#) can then be *efficiently* implemented with $\text{poly}(\frac{1}{\gamma}) \cdot \frac{1}{\varepsilon}$ oracle calls.

While the label complexity obtained in [Theorem 7](#) has desired dependence on $\text{polylog}(\frac{1}{\varepsilon})$, its dependence on γ can be of order $\gamma^{-\text{poly}(d)}$. Our next result shows that, however, such dependence is unavoidable even in the case of learning a single ReLU function.

Theorem 8. *Fix any $\gamma \in (0, 1/8)$. For any accuracy level ε sufficiently small, there exists a problem instance such that (1) $\eta \in \mathcal{W}_1^{1,\infty}(\mathcal{X})$ and is of the form $\eta(x) := \text{ReLU}(\langle w, x \rangle + a) + b$; and (2) for any active learning algorithm, it takes at least $\gamma^{-\Omega(d)}$ labels to identify an ε -optimal classifier, for either standard excess error or Chow’s excess error (with parameter γ).*

4 Extensions

Previous results are developed in the commonly studied Sobolev/Hölder spaces. Our techniques, however, are generic and can be adapted to other function spaces, given neural network approximation results. In this section, we provide extensions of our results to the Radon BV^2 space, which was recently proposed as the natural function space associated with ReLU neural networks ([Ongie et al., 2020](#); [Parhi and Nowak, 2021, 2022a,b](#); [Unser, 2022](#)).⁶

The Radon BV^2 space. The Radon BV^2 unit ball over domain \mathcal{X} is defined as $\mathcal{R}\text{BV}_1^2(\mathcal{X}) := \{f : \|f\|_{\mathcal{R}\text{BV}^2(\mathcal{X})} \leq 1\}$, where $\|f\|_{\mathcal{R}\text{BV}^2(\mathcal{X})}$ denotes the Radon BV^2 norm of f over domain \mathcal{X} .⁷ Following [Parhi and Nowak \(2022b\)](#), we assume $\mathcal{X} = \{x \in \mathbb{R}^d : \|x\|_2 \leq 1\}$ and $\eta \in \mathcal{R}\text{BV}_1^2(\mathcal{X})$.

⁵In practice, one can approximate this oracle by running stochastic gradient descent.

⁶Other extensions are also possible given neural network approximation results, e.g., recent results established in [Lu et al. \(2021\)](#).

⁷We provide more mathematical backgrounds and associated definitions in [Appendix G](#).

The Radon BV^2 space naturally contains neural networks of the form $f_{\text{dnn}}(x) = \sum_{k=1}^K v_k \cdot \text{ReLU}(w_k^\top x + b_k)$. On the contrary, such f_{dnn} doesn't lie in any Sobolev space of order $\alpha \geq 2$ (since f_{dnn} doesn't have second order *weak* derivative). Thus, if η takes the form of the aforementioned neural network (e.g., $\eta = f_{\text{dnn}}$), approximating η up to κ from a Sobolev perspective requires $\tilde{O}(\kappa^{-d})$ total parameters, which suffers from the curse of dimensionality. On the other side, however, such bad dependence on dimensionality goes away when approximating from a Radon BV^2 perspective, as shown in the following theorem.

Theorem 9 (Parhi and Nowak (2022b)). *Fix any $\kappa > 0$. For any $f^* \in \mathcal{R}BV_1^2(\mathcal{X})$, there exists a one-hidden layer neural network f_{dnn} of width $K = O(\kappa^{-\frac{2d}{d+3}})$ such that $\|f^* - f_{\text{dnn}}\|_\infty \leq \kappa$.*

Equipped with this approximation result, we provide the active learning guarantees for learning a smooth function within the Radon BV^2 unit ball as follows.

Theorem 10. *Suppose $\eta \in \mathcal{R}BV_1^2(\mathcal{X})$ and the Tsybakov noise condition is satisfied with parameter $\beta \geq 0$. Fix any $\varepsilon, \delta > 0$. There exists an algorithm such that, with probability at least $1 - \delta$, it learns a classifier $\hat{h} \in \mathcal{H}_{\text{dnn}}$ with excess error $\tilde{O}(\varepsilon)$ after querying $\tilde{O}(\theta_{\mathcal{H}_{\text{dnn}}}(\varepsilon^{\frac{\beta}{1+\beta}}) \cdot \varepsilon^{-\frac{4d+6}{(1+\beta)(d+3)}})$ labels.*

Compared to the label complexity obtained in Theorem 6, the label complexity obtained in the above theorem doesn't suffer from the curse of dimensionality: For d large enough, the above label complexity scales as $\varepsilon^{-O(1)}$ yet label complexity in Theorem 6 scales as $\varepsilon^{-O(d)}$. Active learning guarantees under Chow's excess error in the Radon BV^2 space are similar to results presented in Theorem 7, and are thus deferred to Appendix G.

5 Discussion

We provide the first near-optimal deep active learning guarantees, under both standard excess error and Chow's excess error. Our results are powered by generic algorithms and analyses developed for active learning that bridge approximation guarantees into label complexity guarantees. We outline some natural directions for future research below.

- **Disagreement coefficients for neural networks.** While we have provided some results regarding the disagreement coefficients for neural networks, we believe a comprehensive investigation on this topic is needed. For instance, can we discover more general settings where the classifier-based disagreement coefficient can be upper bounded by $O(1)$? It is also interesting to explore sharper analyses on the value function disagreement coefficient.
- **Adaptivity in deep active learning.** Our current results are established with the knowledge of some problem-dependent parameters, e.g., the smoothness parameters regarding the function spaces and the noise levels. It will be interesting to see if one can develop algorithms that can automatically adapt to unknown parameters, e.g., by leveraging techniques developed in Locatelli et al. (2017, 2018).

Acknowledgments and Disclosure of Funding

The authors would like to thank Rahul Parhi for many helpful discussions regarding his papers. We also would like to thank anonymous reviewers for their constructive comments. This work is partially supported by NSF grant 1934612 and AFOSR grant FA9550-18-1-0166.

References

- Alekh Agarwal, Daniel Hsu, Satyen Kale, John Langford, Lihong Li, and Robert Schapire. Taming the monster: A fast and simple algorithm for contextual bandits. In *International Conference on Machine Learning*, pages 1638–1646. PMLR, 2014.
- Martin Anthony. Uniform glivenko-cantelli theorems and concentration of measure in the mathematical modelling of learning. *Research Report LSE-CDAM-2002-07*, 2002.
- Jordan Ash, Surbhi Goel, Akshay Krishnamurthy, and Sham Kakade. Gone fishing: Neural active learning with fisher embeddings. *Advances in Neural Information Processing Systems*, 34, 2021.

- Jordan T Ash, Chicheng Zhang, Akshay Krishnamurthy, John Langford, and Alekh Agarwal. Deep batch active learning by diverse, uncertain gradient lower bounds. *arXiv preprint arXiv:1906.03671*, 2019.
- Jean-Yves Audibert and Alexandre B Tsybakov. Fast learning rates for plug-in classifiers. *The Annals of statistics*, 35(2):608–633, 2007.
- Maria-Florina Balcan, Alina Beygelzimer, and John Langford. Agnostic active learning. In *Proceedings of the 23rd international conference on Machine learning*, pages 65–72, 2006.
- Peter L Bartlett, Nick Harvey, Christopher Liaw, and Abbas Mehrabian. Nearly-tight vc-dimension and pseudodimension bounds for piecewise linear neural networks. *The Journal of Machine Learning Research*, 20(1):2285–2301, 2019.
- Alina Beygelzimer, Sanjoy Dasgupta, and John Langford. Importance weighted active learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 49–56, 2009.
- Alina Beygelzimer, Daniel J Hsu, John Langford, and Tong Zhang. Agnostic active learning without constraints. *Advances in neural information processing systems*, 23, 2010.
- Stéphane Boucheron, Olivier Bousquet, and Gábor Lugosi. Theory of classification: A survey of some recent advances. *ESAIM: probability and statistics*, 9:323–375, 2005.
- Rui M Castro and Robert D Nowak. Minimax bounds for active learning. *IEEE Transactions on Information Theory*, 54(5):2339–2353, 2008.
- CK Chow. On optimum recognition error and reject tradeoff. *IEEE Transactions on information theory*, 16(1):41–46, 1970.
- Gui Citovsky, Giulia DeSalvo, Claudio Gentile, Lazaros Karydas, Anand Rajagopalan, Afshin Rostamizadeh, and Sanjiv Kumar. Batch active learning at scale. *Advances in Neural Information Processing Systems*, 34, 2021.
- David Cohn, Les Atlas, and Richard Ladner. Improving generalization with active learning. *Machine learning*, 15(2):201–221, 1994.
- Corinna Cortes, Giulia DeSalvo, Mehryar Mohri, Ningshan Zhang, and Claudio Gentile. Active learning with disagreement graphs. In *International Conference on Machine Learning*, pages 1379–1387. PMLR, 2019.
- George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.
- Sanjoy Dasgupta, Daniel J Hsu, and Claire Monteleoni. A general agnostic active learning algorithm. *Advances in neural information processing systems*, 20, 2007.
- Zeyad Ali Sami Emam, Hong-Min Chu, Ping-Yeh Chiang, Wojciech Czaja, Richard Leapman, Micah Goldblum, and Tom Goldstein. Active learning at the imagenet scale. *arXiv preprint arXiv:2111.12880*, 2021.
- Dylan Foster, Alekh Agarwal, Miroslav Dudík, Haipeng Luo, and Robert Schapire. Practical contextual bandits with regression oracles. In *International Conference on Machine Learning*, pages 1539–1548. PMLR, 2018.
- Dylan J Foster, Alexander Rakhlin, David Simchi-Levi, and Yunzong Xu. Instance-dependent complexity of contextual bandits and reinforcement learning: A disagreement-based perspective. *arXiv preprint arXiv:2010.03104*, 2020.
- David A Freedman. On tail probabilities for martingales. *the Annals of Probability*, pages 100–118, 1975.
- Eric Friedman. Active learning for smooth problems. In *COLT*. Citeseer, 2009.
- 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. Theory of active learning. *Foundations and Trends in Machine Learning*, 7(2-3), 2014.
- David Haussler. Decision theoretic generalizations of the pac model for neural net and other learning applications. 1989.
- David Haussler. Sphere packing numbers for subsets of the boolean n-cube with bounded vavnik-chervonenkis dimension. *Journal of Combinatorial Theory, Series A*, 69(2):217–232, 1995.

- Juha Heinonen. *Lectures on Lipschitz analysis*. Number 100. University of Jyväskylä, 2005.
- Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural networks*, 4(2): 251–257, 1991.
- Tzu-Kuo Huang, Alekh Agarwal, Daniel J Hsu, John Langford, and Robert E Schapire. Efficient and parsimonious agnostic active learning. *Advances in Neural Information Processing Systems*, 28, 2015.
- Matti Kääriäinen. Active learning in the non-realizable case. In *International Conference on Algorithmic Learning Theory*, pages 63–77. Springer, 2006.
- Mina Karzand and Robert D Nowak. Maximin active learning in overparameterized model classes. *IEEE Journal on Selected Areas in Information Theory*, 1(1):167–177, 2020.
- Yongdai Kim, Ilsang Ohn, and Dongha Kim. Fast convergence rates of deep neural networks for classification. *Neural Networks*, 138:179–197, 2021.
- Suraj Kothawade, Nathan Beck, Krishnateja Killamsetty, and Rishabh Iyer. Similar: Submodular information measures based active learning in realistic scenarios. *Advances in Neural Information Processing Systems*, 34, 2021.
- Samory Kpotufe, Gan Yuan, and Yunfan Zhao. Nuances in margin conditions determine gains in active learning. *arXiv preprint arXiv:2110.08418*, 2021.
- Akshay Krishnamurthy, Alekh Agarwal, Tzu-Kuo Huang, Hal Daumé III, and John Langford. Active learning for cost-sensitive classification. In *International Conference on Machine Learning*, pages 1915–1924. PMLR, 2017.
- Akshay Krishnamurthy, Alekh Agarwal, Tzu-Kuo Huang, Hal Daumé III, and John Langford. Active learning for cost-sensitive classification. *Journal of Machine Learning Research*, 20:1–50, 2019.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012.
- Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.
- Gene Li, Pritish Kamath, Dylan J Foster, and Nathan Srebro. Eluder dimension and generalized rank. *arXiv preprint arXiv:2104.06970*, 2021.
- Andrea Locatelli, Alexandra Carpentier, and Samory Kpotufe. Adaptivity to noise parameters in nonparametric active learning. In *Proceedings of the 2017 Conference on Learning Theory*, PMLR, 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. PMLR, 2018.
- Jianfeng Lu, Zuwei Shen, Haizhao Yang, and Shijun Zhang. Deep network approximation for smooth functions. *SIAM Journal on Mathematical Analysis*, 53(5):5465–5506, 2021.
- Pascal Massart and Élodie Nédélec. Risk bounds for statistical learning. *The Annals of Statistics*, 34(5):2326–2366, 2006.
- Stanislav Minsker. Plug-in approach to active learning. *Journal of Machine Learning Research*, 13(1), 2012.
- Greg Ongie, Rebecca Willett, Daniel Soudry, and Nathan Srebro. A function space view of bounded norm infinite width relu nets: The multivariate case. In *International Conference on Learning Representations*, 2020.
- Rahul Parhi and Robert D Nowak. Banach space representer theorems for neural networks and ridge splines. *J. Mach. Learn. Res.*, 22(43):1–40, 2021.
- Rahul Parhi and Robert D Nowak. What kinds of functions do deep neural networks learn? insights from variational spline theory. *SIAM Journal on Mathematics of Data Science*, 4(2):464–489, 2022a.
- Rahul Parhi and Robert D Nowak. Near-minimax optimal estimation with shallow relu neural networks. *IEEE Transactions on Information Theory*, 2022b.
- D Pollard. *Convergence of Stochastic Processes*. David Pollard, 1984.

- Nikita Puchkin and Nikita Zhivotovskiy. Exponential savings in agnostic active learning through abstention. *arXiv preprint arXiv:2102.00451*, 2021.
- Pengzhen Ren, Yun Xiao, Xiaojun Chang, Po-Yao Huang, Zihui Li, Brij B Gupta, Xiaojiang Chen, and Xin Wang. A survey of deep active learning. *ACM Computing Surveys (CSUR)*, 54(9):1–40, 2021.
- Daniel Russo and Benjamin Van Roy. Eluder dimension and the sample complexity of optimistic exploration. In *NIPS*, pages 2256–2264. Citeseer, 2013.
- Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach. In *International Conference on Learning Representations*, 2018.
- Burr Settles. Active learning literature survey. 2009.
- Shubhanshu Shekhar, Mohammad Ghavamzadeh, and Tara Javidi. Active learning for classification with abstention. *IEEE Journal on Selected Areas in Information Theory*, 2(2):705–719, 2021.
- Alexander B Tsybakov. Optimal aggregation of classifiers in statistical learning. *The Annals of Statistics*, 32(1):135–166, 2004.
- Michael Unser. Ridges, neural networks, and the radon transform. *arXiv preprint arXiv:2203.02543*, 2022.
- VN Vapnik and A Ya Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16(2):264, 1971.
- Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.
- Liwei Wang. Smoothness, disagreement coefficient, and the label complexity of agnostic active learning. *Journal of Machine Learning Research*, 12(7), 2011.
- Zhilei Wang, Pranjali Awasthi, Christoph Dann, Ayush Sekhari, and Claudio Gentile. Neural active learning with performance guarantees. *Advances in Neural Information Processing Systems*, 34, 2021.
- Andrew Chi-Chin Yao. Probabilistic computations: Toward a unified measure of complexity. In *18th Annual Symposium on Foundations of Computer Science (sfcs 1977)*, pages 222–227. IEEE Computer Society, 1977.
- Dmitry Yarotsky. Error bounds for approximations with deep relu networks. *Neural Networks*, 94: 103–114, 2017.
- Dmitry Yarotsky. Optimal approximation of continuous functions by very deep relu networks. In *Conference on learning theory*, pages 639–649. PMLR, 2018.
- Yinglun Zhu and Robert Nowak. Efficient active learning with abstention. *arXiv preprint arXiv:2204.00043*, 2022.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes] See [Section 5](#) for discussions on limitations and directions for future work.
 - (c) Did you discuss any potential negative societal impacts of your work? [N/A] Our paper is theoretical in nature, and there is no negative societal impact of our work in the foreseeable future.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes] Assumptions are clearly stated in the statement of each theorem.
 - (b) Did you include complete proofs of all theoretical results? [Yes] Complete proofs are provided in the Appendix.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [N/A]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Omitted details for Section 2.2

Proposition 1. *Suppose $\mathcal{D}_{\mathcal{X}\mathcal{Y}} \in \mathcal{P}(\alpha, \beta)$. One can construct a set of neural network classifier \mathcal{H}_{dnn} such that the following two properties hold simultaneously:*

$$\inf_{h \in \mathcal{H}_{\text{dnn}}} \text{err}(h) - \text{err}(h^*) = O(\varepsilon) \quad \text{and} \quad \text{VCdim}(\mathcal{H}_{\text{dnn}}) = \tilde{O}(\varepsilon^{-\frac{d}{\alpha(1+\beta)}}).$$

Proof. We take $\kappa = \varepsilon^{\frac{1}{1+\beta}}$ in [Theorem 3](#) to construct a set of neural network classifiers \mathcal{H}_{dnn} with $W = O(\varepsilon^{-\frac{d}{\alpha(1+\beta)}} \log \frac{1}{\varepsilon})$ total parameters arranged in $L = O(\log \frac{1}{\varepsilon})$ layers. According to [Theorem 4](#), we know

$$\text{VCdim}(\mathcal{H}_{\text{dnn}}) = O(\varepsilon^{-\frac{d}{\alpha(1+\beta)}} \cdot \log^2(\varepsilon^{-1})) = \tilde{O}(\varepsilon^{-\frac{d}{\alpha(1+\beta)}}).$$

We now show that there exists a classifier $\bar{h} \in \mathcal{H}_{\text{dnn}}$ with small excess error. Let $\bar{h} = h_{\bar{f}}$ be the classifier such that $\|\bar{f} - \eta\|_{\infty} \leq \kappa$. We can see that

$$\begin{aligned} \text{excess}(\bar{h}) &= \mathbb{E}[\mathbb{1}(\bar{h}(x) \neq y) - \mathbb{1}(h^*(x) \neq y)] \\ &= \mathbb{E}[|2\eta(x) - 1| \cdot \mathbb{1}(\bar{h}(x) \neq h^*(x))] \\ &\leq 2\kappa \cdot \mathbb{P}_{x \sim \mathcal{D}_{\mathcal{X}}}(x \in \mathcal{X} : |\eta(x) - 1/2| \leq \kappa) \\ &= O(\kappa^{1+\beta}) \\ &= O(\varepsilon), \end{aligned}$$

where the third line follows from the fact that \bar{h} and h^* disagrees only within region $\{x \in \mathcal{X} : |\eta(x) - 1/2| \leq \kappa\}$ and the incurred error is at most 2κ on each disagreed data point. The fourth line follows from the Tsybakov noise condition and the last line follows from the selection of κ . \square

Before proving [Theorem 5](#), we first recall the excess error guarantee for empirical risk minimization under Tsybakov noise condition.

Theorem 11 ([Boucheron et al. \(2005\)](#)). *Suppose $\mathcal{D}_{\mathcal{X}\mathcal{Y}}$ satisfies Tsybakov noise condition with parameter $\beta \geq 0$. Consider a dataset $D_n = \{(x_i, y_i)\}_{i=1}^n$ of n points i.i.d. sampled from $\mathcal{D}_{\mathcal{X}\mathcal{Y}}$. Let $\hat{h} \in \mathcal{H}$ be the empirical risk minimizer on D_n . For any constant $\rho > 0$, we have*

$$\begin{aligned} &\text{err}(\hat{h}) - \min_{h \in \mathcal{H}} \text{err}(h) \\ &\leq \rho \cdot (\min_{h \in \mathcal{H}} \text{err}(h) - \text{err}(h^*)) + O\left(\frac{(1+\rho)^2}{\rho} \cdot \left(\frac{\text{VCdim}(\mathcal{H}) \cdot \log n}{n}\right)^{\frac{1+\beta}{2+\beta}} + \frac{\log \delta^{-1}}{n}\right), \end{aligned}$$

with probability at least $1 - \delta$.

Theorem 5. *Suppose $\mathcal{D}_{\mathcal{X}\mathcal{Y}} \in \mathcal{P}(\alpha, \beta)$. Fix any $\varepsilon, \delta > 0$. Let \mathcal{H}_{dnn} be the set of neural network classifiers constructed in [Proposition 1](#). With $n = \tilde{O}(\varepsilon^{-\frac{d+2\alpha+\alpha\beta}{\alpha(1+\beta)}})$ i.i.d. sampled points, with probability at least $1 - \delta$, the empirical risk minimizer $\hat{h} \in \mathcal{H}_{\text{dnn}}$ achieves excess error $O(\varepsilon)$.*

Proof. [Proposition 1](#) certifies $\min_{h \in \mathcal{H}_{\text{dnn}}} \text{err}(h) - \text{err}(h^*) = O(\varepsilon)$ and $\text{VCdim}(\mathcal{H}_{\text{dnn}}) = O(\varepsilon^{-\frac{d}{\alpha(1+\beta)}} \cdot \log^2(\varepsilon^{-1}))$. Take $\rho = 1$ in [Theorem 11](#), leads to

$$\text{err}(\hat{h}) - \text{err}(h^*) \leq O\left(\varepsilon + \left(\varepsilon^{-\frac{d}{\alpha(1+\beta)}} \cdot \log^2(\varepsilon^{-1}) \cdot \frac{\log n}{n}\right)^{\frac{1+\beta}{2+\beta}} + \frac{\log \delta^{-1}}{n}\right),$$

Taking $n = O(\varepsilon^{-\frac{d+2\alpha+\alpha\beta}{\alpha(1+\beta)}} \cdot \log(\varepsilon^{-1}) + \varepsilon^{-1} \cdot \log(\delta^{-1})) = \tilde{O}(\varepsilon^{-\frac{d+2\alpha+\alpha\beta}{\alpha(1+\beta)}})$ thus ensures that $\text{err}(\hat{h}) - \text{err}(h^*) = O(\varepsilon)$. \square

B Generic version of Algorithm 1 and its guarantees

We present Algorithm 3 below, a generic version of Algorithm 1 that doesn't require the approximating classifiers to be neural networks. The guarantees of Algorithm 3 are provided in Theorem 12, which is proved in Appendix B.2 based on supporting lemmas provided in Appendix B.1.

Algorithm 3 RobustCAL with Approximation

Input: Accuracy level $\varepsilon \in (0, 1)$, confidence level $\delta \in (0, 1)$.

- 1: Let \mathcal{H} be a set of approximating classifiers such that $\inf_{h \in \mathcal{H}} \text{err}(h) - \text{err}(h^*) = O(\varepsilon)$.
 - 2: Define $T := \varepsilon^{-\frac{2+\beta}{1+\beta}} \cdot \text{VCdim}(\mathcal{H})$, $M := \lceil \log_2 T \rceil$, $\tau_m := 2^m$ for $m \geq 1$ and $\tau_0 := 0$.
 - 3: Define $\rho_m := O\left(\left(\frac{\text{VCdim}(\mathcal{H}) \cdot \log(\tau_{m-1}) \cdot \log(M/\delta)}{\tau_{m-1}}\right)^{\frac{1+\beta}{2+\beta}}\right)$ for $m \geq 2$ and $\rho_1 := 1$.
 - 4: Define $\widehat{R}_m(h) := \sum_{t=1}^{\tau_{m-1}} Q_t \mathbb{1}(h(x_t) \neq y_t)$ with the convention that $\sum_{t=1}^0 \dots = 0$.
 - 5: Initialize $\mathcal{H}_0 := \mathcal{H}$.
 - 6: **for** epoch $m = 1, 2, \dots, M$ **do**
 - 7: Update active set $\mathcal{H}_m := \left\{ h \in \mathcal{H}_{m-1} : \widehat{R}_m(h) \leq \inf_{h \in \mathcal{H}_{m-1}} \widehat{R}_m(h) + \tau_{m-1} \cdot \rho_m \right\}$
 - 8: **if** epoch $m = M$ **then**
 - 9: **Return** any classifier $\widehat{h} \in \mathcal{H}_M$.
 - 10: **for** time $t = \tau_{m-1} + 1, \dots, \tau_m$ **do**
 - 11: Observe $x_t \sim \mathcal{D}_{\mathcal{X}}$. Set $Q_t := \mathbb{1}(x_t \in \text{DIS}(\mathcal{H}_m))$.
 - 12: **if** $Q_t = 1$ **then**
 - 13: Query the label y_t of x_t .
-

We provide guarantees for Algorithm 3, and then specialize them to the settings with neural network approximation, i.e., in Theorem 6 and Theorem 10. As discussed before, our analysis is based on the analysis RobustCAL, but with novel extensions in removing the requirements that the Bayes classifier is in the class (or a Bernstein condition for every $h \in \mathcal{H}$).

Theorem 12. Fix $\varepsilon, \delta > 0$. With probability at least $1 - \delta$, Algorithm 3 returns a classifier $\widehat{h} \in \mathcal{H}$ with excess error $\widetilde{O}(\varepsilon)$ after querying

$$\widetilde{O}\left(\theta_{\mathcal{H}}(\varepsilon^{\frac{\beta}{1+\beta}}) \cdot \varepsilon^{-\frac{2}{1+\beta}} \cdot \text{VCdim}(\mathcal{H})\right)$$

labels.

B.1 Supporting lemmas

We first recall that Tsybakov noise condition leads to the so-called Bernstein condition (with respect to Bayes classifier h^*).

Lemma 1 (Tsybakov (2004)). Suppose $\mathcal{D}_{\mathcal{X}\mathcal{Y}}$ satisfies the Tsybakov noise condition with parameter $\beta \geq 0$, then there exists an universal constant $c' > 0$ such that we have

$$\mathbb{P}_{x \sim \mathcal{D}_{\mathcal{X}}}(h(x) \neq h^*(x)) \leq c'(\text{err}(h) - \text{err}(h^*))^{\frac{\beta}{1+\beta}}$$

for any $h : \mathcal{X} \rightarrow \mathcal{Y}$.

We next present a lemma in the passive learning setting, which will later be incorporated into the active learning setting. We first define some notations. Suppose $D_n = \{(x_i, y_i)\}_{i=1}^n$ are n i.i.d. data points drawn from $\mathcal{D}_{\mathcal{X}\mathcal{Y}}$. For any $h : \mathcal{X} \rightarrow \mathcal{Y}$, we denote $\bar{R}_n(h) := \sum_{i=1}^n \mathbb{1}(h(x_i) \neq y_i)$ as the empirical error of h over dataset D_n . We clearly have $\mathbb{E}[\bar{R}_n(h)] = n \cdot \text{err}(h)$ by i.i.d. assumption.

Lemma 2. Fix $\varepsilon, \bar{\delta} > 0$. Suppose $\mathcal{D}_{\mathcal{X}\mathcal{Y}}$ satisfies Tsybakov noise condition with parameter $\beta \geq 0$ and $\text{err}(\bar{h}) - \text{err}(h^*) = O(\varepsilon)$, where $\bar{h} = \arg \max_{h \in \mathcal{H}} \text{err}(h)$ and h^* is the Bayes classifier. Let $D_n = \{(x_i, y_i)\}_{i=1}^n$ be a set of n i.i.d. data points drawn from $\mathcal{D}_{\mathcal{X}\mathcal{Y}}$. If $\beta > 0$, suppose n satisfies

$$n \leq \varepsilon^{-\frac{2+\beta}{1+\beta}} \cdot \text{VCdim}(\mathcal{H})^{\frac{2+\beta}{\beta}} \cdot \log(\bar{\delta}^{-1}) \cdot (\log n)^{\frac{2+\beta}{\beta}}.$$

With probability at least $1 - \bar{\delta}$, we have the following inequalities hold:

$$n \cdot (\text{err}(h) - \text{err}(h^*)) \leq 2 \cdot (\bar{R}_n(h) - \bar{R}_n(\check{h})) + n \cdot \rho(n, \bar{\delta}), \quad \forall h \in \mathcal{H}, \quad (1)$$

$$\bar{R}_n(\check{h}) - \min_{h \in \mathcal{H}} \bar{R}_n(h) \leq n \cdot \rho(n, \bar{\delta}), \quad (2)$$

where $\rho(n, \bar{\delta}) = C \cdot \left(\left(\frac{\text{VCdim}(\mathcal{H}) \cdot \log n \cdot \log \bar{\delta}^{-1}}{n} \right)^{\frac{1+\beta}{2+\beta}} + \varepsilon \right)$ with a universal constant $C > 0$.⁸

Proof. Denote $\bar{\mathcal{H}} := \mathcal{H} \cup \{h^*\}$. We know that $\text{VCdim}(\bar{\mathcal{H}}) \leq \text{VCdim}(\mathcal{H}) + 1 = O(\text{VCdim}(\mathcal{H}))$. From [Lemma 1](#), we know Bernstein condition is satisfied with respect to $\bar{\mathcal{H}}$ and $h^* \in \bar{\mathcal{H}}$. Invoking [Lemma 3.1 in Hanneke \(2014\)](#), with probability at least $1 - \frac{\bar{\delta}}{2}$, $\forall h \in \bar{\mathcal{H}}$, we have

$$n \cdot (\text{err}(h) - \text{err}(h^*)) \leq \max\{2 \cdot (\bar{R}_n(h) - \bar{R}_n(h^*)), n \cdot \bar{\rho}(n, \bar{\delta})\}, \quad (3)$$

$$\bar{R}_n(h) - \min_{h \in \bar{\mathcal{H}}} \bar{R}_n(h) \leq \max\{2n \cdot (\text{err}(h) - \text{err}(h^*)), n \cdot \bar{\rho}(n, \bar{\delta})\}, \quad (4)$$

where $\bar{\rho}(n, \bar{\delta}) = O\left(\left(\frac{\text{VCdim}(\bar{\mathcal{H}}) \cdot \log n + \log \bar{\delta}^{-1}}{n}\right)^{\frac{1+\beta}{2+\beta}}\right) = O\left(\left(\frac{\text{VCdim}(\mathcal{H}) \cdot \log n \cdot \log \bar{\delta}^{-1}}{n}\right)^{\frac{1+\beta}{2+\beta}}\right)$.

[Eq. \(2\)](#) follows by taking $h = \check{h}$ in [Eq. \(4\)](#) and noticing that

$$\begin{aligned} \bar{R}_n(\check{h}) - \min_{h \in \mathcal{H}} \bar{R}_n(h) &\leq \bar{R}_n(\check{h}) - \min_{h \in \bar{\mathcal{H}}} \bar{R}_n(h) \\ &\leq \max\{2n \cdot O(\varepsilon), n \cdot \bar{\rho}(n, \bar{\delta})\}, \end{aligned}$$

where we use the assumption that $\text{err}(\check{h}) - \text{err}(h^*) = O(\varepsilon)$.

To derive [Eq. \(1\)](#), we first notice that applying [Eq. \(3\)](#) for any $h \in \mathcal{H}$, we have

$$n \cdot (\text{err}(h) - \text{err}(h^*)) \leq 2 \cdot (\bar{R}_n(h) - \bar{R}_n(\check{h}) + \bar{R}_n(\check{h}) - \bar{R}_n(h^*)) + n \cdot \bar{\rho}(n, \bar{\delta}).$$

We next only need to upper bound $\bar{R}_n(\check{h}) - \bar{R}_n(h^*)$, and show that it is order-wise smaller than $n \cdot \rho(n, \bar{\delta})$. We consider random variable $g_i := \mathbb{1}(\check{h}(x_i) \neq y_i) - \mathbb{1}(h^*(x_i) \neq y_i)$. We have

$$\begin{aligned} \mathbb{V}(g_i) &\leq \mathbb{E}[g_i^2] \\ &= \mathbb{E}[\mathbb{1}(\check{h}(x_i) \neq h^*(x_i))] \\ &= O\left(\varepsilon^{\frac{\beta}{1+\beta}}\right), \end{aligned}$$

where the last line follows from [Lemma 1](#) and the assumption that $\text{err}(\check{h}) - \text{err}(h^*) = O(\varepsilon)$. Denote $g = \frac{1}{n} \sum_{i=1}^n g_i = \frac{1}{n} (\bar{R}_n(\check{h}) - \bar{R}_n(h^*))$, and notice that $\mathbb{E}[g] = \text{err}(\check{h}) - \text{err}(h^*)$. Applying Bernstein inequality on $-g$, with probability at least $1 - \frac{\bar{\delta}}{2}$, we have

$$g - \mathbb{E}[g] \leq O\left(\left(\frac{\varepsilon^{\frac{\beta}{1+\beta}} \log \bar{\delta}^{-1}}{n}\right)^{\frac{1}{2}} + \frac{\log \bar{\delta}^{-1}}{n}\right),$$

which further leads to

$$\bar{R}_n(\check{h}) - \bar{R}_n(h^*) \leq n \cdot O\left(\varepsilon + \left(\frac{\varepsilon^{\frac{\beta}{1+\beta}} \log \bar{\delta}^{-1}}{n}\right)^{\frac{1}{2}} + \frac{\log \bar{\delta}^{-1}}{n}\right).$$

The RHS is order-wise smaller than ρ_n when $\beta = 0$. We consider the case when $\beta > 0$ next. Since $\log(\bar{\delta}^{-1})/n$ is clearly a lower-order term compared to ρ_n , we only need to show that $\left(\frac{\varepsilon^{\frac{\beta}{1+\beta}} \log \bar{\delta}^{-1}}{n}\right)^{\frac{1}{2}}$

⁸The logarithmic factors in this bound might be further optimized. We don't focus on optimizing logarithmic factors.

is order-wise smaller than ρ_n . We can easily check that

$$\left(\frac{\varepsilon^{\frac{\beta}{1+\beta}} \log \bar{\delta}^{-1}}{n} \right)^{\frac{1}{2}} \leq \left(\frac{\text{VCdim}(\mathcal{H}) \cdot \log n \cdot \log \bar{\delta}^{-1}}{n} \right)^{\frac{1+\beta}{2+\beta}}$$

whenever n satisfies the following condition

$$n \leq \varepsilon^{-\frac{2+\beta}{1+\beta}} \cdot \text{VCdim}(\mathcal{H})^{\frac{2+2\beta}{\beta}} \cdot \log(\bar{\delta}^{-1}) \cdot (\log n)^{\frac{2+2\beta}{\beta}}.$$

□

We denote $\check{h} = \arg \min_{h \in \mathcal{H}} \text{err}(h)$. By assumption of [Theorem 12](#), we have $\text{err}(\check{h}) - \text{err}(h^*) = O(\varepsilon)$. For any $h \in \mathcal{H}$, we also use the shorthand $\bar{R}_m(h) = \bar{R}_{\tau_{m-1}}(h) := \sum_{t=1}^{\tau_{m-1}} \mathbb{1}(h(x_t) \neq y_t)$. Note that \bar{R}_m is only used in analysis since some y_t are not observable.

Lemma 3. *With probability at least $1 - \frac{\delta}{2}$, the following holds true for all epochs $m \in [M]$:*

1. $\check{h} \in \mathcal{H}_m$.
2. $\text{err}(h) - \text{err}(h^*) \leq 3\rho_m, \forall h \in \mathcal{H}_m$.

Proof. For each $m = 2, 3, \dots, M$, we invoke [Lemma 2](#) with $n = \tau_{m-1}$ and $\bar{\delta} = \delta/2M$, which guarantees that

$$\tau_{m-1} \cdot (\text{err}(h) - \text{err}(h^*)) \leq 2 \cdot (\bar{R}_m(h) - \bar{R}_m(\check{h})) + \tau_{m-1} \cdot \rho_m, \quad \forall h \in \mathcal{H}, \quad (5)$$

$$\bar{R}_m(\check{h}) - \min_{h \in \mathcal{H}} \bar{R}_m(h) \leq \tau_{m-1} \cdot \rho_m. \quad (6)$$

Note that the choice T chosen in [Algorithm 3](#) clearly satisfies the requirement needed (for $n = \tau_{m-1}$) in [Lemma 2](#) when $\beta > 0$; and ensures that the second term in $\rho(\tau_{m-1}, \delta/2M)$ (i.e., ε , see [Lemma 2](#) for definition of $\rho(\tau_{m-1}, \delta/2M)$) is a lower-order term compared to the first term.

We use \mathcal{E} to denote the good event where [Eq. \(5\)](#) and [Eq. \(6\)](#) hold true across $m = 2, 3, \dots, M$. This good event happens with probability at least $1 - \frac{\delta}{2}$. We analyze under \mathcal{E} in the following.

We prove [Lemma 3](#) through induction. The statements clearly hold true for $m = 1$. Suppose the statements hold true up to epoch m , we next prove the correctness for epoch $m + 1$.

We know that $\check{h} \in \mathcal{H}_m$ by assumption. Based on the querying criteria of [Algorithm 3](#), we know that

$$\widehat{R}_{m+1}(\check{h}) - \widehat{R}_{m+1}(h) = \bar{R}_{m+1}(\check{h}) - \bar{R}_{m+1}(h), \quad \forall h \in \mathcal{H}_m \quad (7)$$

From [Eq. \(6\)](#), we also have

$$\begin{aligned} \bar{R}_{m+1}(\check{h}) - \min_{h \in \mathcal{H}_m} \bar{R}_{m+1}(h) &\leq \bar{R}_{m+1}(\check{h}) - \min_{h \in \mathcal{H}} \bar{R}_{m+1}(h) \\ &\leq \tau_m \cdot \rho_{m+1}. \end{aligned}$$

Combining the above two inequalities shows that

$$\widehat{R}_{m+1}(\check{h}) - \widehat{R}_{m+1}(h) \leq \tau_m \cdot \rho_{m+1},$$

implying that $\check{h} \in \mathcal{H}_{m+1}$ (due to the construction of \mathcal{H}_{m+1} in [Algorithm 3](#)).

Based on [Eq. \(7\)](#), the construction \mathcal{H}_{m+1} and the fact that $\check{h} \in \mathcal{H}_m$, we know that, for any $h \in \mathcal{H}_{m+1} \subseteq \mathcal{H}_m$,

$$\begin{aligned} \bar{R}_{m+1}(h) - \bar{R}_{m+1}(\check{h}) &= \widehat{R}_{m+1}(h) - \widehat{R}_{m+1}(\check{h}) \\ &\leq \widehat{R}_{m+1}(h) - \min_{h \in \mathcal{H}_m} \widehat{R}_{m+1}(h) \\ &\leq \tau_m \cdot \rho_{m+1}. \end{aligned}$$

Plugging the above inequality into [Eq. \(5\)](#) (at epoch $m + 1$) leads to $\text{err}(h) - \text{err}(h^*) \leq 3\rho_{m+1}$ for any $h \in \mathcal{H}_{m+1}$. We thus prove the desired statements at epoch $m + 1$. □

B.2 Proof of Theorem 12

Theorem 12. Fix $\varepsilon, \delta > 0$. With probability at least $1 - \delta$, [Algorithm 3](#) returns a classifier $\hat{h} \in \mathcal{H}$ with excess error $\tilde{O}(\varepsilon)$ after querying

$$\tilde{O}\left(\theta_{\mathcal{H}}(\varepsilon^{\frac{\beta}{1+\beta}}) \cdot \varepsilon^{-\frac{2}{1+\beta}} \cdot \text{VCdim}(\mathcal{H})\right)$$

labels.

Proof. Based on [Lemma 3](#), we know that, with probability at least $1 - \frac{\delta}{2}$, we have

$$\begin{aligned} \text{err}(\hat{h}) - \text{err}(h^*) &\leq 3\rho_M \\ &= O\left(\left(\frac{\text{VCdim}(\mathcal{H}) \cdot \log(\tau_{M-1}) \cdot \log(M/\delta)}{\tau_{M-1}}\right)^{\frac{1+\beta}{2+\beta}}\right) \\ &= \tilde{O}(\varepsilon), \end{aligned}$$

where we use the definition of T and τ_M .

We next analyze the label complexity of [Algorithm 3](#). Since [Algorithm 3](#) stops at the beginning of epoch M , we only need to calculate the label complexity in the first $M - 1$ epochs. We have

$$\begin{aligned} \sum_{t=1}^{\tau_{M-1}} Q_t &= \sum_{m=1}^{M-1} (\tau_m - \tau_{m-1}) \cdot \mathbb{1}(x_t \in \text{DIS}(\mathcal{H}_m)) \\ &\leq \sum_{m=1}^{M-1} (\tau_m - \tau_{m-1}) \cdot \mathbb{1}\left(x_t \in \text{DIS}(\mathcal{B}_{\mathcal{H}}(h^*, c'(3\rho_m)^{\frac{\beta}{1+\beta}}))\right), \end{aligned}$$

where on the last line we use the facts (1) $\text{err}(h) - \text{err}(h^*) \leq 3\rho_m, \forall h \in \mathcal{H}_m$ from [Lemma 3](#); and (2) $\mathbb{P}(x : h(x) \neq h^*(x)) \leq c'(\text{err}(h) - \text{err}(h^*))^{\frac{\beta}{1+\beta}}$ from [Lemma 2](#) (with the same constant c'). Suppose $\text{err}(\check{h}) - \text{err}(h^*) = c''\varepsilon$ (with another universal constant c'' by assumption). Applying [Lemma 2](#) on \check{h} leads to the fact that $h^* \in \mathcal{B}_{\mathcal{H}}(\check{h}, c''\varepsilon)^{\frac{\beta}{1+\beta}}$. Since $\mathbb{P}(x : h(x) \neq \check{h}(x)) \leq \mathbb{P}(x : h(x) \neq h^*(x)) + \mathbb{P}(x : h^*(x) \neq \check{h}(x))$, we further have

$$\sum_{t=1}^{\tau_{M-1}} Q_t \leq \sum_{m=1}^{M-1} (\tau_m - \tau_{m-1}) \cdot \mathbb{1}\left(x_t \in \text{DIS}(\mathcal{B}_{\mathcal{H}}(\check{h}, \bar{c} \cdot \rho_m^{\frac{\beta}{1+\beta}}))\right),$$

with a universal constant $\bar{c} > 0$. Noticing that the RHS is a sum of independent Bernoulli random variables, applying a Bernstein-type bound (e.g., [Lemma 5](#)), on a good event \mathcal{E}' that happens with probability at least $1 - \frac{\delta}{2}$, we have

$$\begin{aligned} \sum_{t=1}^{\tau_{M-1}} Q_t &\leq 2 \sum_{m=1}^{M-1} (\tau_m - \tau_{m-1}) \cdot \mathbb{P}\left(x \in \text{DIS}(\mathcal{B}_{\mathcal{H}}(\check{h}, \bar{c} \cdot \rho_m^{\frac{\beta}{1+\beta}}))\right) + 4 \log(4/\delta) \\ &\leq 2 \sum_{m=2}^{M-1} \tau_{m-1} \cdot \theta_{\mathcal{H}, \check{h}}\left(\bar{c} \cdot \rho_m^{\frac{\beta}{1+\beta}}\right) \cdot \bar{c} \cdot \rho_m^{\frac{\beta}{1+\beta}} + 4 \log(4/\delta) + 4 \\ &\leq 2M \cdot \theta_{\mathcal{H}, \check{h}}\left(\bar{c} \cdot \rho_M^{\frac{\beta}{1+\beta}}\right) \cdot \left(\bar{c} \cdot \tau_{M-1} \cdot \rho_M^{\frac{\beta}{1+\beta}}\right) + 4 \log(4/\delta) + 4, \end{aligned}$$

where the second using the definition of disagreement coefficient; and the last line follows from the fact that ρ_m is non-increasing and $\tau_{m-1} \cdot \rho_m$ is increasing. Basic algebra and basic properties of the disagreement coefficient (i.e., Theorem 7.1 and Corollary 7.2 in [Hanneke \(2014\)](#)) shows that

$$\sum_{t=1}^{\tau_{M-1}} Q_t \leq \tilde{O}\left(\theta_{\mathcal{H}}(\varepsilon^{\frac{\beta}{1+\beta}}) \cdot \varepsilon^{-\frac{2}{1+\beta}} \cdot \text{VCdim}(\mathcal{H})\right),$$

under event $\mathcal{E} \cap \mathcal{E}'$, which happen with probability at least $1 - \delta$. \square

C Omitted details for Section 2.3

We prove [Theorem 6](#) in [Appendix C.1](#) and discuss the disagreement coefficient in [Appendix C.2](#).

C.1 Proof of [Theorem 6](#)

Theorem 6. *Suppose $\mathcal{D}_{\mathcal{X}\mathcal{Y}} \in \mathcal{P}(\alpha, \beta)$. Fix any $\varepsilon, \delta > 0$. With probability at least $1 - \delta$, [Algorithm 1](#) returns a classifier $\hat{h} \in \mathcal{H}_{\text{dnn}}$ with excess error $\tilde{O}(\varepsilon)$ after querying $\tilde{O}(\theta_{\mathcal{H}_{\text{dnn}}}(\varepsilon^{\frac{\beta}{1+\beta}}) \cdot \varepsilon^{-\frac{d+2\alpha}{\alpha+\alpha\beta}})$ labels.*

Proof. Construct \mathcal{H}_{dnn} based on [Proposition 1](#) such that $\min_{h \in \mathcal{H}_{\text{dnn}}} \text{err}(h) - \text{err}(h^*) = O(\varepsilon)$ and $\text{VCdim}(\mathcal{H}_{\text{dnn}}) = \tilde{O}(\varepsilon^{-\frac{d}{\alpha(1+\beta)}})$. Taking such \mathcal{H}_{dnn} into [Theorem 12](#) leads to the desired result. \square

C.2 Discussion on disagreement coefficient in [Theorem 6](#)

We discuss cases when the (classifier-based) disagreement coefficient with respect to a set of neural networks is well-bounded. As mentioned before, even for simple classifiers such as linear functions, the disagreement coefficient has been analyzed under additional assumptions ([Friedman, 2009](#); [Hanneke, 2014](#)). In this section, we analyze the disagreement coefficient for a set of neural networks under additional assumptions on $\mathcal{D}_{\mathcal{X}\mathcal{Y}}$ and \mathcal{H}_{dnn} (assumptions on \mathcal{H}_{dnn} can be implemented via proper preprocessing steps). We leave a more comprehensive investigation of the disagreement coefficient for future work.

The first case is when $\mathcal{D}_{\mathcal{X}}$ is supported on countably many data points. The following result shows strict improvement over passive learning.

Definition 3 (Disagreement core). *For any hypothesis class \mathcal{H} and classifier h , the disagreement core of h with respect to \mathcal{H} under $\mathcal{D}_{\mathcal{X}\mathcal{Y}}$ is defined as*

$$\partial_{\mathcal{H}} h := \lim_{r \rightarrow 0} \text{DIS}(\mathcal{B}_{\mathcal{H}}(h, r)). \quad (8)$$

Proposition 2 (Lemma 7.12 and Theorem 7.14 in [Hanneke \(2014\)](#)). *For any hypothesis class \mathcal{H} and classifier h , we have $\theta_h(\varepsilon) = o(1/\varepsilon)$ if and only if $\mathcal{D}_{\mathcal{X}}(\partial_{\mathcal{H}} h) = 0$. In particular, this implies that $\theta_{\mathcal{H}}(\varepsilon) = o(1/\varepsilon)$ whenever $\mathcal{D}_{\mathcal{X}}$ is supported on countably many data points.*

We now discuss conditions under which we can upper bound the disagreement coefficient by $O(1)$, which ensures results in [Theorem 6](#) matching the minimax lower bound for active learning, up to logarithmic factors. We introduce the following *decomposable* condition.

Definition 4. *A marginal distribution $\mathcal{D}_{\mathcal{X}}$ is ε -decomposable if its (known) support $\text{supp}(\mathcal{D}_{\mathcal{X}})$ can be decomposed into connected subsets, i.e., $\text{supp}(\mathcal{D}_{\mathcal{X}}) = \cup_{i \in \mathcal{I}} \mathcal{X}_i$, such that*

$$\mathcal{D}_{\mathcal{X}}(\cup_{i \in \mathcal{I}'} \mathcal{X}_i) = O(\varepsilon),$$

where $\mathcal{I}' := \{i \in \mathcal{I} : \mathcal{D}_{\mathcal{X}}(\mathcal{X}_i) \leq \varepsilon\}$.

Remark 1. *Note that [Definition 4](#) permits a decomposition such that $|\bar{\mathcal{I}}| = \Omega(\frac{1}{\varepsilon})$ where $\bar{\mathcal{I}} = \mathcal{I} \setminus \mathcal{I}'$. [Definition 4](#) requires no knowledge of the index set \mathcal{I} or any \mathcal{X}_i ; it also places no restrictions on the conditional probability on each \mathcal{X}_i .*

We first give results for a general hypothesis class \mathcal{H} as follows, and then discuss how to bound the disagreement coefficient for a set of neural networks.

Proposition 3. *Suppose $\mathcal{D}_{\mathcal{X}}$ is decomposable (into $\cup_{i \in \mathcal{I}} \mathcal{X}_i$) and the hypothesis class \mathcal{H} consists of classifiers whose predication on each \mathcal{X}_i is the same, i.e., $|\{h(x) : x \in \mathcal{X}_i\}| = 1$ for any $h \in \mathcal{H}$ and $i \in \mathcal{I}$. We then have $\theta_{\mathcal{H}}(\varepsilon) = O(1)$ for ε sufficiently small.*

Proof. Fix any $h \in \mathcal{H}$. we know that for any $h' \in \mathcal{B}_{\mathcal{H}}(h, \varepsilon)$, we must have $\text{DIS}(\{h, h'\}) \subseteq \cup_{i \in \mathcal{I}'} \mathcal{X}_i$ since $\mathcal{D}_{\mathcal{X}}(x \in \mathcal{X} : h(x) \neq h'(x)) \leq \varepsilon$, and $|\{h(x) : x \in \mathcal{X}_i\}| = 1$ for any $h \in \mathcal{H}$ and any \mathcal{X}_i . This further implies that $\mathbb{P}(\text{DIS}(\mathcal{B}_{\mathcal{H}}(h, \varepsilon))) = O(\varepsilon)$, and thus $\theta_{\mathcal{H}}(\varepsilon) = O(1)$. \square

We next discuss conditions under which we can satisfy the prerequisites of [Proposition 3](#). Suppose $\mathcal{D}_{\mathcal{X}\mathcal{Y}} \in \mathcal{P}(\alpha, \beta)$. We assume that $\mathcal{D}_{\mathcal{X}}$ is $(\varepsilon^{\frac{\beta}{1+\beta}})$ -decomposable, and, for the desired accuracy level ε , we have

$$|\eta(x) - 1/2| \geq 2\varepsilon^{\frac{1}{1+\beta}}, \quad \forall x \in \text{supp}(\mathcal{D}_{\mathcal{X}}). \quad (9)$$

With the above conditions satisfied, we can filter out neural networks that are clearly not “close” to η . Specifically, with $\kappa = \varepsilon^{\frac{1}{1+\beta}}$ and \mathcal{F}_{dnn} be the set of neural networks constructed from [Proposition 1](#), we consider

$$\tilde{\mathcal{F}}_{\text{dnn}} := \{f \in \mathcal{F}_{\text{dnn}} : |f(x) - 1/2| \geq \varepsilon^{\frac{1}{1+\beta}}, \forall x \in \text{supp}(\mathcal{D}_{\mathcal{X}})\}, \quad (10)$$

which is guaranteed to contain $\bar{f} \in \mathcal{F}_{\text{dnn}}$ such that $\|\bar{f} - \eta\|_{\infty} \leq \varepsilon^{\frac{1}{1+\beta}}$. Now focus on the subset

$$\tilde{\mathcal{H}}_{\text{dnn}} := \{h_f : f \in \tilde{\mathcal{F}}_{\text{dnn}}\}. \quad (11)$$

We clearly have $h_{\bar{f}} \in \tilde{\mathcal{H}}_{\text{dnn}}$ (which ensures an $O(\varepsilon)$ -optimal classifier) and $\text{VCdim}(\tilde{\mathcal{H}}_{\text{dnn}}) \leq \text{VCdim}(\mathcal{H}_{\text{dnn}})$ (since $\tilde{\mathcal{H}}_{\text{dnn}} \subseteq \mathcal{H}_{\text{dnn}}$). We upper bound the disagreement coefficient $\theta_{\tilde{\mathcal{H}}_{\text{dnn}}}(\varepsilon^{\frac{\beta}{1+\beta}})$ next.

Proposition 4. *Suppose $\mathcal{D}_{\mathcal{X}\mathcal{Y}} \in \mathcal{P}(\alpha, \beta)$ such that $\mathcal{D}_{\mathcal{X}}$ is $(\varepsilon^{\frac{\beta}{1+\beta}})$ -decomposable and [Eq. \(9\)](#) is satisfied (with the desired accuracy level ε). We then have $\theta_{\tilde{\mathcal{H}}_{\text{dnn}}}(\varepsilon^{\frac{\beta}{1+\beta}}) = O(1)$.*

Proof. The proof is similar to the proof of [Proposition 3](#). Fix any $h = h_f \in \tilde{\mathcal{H}}_{\text{dnn}}$. We first argue that, for any $i \in \mathcal{I}$, under [Eq. \(9\)](#), $|\{h_f(x) : x \in \mathcal{X}_i\}| = 1$, i.e., for $x \in \mathcal{X}_i$, $h_f(x)$ equals either 1 or 0, but not both: This can be seen from the fact that any $f \in \tilde{\mathcal{F}}_{\text{dnn}}$ is continuous and satisfies $|f(x) - 1/2| \geq \varepsilon^{\frac{1}{1+\beta}}$ for any $x \in \mathcal{X}_i$.

Fix any $h \in \tilde{\mathcal{H}}_{\text{dnn}}$. We know that for any $h' \in \mathcal{H}_{\tilde{\mathcal{H}}_{\text{dnn}}}(h, \varepsilon^{\frac{\beta}{1+\beta}})$, we must have $\text{DIS}(\{h, h'\}) \subseteq \cup_{i \in \mathcal{I}'} \mathcal{X}_i$ due to similar reasons argued in the proof of [Proposition 3](#). This further implies that $\mathbb{P}(\text{DIS}(\mathcal{B}_{\tilde{\mathcal{H}}_{\text{dnn}}}(h, \varepsilon^{\frac{\beta}{1+\beta}})) = O(\varepsilon^{\frac{\beta}{1+\beta}})$, and thus $\theta_{\tilde{\mathcal{H}}_{\text{dnn}}}(\varepsilon^{\frac{\beta}{1+\beta}}) = O(1)$. \square

We next argue that [Eq. \(9\)](#) is only needed in an approximate sense. We define the approximate decomposable condition in the following.

Definition 5. *A marginal distribution $\mathcal{D}_{\mathcal{X}}$ is (ε, δ) -decomposable if there exists a known subset $\bar{\mathcal{X}} \subseteq \text{supp}(\mathcal{D}_{\mathcal{X}})$ such that*

$$\mathcal{D}_{\mathcal{X}}(\bar{\mathcal{X}}) \geq 1 - \delta, \quad (12)$$

and it can be decomposed into connected subsets, i.e., $\bar{\mathcal{X}} = \cup_{i \in \mathcal{I}'} \mathcal{X}_i$, such that

$$\mathcal{D}_{\mathcal{X}}(\cup_{i \in \mathcal{I}'} \mathcal{X}_i) = O(\varepsilon),$$

where $\mathcal{I}' := \{i \in \mathcal{I} : \mathcal{D}_{\mathcal{X}}(\mathcal{X}_i) \leq \varepsilon\}$.

Suppose $\mathcal{D}_{\mathcal{X}\mathcal{Y}} \in \mathcal{P}(\alpha, \beta)$. We assume that $\mathcal{D}_{\mathcal{X}}$ is $(\varepsilon^{\frac{\beta}{1+\beta}}, \varepsilon^{\frac{\beta}{1+\beta}})$ -decomposable (wrt $\bar{\mathcal{X}} \subseteq \mathcal{D}_{\mathcal{X}}$), and, for the desired accuracy level ε , we have

$$|\eta(x) - 1/2| \geq 2\varepsilon^{\frac{1}{1+\beta}}, \quad \forall x \in \bar{\mathcal{X}}. \quad (13)$$

With the above conditions satisfied, we can filter out neural networks that are clearly not “close” to η . Specifically, with $\kappa = \varepsilon^{\frac{1}{1+\beta}}$ and \mathcal{F}_{dnn} be the set of neural networks constructed from [Proposition 1](#), we consider

$$\bar{\mathcal{F}}_{\text{dnn}} := \{f \in \mathcal{F}_{\text{dnn}} : |f(x) - 1/2| \geq \varepsilon^{\frac{1}{1+\beta}}, \forall x \in \bar{\mathcal{X}}\}, \quad (14)$$

which is guaranteed to contain $\bar{f} \in \mathcal{F}_{\text{dnn}}$ such that $\|\bar{f} - \eta\|_{\infty} \leq \varepsilon^{\frac{1}{1+\beta}}$. Now focus on the subset

$$\bar{\mathcal{H}}_{\text{dnn}} := \{h_f : f \in \bar{\mathcal{F}}_{\text{dnn}}\}. \quad (15)$$

We clearly have $h_{\bar{f}} \in \bar{\mathcal{H}}_{\text{dnn}}$ (which ensures an $O(\varepsilon)$ -optimal classifier) and $\text{VCdim}(\bar{\mathcal{H}}_{\text{dnn}}) \leq \text{VCdim}(\mathcal{H}_{\text{dnn}})$ (since $\bar{\mathcal{H}}_{\text{dnn}} \subseteq \mathcal{H}_{\text{dnn}}$). We upper bound the disagreement coefficient $\theta_{\bar{\mathcal{H}}_{\text{dnn}}}(\varepsilon^{\frac{\beta}{1+\beta}})$ next.

Proposition 5. *Suppose $\mathcal{D}_{\mathcal{X}\mathcal{Y}} \in \mathcal{P}(\alpha, \beta)$ such that $\mathcal{D}_{\mathcal{X}}$ is $(\varepsilon^{\frac{1}{1+\beta}}, \varepsilon)$ -decomposable (wrt known $\bar{\mathcal{X}} \subseteq \text{supp}(\mathcal{D}_{\mathcal{X}})$) and [Eq. \(13\)](#) is satisfied (with the desired accuracy level ε). We then have $\theta_{\bar{\mathcal{H}}_{\text{dnn}}}(\varepsilon^{\frac{\beta}{1+\beta}}) = O(1)$.*

Proof. The proof is the same as the proof of [Proposition 5](#) except for any $h' \in \mathcal{H}_{\overline{\mathcal{H}}_{\text{dnn}}}(h, \varepsilon^{\frac{\beta}{1+\beta}})$, we must have $\text{DIS}(\{h, h'\}) \subseteq (\cup_{i \in \mathcal{I}'} \mathcal{X}_i) \cup (\text{supp}(\mathcal{D}_{\mathcal{X}}) \setminus \overline{\mathcal{X}})$. Based on the assumption that $\mathcal{D}_{\mathcal{X}}$ is $(\varepsilon^{\frac{1}{1+\beta}}, \varepsilon)$ -decomposable, this also leads to $\theta_{\overline{\mathcal{H}}_{\text{dnn}}}(\varepsilon^{\frac{\beta}{1+\beta}}) = O(1)$ \square

D Generic version of [Algorithm 2](#) and its guarantees

This section is organized as follows. We first introduce some complexity measures in [Appendix D.1](#). We then provide the generic algorithm ([Algorithm 4](#)) and state its theoretical guarantees ([Theorem 14](#)) in [Appendix D.2](#).

D.1 Complexity measures

We first introduce *pseudo dimension* ([Pollard, 1984](#); [Haussler, 1989, 1995](#)), a complexity measure used to analyze real-valued functions.

Definition 6 (Pseudo dimension). *Consider a set of real-valued function $\mathcal{F} : \mathcal{X} \rightarrow \mathbb{R}$. The pseudo dimension $\text{Pdim}(\mathcal{F})$ of \mathcal{F} is defined as the VC dimension of the set of threshold functions $\{(x, \zeta) \mapsto \mathbb{1}(f(x) > \zeta) : f \in \mathcal{F}\}$.*

As discussed in [Bartlett et al. \(2019\)](#), similar results as in [Theorem 4](#) holds true for $\text{Pdim}(\mathcal{F})$ as well.

Theorem 13 ([Bartlett et al. \(2019\)](#)). *Let \mathcal{F}_{dnn} be a set of neural network regression functions of the same architecture and with W parameters arranged in L layers. We then have*

$$\Omega(WL \log(W/L)) \leq \text{Pdim}(\mathcal{F}_{\text{dnn}}) \leq O(WL \log(W)).$$

We now introduce *value function disagreement coefficient*, which is proposed by [Foster et al. \(2020\)](#) in contextual bandits and then adapted to active learning by [Zhu and Nowak \(2022\)](#) with additional supreme over the marginal distribution $\mathcal{D}_{\mathcal{X}}$ to deal with distributional shifts caused by selective sampling.

Definition 7 (Value function disagreement coefficient). *For any $f^* \in \mathcal{F}$ and $\gamma_0, \varepsilon_0 > 0$, the value function disagreement coefficient $\theta_{f^*}^{\text{val}}(\mathcal{F}, \gamma_0, \varepsilon_0)$ is defined as*

$$\sup_{\mathcal{D}_{\mathcal{X}}} \sup_{\gamma > \gamma_0, \varepsilon > \varepsilon_0} \left\{ \frac{\gamma^2}{\varepsilon^2} \cdot \mathbb{P}_{\mathcal{D}_{\mathcal{X}}}(\exists f \in \mathcal{F} : |f(x) - f^*(x)| > \gamma, \|f - f^*\|_{\mathcal{D}_{\mathcal{X}}} \leq \varepsilon) \right\} \vee 1,$$

where $\|f\|_{\mathcal{D}_{\mathcal{X}}}^2 := \mathbb{E}_{x \sim \mathcal{D}_{\mathcal{X}}}[f^2(x)]$. We also define $\theta_{\mathcal{F}}^{\text{val}}(\gamma_0) := \sup_{f^* \in \mathcal{F}, \varepsilon_0 > 0} \theta_{f^*}^{\text{val}}(\mathcal{F}, \gamma_0, \varepsilon_0)$.

D.2 The generic algorithm and its guarantees

We present [Algorithm 4](#), a generic version of [Algorithm 2](#) that doesn't require the approximating classifiers to be neural networks.

Algorithm 4 NeuralCAL++ (Generic Version)

Input: Accuracy level $\varepsilon \in (0, 1)$, confidence level $\delta \in (0, 1)$, abstention parameter $\gamma \in (0, 1/2)$.

- 1: Let $\mathcal{F} : \mathcal{X} \rightarrow [0, 1]$ be a set of regression functions such that there exists a regression function $\bar{f} \in \mathcal{F}$ with $\|\bar{f} - \eta\|_\infty \leq \kappa \leq \gamma/4$.
- 2: Define $T := \frac{\theta_{\mathcal{F}}^{\text{val}}(\gamma/4) \cdot \text{Pdim}(\mathcal{F})}{\varepsilon \gamma}$, $M := \lceil \log_2 T \rceil$, and $C_\delta := O(\text{Pdim}(\mathcal{F}) \cdot \log(T/\delta))$.
- 3: Define $\tau_m := 2^m$ for $m \geq 1$, $\tau_0 := 0$, and $\beta_m := 3(M - m + 1)C_\delta$.
- 4: Define $\widehat{R}_m(f) := \sum_{t=1}^{\tau_m-1} Q_t(\widehat{f}(x_t) - y_t)^2$ with the convention that $\sum_{t=1}^0 \dots = 0$.
- 5: **for** epoch $m = 1, 2, \dots, M$ **do**
- 6: Get $\widehat{f}_m := \arg \min_{f \in \mathcal{F}} \sum_{t=1}^{\tau_m-1} Q_t(f(x_t) - y_t)^2$.
- 7: (Implicitly) Construct active set $\mathcal{F}_m := \{f \in \mathcal{F} : \widehat{R}_m(f) \leq \widehat{R}_m(\widehat{f}_m) + \beta_m\}$.
- 8: Construct classifier $\widehat{h}_m : \mathcal{X} \rightarrow \{+1, -1, \perp\}$ as

$$\widehat{h}_m(x) := \begin{cases} \perp, & \text{if } [\text{lcb}(x; \mathcal{F}_m) - \frac{\gamma}{4}, \text{ucb}(x; \mathcal{F}_m) + \frac{\gamma}{4}] \subseteq [\frac{1}{2} - \gamma, \frac{1}{2} + \gamma]; \\ \text{sign}(2\widehat{f}_m(x) - 1), & \text{o.w.} \end{cases}$$
 and query function $g_m(x) := \mathbb{1}(\frac{1}{2} \in (\text{lcb}(x; \mathcal{F}_m) - \frac{\gamma}{4}, \text{ucb}(x; \mathcal{F}_m) + \frac{\gamma}{4})) \cdot \mathbb{1}(\widehat{h}_m(x) \neq \perp)$.
- 9: **if** epoch $m = M$ **then**
- 10: **Return** classifier \widehat{h}_M .
- 11: **for** time $t = \tau_{m-1} + 1, \dots, \tau_m$ **do**
- 12: Observe $x_t \sim \mathcal{D}_{\mathcal{X}}$. Set $Q_t := g_m(x_t)$.
- 13: **if** $Q_t = 1$ **then**
- 14: Query the label y_t of x_t .

We next state the theoretical guarantees for [Algorithm 4](#).

Theorem 14. Suppose $\theta_{\mathcal{F}}^{\text{val}}(\gamma/4) \leq \bar{\theta}$ and the approximation level $\kappa \in (0, \gamma/4]$ satisfies

$$\left(\frac{432\bar{\theta} \cdot M^2}{\gamma^2} \right) \cdot \kappa^2 \leq \frac{1}{10}. \quad (16)$$

With probability at least $1 - \delta$, [Algorithm 4](#) returns a classifier $\widehat{h} : \mathcal{X} \rightarrow \{+1, -1, \perp\}$ with Chow's excess error

$$\text{excess}_\gamma(\widehat{h}) = O\left(\varepsilon \cdot \log\left(\frac{\bar{\theta} \cdot \text{Pdim}(\mathcal{F})}{\varepsilon \gamma \delta}\right)\right),$$

after querying at most

$$O\left(\frac{M^2 \cdot \text{Pdim}(\mathcal{F}) \cdot \log(T/\delta) \cdot \bar{\theta}}{\gamma^2}\right)$$

labels.

[Theorem 14](#) is proved in [Appendix D.3](#), based on supporting lemmas and theorems established in [Appendix D.2.1](#) and [Appendix D.2.2](#). The general result ([Theorem 14](#)) will be used to prove results in specific settings (e.g., [Theorem 7](#) and [Theorem 18](#)).

D.2.1 Concentration results

The Freedman's inequality is commonly used in the field of active learning and contextual bandits, e.g., ([Freedman, 1975](#); [Agarwal et al., 2014](#); [Krishnamurthy et al., 2019](#); [Foster et al., 2020](#)). We thus state the result without proof.

Lemma 4 (Freedman's inequality). *Let $(X_t)_{t \leq T}$ be a real-valued martingale difference sequence adapted to a filtration \mathfrak{F}_t , and let $\mathbb{E}_t[\cdot] := \mathbb{E}[\cdot \mid \mathfrak{F}_{t-1}]$. If $|X_t| \leq B$ almost surely, then for any $\eta \in (0, 1/B)$ it holds with probability at least $1 - \delta$,*

$$\sum_{t=1}^T X_t \leq \eta \sum_{t=1}^T \mathbb{E}_t[X_t^2] + \frac{\log \delta^{-1}}{\eta}.$$

Lemma 5. Let $(X_t)_{t \leq T}$ to be real-valued sequence of random variables adapted to a filtration \mathfrak{F}_t . If $|X_t| \leq B$ almost surely, then with probability at least $1 - \delta$,

$$\sum_{t=1}^T X_t \leq \frac{3}{2} \sum_{t=1}^T \mathbb{E}_t[X_t] + 4B \log(2\delta^{-1}),$$

and

$$\sum_{t=1}^T \mathbb{E}_t[X_t] \leq 2 \sum_{t=1}^T X_t + 8B \log(2\delta^{-1}).$$

Proof. This is a direct consequence of [Lemma 4](#). \square

We now define/recall some notations. Denote $n_m := \tau_m - \tau_{m-1}$. Fix any epoch $m \in [M]$ and any time step t within epoch m . We have $f^* = \eta$. For any $f \in \mathcal{F}$, we denote $M_t(f) := Q_t((f(x_t) - y_t)^2 - (f^*(x_t) - y_t)^2)$, and $\hat{R}_m(f) := \sum_{t=\tau_{m-1}}^{\tau_m-1} Q_t(f(x_t) - y_t)^2$. Recall that we have $Q_t = g_m(x_t)$. We define filtration $\mathfrak{F}_t := \sigma((x_1, y_1), \dots, (x_t, y_t))$,⁹ and denote $\mathbb{E}_t[\cdot] := \mathbb{E}[\cdot | \mathfrak{F}_{t-1}]$. We next present concentration results with respect to a general set of regression function \mathcal{F} with finite pseudo dimension.

Lemma 6 ([Krishnamurthy et al. \(2019\)](#)). Consider an infinite set of regression function \mathcal{F} . Fix any $\delta \in (0, 1)$. For any $\tau, \tau' \in [T]$ such that $\tau < \tau'$, with probability at least $1 - \frac{\delta}{2}$, we have

$$\sum_{t=\tau}^{\tau'} M_t(f) \leq \sum_{t=\tau}^{\tau'} \frac{3}{2} \mathbb{E}_t[M_t(f)] + C_\delta,$$

and

$$\sum_{t=\tau}^{\tau'} \mathbb{E}_t[M_t(f)] \leq 2 \sum_{t=\tau}^{\tau'} M_t(f) + C_\delta,$$

where $C_\delta = C \cdot \left(\text{Pdim}(\mathcal{F}) \cdot \log T + \log \left(\frac{\text{Pdim}(\mathcal{F}) \cdot T}{\delta} \right) \right)$ with a universal constant $C > 0$.

D.2.2 Supporting lemmas for [Theorem 14](#)

Fix any classifier $\hat{h} : \mathcal{X} \rightarrow \{+1, -1, \perp\}$. For any $x \in \mathcal{X}$, we use the notion

$$\begin{aligned} \text{excess}_\gamma(\hat{h}; x) &:= \\ &\mathbb{P}_{y|x}(y \neq \text{sign}(\hat{h}(x))) \cdot \mathbb{1}(\hat{h}(x) \neq \perp) + (1/2 - \gamma) \cdot \mathbb{1}(\hat{h}(x) = \perp) - \mathbb{P}_{y|x}(y \neq \text{sign}(h^*(x))) \\ &= \mathbb{1}(\hat{h}(x) \neq \perp) \cdot (\mathbb{P}_{y|x}(y \neq \text{sign}(\hat{h}(x))) - \mathbb{P}_{y|x}(y \neq \text{sign}(h^*(x)))) \\ &\quad + \mathbb{1}(\hat{h}(x) = \perp) \cdot ((1/2 - \gamma) - \mathbb{P}_{y|x}(y \neq \text{sign}(h^*(x)))) \end{aligned} \tag{17}$$

to represent the excess error of \hat{h} at point $x \in \mathcal{X}$. Excess error of classifier \hat{h} can be then written as $\text{excess}_\gamma(\hat{h}) := \text{err}_\gamma(\hat{h}) - \text{err}(h^*) = \mathbb{E}_{x \sim \mathcal{D}_x}[\text{excess}_\gamma(\hat{h}; x)]$.

We let \mathcal{E} denote the good event considered in [Lemma 6](#), we analyze under this event through out the rest of this section. Most lemmas presented in this section are inspired by results provided [Zhu and Nowak \(2022\)](#). Our main innovation is an inductive analysis of lemmas that eventually relaxes the requirements for approximation error for [Theorem 14](#).

General lemmas. We introduce some general lemmas for [Theorem 14](#).

Lemma 7. For any $m \in [M]$, we have $g_m(x) = 1 \implies w(x; \mathcal{F}_m) > \frac{\gamma}{2}$.

⁹ y_t is not observed (and thus not included in the filtration) when $Q_t = 0$. Note that Q_t is measurable with respect to $\sigma((\mathfrak{F}_{t-1}, x_t))$.

Proof. We only need to show that $\text{ucb}(x; \mathcal{F}_m) - \text{lcb}(x; \mathcal{F}_m) \leq \frac{\gamma}{2} \implies g_m(x) = 0$. Suppose otherwise $g_m(x) = 1$, which implies that both

$$\begin{aligned} \frac{1}{2} &\in \left(\text{lcb}(x; \mathcal{F}_m) - \frac{\gamma}{4}, \text{ucb}(x; \mathcal{F}_m) + \frac{\gamma}{4} \right) \quad \text{and} \\ &\left[\text{lcb}(x; \mathcal{F}_m) - \frac{\gamma}{4}, \text{ucb}(x; \mathcal{F}_m) + \frac{\gamma}{4} \right] \not\subseteq \left[\frac{1}{2} - \gamma, \frac{1}{2} + \gamma \right]. \end{aligned} \quad (18)$$

If $\frac{1}{2} \in (\text{lcb}(x; \mathcal{F}_m) - \frac{\gamma}{4}, \text{ucb}(x; \mathcal{F}_m) + \frac{\gamma}{4})$ and $\text{ucb}(x; \mathcal{F}_m) - \text{lcb}(x; \mathcal{F}_m) \leq \frac{\gamma}{2}$, we must have $\text{lcb}(x; \mathcal{F}_m) \geq \frac{1}{2} - \frac{3}{4}\gamma$ and $\text{ucb}(x; \mathcal{F}_m) \leq \frac{1}{2} + \frac{3}{4}\gamma$, which contradicts with Eq. (18). \square

Lemma 8. Fix any $m \in [M]$. Suppose $\bar{f} \in \mathcal{F}_m$, we have $\text{excess}_\gamma(\hat{h}_m; x) \leq 0$ if $g_m(x) = 0$.

Proof. Recall that

$$\begin{aligned} \text{excess}_\gamma(\hat{h}; x) &= \mathbb{1}(\hat{h}(x) \neq \perp) \cdot (\mathbb{P}_{y|x}(y \neq \text{sign}(\hat{h}(x))) - \mathbb{P}_{y|x}(y \neq \text{sign}(h^*(x)))) \\ &\quad + \mathbb{1}(\hat{h}(x) = \perp) \cdot ((1/2 - \gamma) - \mathbb{P}_{y|x}(y \neq \text{sign}(h^*(x)))). \end{aligned}$$

We now analyze the event $\{g_m(x) = 0\}$ in two cases.

Case 1: $\hat{h}_m(x) = \perp$.

Since $\bar{f}(x) \in [\text{lcb}(x; \mathcal{F}_m), \text{ucb}(x; \mathcal{F}_m)]$ and $\kappa \leq \frac{\gamma}{4}$ by assumption, we know that $\eta(x) = f^*(x) \in [\frac{1}{2} - \gamma, \frac{1}{2} + \gamma]$ and thus $\mathbb{P}_y(y \neq \text{sign}(h^*(x))) \geq \frac{1}{2} - \gamma$. As a result, we have $\text{excess}_\gamma(\hat{h}_m; x) \leq 0$.

Case 2: $\hat{h}_m(x) \neq \perp$ but $\frac{1}{2} \notin (\text{lcb}(x; \mathcal{F}_m) - \frac{\gamma}{4}, \text{ucb}(x; \mathcal{F}_m) + \frac{\gamma}{4})$.

Since $\bar{f}(x) \in [\text{lcb}(x; \mathcal{F}_m), \text{ucb}(x; \mathcal{F}_m)]$ and $\kappa \leq \frac{\gamma}{4}$ by assumption, we clearly have $\text{sign}(\hat{h}_m(x)) = \text{sign}(h^*(x))$ when $\frac{1}{2} \notin (\text{lcb}(x; \mathcal{F}_m) - \frac{\gamma}{4}, \text{ucb}(x; \mathcal{F}_m) + \frac{\gamma}{4})$. We thus have $\text{excess}_\gamma(\hat{h}_m; x) \leq 0$. \square

Inductive lemmas. We prove a set of statements for Theorem 14 in an inductive way. Fix any epoch $m \in [M]$, we consider

$$\left\{ \begin{array}{l} \hat{R}_m(\bar{f}) - \hat{R}_m(f^*) \leq \mathbb{E}_t \left[Q_t(\bar{f}(x_t) - f^*(x_t))^2 \right] + C_\delta \leq \frac{3}{2} C_\delta \\ \bar{f} \in \mathcal{F}_m \\ \sum_{t=1}^{\tau_{m-1}} \mathbb{E}_t[M_t(f)] \leq 4\beta_m, \forall f \in \mathcal{F}_m \\ \sum_{t=1}^{\tau_{m-1}} \mathbb{E}[Q_t(x_t)(f(x_t) - \bar{f}(x_t))^2] \leq 9\beta_m, \forall f \in \mathcal{F}_m \\ \mathcal{F}_m \subseteq \mathcal{F}_{m-1} \end{array} \right., \quad (19)$$

$$\mathbb{E}_{x \sim \mathcal{D}_X}[\mathbb{1}(g_m(x) = 1)] \leq \frac{144\beta_m}{\tau_{m-1}\gamma^2} \cdot \theta_{\bar{f}}^{\text{val}}(\mathcal{F}, \gamma/4, \sqrt{\beta_m/\tau_{m-1}}) \leq \frac{144\beta_m}{\tau_{m-1}\gamma^2} \cdot \bar{\theta}, \quad (20)$$

and

$$\mathbb{E}_{x \sim \mathcal{D}_X}[\mathbb{1}(g_m(x) = 1) \cdot w(x; \mathcal{F}_m)] \leq \frac{72\beta_m}{\tau_{m-1}\gamma} \cdot \theta_{\bar{f}}^{\text{val}}(\mathcal{F}, \gamma/4, \sqrt{\beta_m/\tau_{m-1}}) \leq \frac{72\beta_m}{\tau_{m-1}\gamma} \cdot \bar{\theta}. \quad (21)$$

Lemma 9. Fix any $\bar{m} = [M]$. When $\bar{m} = 1, 2$ or when Eq. (20) holds true for epochs $m = 2, 3, \dots, \bar{m} - 1$, then Eq. (19) holds true for epoch $m = \bar{m}$.

Proof. The statements in Eq. (19) clearly hold true for $m = \bar{m} = 1$ since, by definition, $\mathcal{F}_0 = \mathcal{F}$ and $\sum_{t=1}^0 \dots = 0$. We thus only need to consider the case when $\bar{m} \geq 2$. We next prove each of the five statements in Eq. (19) for epoch $m = \bar{m}$.

1. In the case when $\bar{m} = 2$, from [Lemma 6](#), we know that

$$\begin{aligned}\widehat{R}_{\bar{m}}(\bar{f}) - \widehat{R}_{\bar{m}}(f^*) &\leq \sum_{t=1}^{\tau_{\bar{m}-1}} \frac{3}{2} \cdot \mathbb{E}_t \left[Q_t(\bar{f}(x_t) - f^*(x_t))^2 \right] + C_\delta \\ &\leq 3 + C_\delta \leq \frac{3}{2} C_\delta,\end{aligned}$$

where the second line follows from the fact that $\tau_1 = 2$ (without loss of generality, we assume $C_\delta \geq 6$ here).

We now focus on the case when $\bar{m} \geq 3$. We have

$$\begin{aligned}\widehat{R}_{\bar{m}}(\bar{f}) - \widehat{R}_{\bar{m}}(f^*) &\leq \sum_{t=1}^{\tau_{\bar{m}-1}} \frac{3}{2} \cdot \mathbb{E}_t \left[Q_t(\bar{f}(x_t) - f^*(x_t))^2 \right] + C_\delta \\ &\leq \frac{3}{2} \sum_{\bar{m}=1}^{\bar{m}-1} n_{\bar{m}} \mathbb{E}_{x \sim \mathcal{D}_X} [\mathbb{1}(g_{\bar{m}}(x) = 1)] \cdot \kappa^2 + C_\delta \\ &\leq \frac{3}{2} \left(2 + \sum_{\bar{m}=2}^{\bar{m}-1} n_{\bar{m}} \frac{144\beta_{\bar{m}} \cdot \bar{\theta}}{\tau_{\bar{m}-1}\gamma^2} \right) \cdot \kappa^2 + C_\delta \\ &\leq \left(3 + \frac{144\bar{\theta}}{\gamma^2} \cdot \left(\sum_{\bar{m}=2}^{\bar{m}-1} \beta_{\bar{m}} \right) \right) \cdot \kappa^2 + C_\delta \\ &\leq \left(3 + \frac{432\bar{\theta} \cdot M^2}{\gamma^2} \cdot C_\delta \right) \cdot \kappa^2 + C_\delta \\ &\leq \frac{3}{2} C_\delta,\end{aligned}$$

where the first line follows from [Lemma 6](#); the second line follows from the fact that $\|\bar{f} - f^*\|_\infty \leq \kappa$; the third line follows from [Eq. \(20\)](#); the fourth line follows from $n_{\bar{m}} = \tau_{\bar{m}-1}$; the fifth line follows from the definition of $\beta_{\bar{m}}$; and the last line follows from the choice of κ in [Eq. \(16\)](#)

2. Since $\mathbb{E}_t[M_t(f)] = \mathbb{E}_t[Q_t(f(x_t) - f^*(x_t))^2]$, by [Lemma 6](#), we have $\widehat{R}_{\bar{m}}(f^*) \leq \widehat{R}_{\bar{m}}(f) + C_\delta/2$ for any $f \in \mathcal{F}$. Combining this with statement 1 leads to

$$\begin{aligned}\widehat{R}_{\bar{m}}(\bar{f}) &\leq \widehat{R}_{\bar{m}}(f) + 2C_\delta \\ &\leq \widehat{R}_{\bar{m}}(f) + \beta_{\bar{m}}\end{aligned}$$

for any $f \in \mathcal{F}$, where the second line follows from the definition of $\beta_{\bar{m}}$. We thus have $\bar{f} \in \mathcal{F}_{\bar{m}}$ based on the elimination rule.

3. Fix any $f \in \mathcal{F}_{\bar{m}}$. We have

$$\begin{aligned}\sum_{t=1}^{\tau_{\bar{m}-1}} \mathbb{E}_t[M_t(f)] &\leq 2 \sum_{t=1}^{\tau_{\bar{m}-1}} M_t(f) + C_\delta \\ &= 2\widehat{R}_{\bar{m}}(f) - 2\widehat{R}_{\bar{m}}(f^*) + C_\delta \\ &\leq 2\widehat{R}_{\bar{m}}(f) - 2\widehat{R}_{\bar{m}}(\bar{f}) + 4C_\delta \\ &\leq 2\widehat{R}_{\bar{m}}(f) - 2\widehat{R}_{\bar{m}}(\widehat{f}_{\bar{m}}) + 4C_\delta \\ &\leq 2\beta_{\bar{m}} + 4C_\delta \\ &\leq 4\beta_{\bar{m}},\end{aligned}$$

where the first line follows from [Lemma 6](#); the third line follows from statement 1; the fourth line follows from the fact that $\widehat{f}_{\bar{m}}$ is the minimizer of $\widehat{R}_{\bar{m}}(\cdot)$; and the fifth line follows from the fact that $f \in \mathcal{F}_{\bar{m}}$.

4. Fix any $f \in \mathcal{F}_{\bar{m}}$. We have

$$\begin{aligned}
\sum_{t=1}^{\tau_{\bar{m}-1}} \mathbb{E}_t[Q_t(x_t)(f(x_t) - \bar{f}(x_t))^2] &= \sum_{t=1}^{\tau_{\bar{m}-1}} \mathbb{E}_t[Q_t(x_t)((f(x_t) - f^*(x_t)) + (f^*(x_t) - \bar{f}(x_t)))^2] \\
&\leq 2 \sum_{t=1}^{\tau_{\bar{m}-1}} \mathbb{E}_t[Q_t(x_t)(f(x_t) - f^*(x_t))^2] + 2C_\delta \\
&= 2 \sum_{t=1}^{\tau_{\bar{m}-1}} \mathbb{E}_t[M_t(f)] + 2C_\delta \\
&\leq 8\beta_{\bar{m}} + 2C_\delta \\
&\leq 9\beta_{\bar{m}},
\end{aligned}$$

where the second line follows from $(a+b)^2 \leq 2(a^2+b^2)$ and (the proof of) statement 1 on the second line; and the fourth line follows from statement 3.

5. Fix any $f \in \mathcal{F}_{\bar{m}}$. We have

$$\begin{aligned}
\widehat{R}_{\bar{m}-1}(f) - \widehat{R}_{\bar{m}-1}(\widehat{f}_{\bar{m}-1}) &\leq \widehat{R}_{\bar{m}-1}(f) - \widehat{R}_{\bar{m}-1}(f^*) + \frac{C_\delta}{2} \\
&= \widehat{R}_{\bar{m}}(f) - \widehat{R}_{\bar{m}}(f^*) - \sum_{t=\tau_{\bar{m}-2}+1}^{\tau_{\bar{m}-1}} M_t(f) + \frac{C_\delta}{2} \\
&\leq \widehat{R}_{\bar{m}}(f) - \widehat{R}_{\bar{m}}(\bar{f}) + \frac{3}{2}C_\delta - \sum_{t=\tau_{\bar{m}-2}+1}^{\tau_{\bar{m}-1}} \mathbb{E}_t[M_t(f)]/2 + C_\delta \\
&\leq \widehat{R}_{\bar{m}}(f) - \widehat{R}_{\bar{m}}(\widehat{f}_{\bar{m}}) + \frac{5}{2}C_\delta \\
&\leq \beta_{\bar{m}} + 3C_\delta \\
&\leq \beta_{\bar{m}-1},
\end{aligned}$$

where the first line follows from [Lemma 6](#); the third line follows from statement 1 and [Lemma 6](#); the fourth line follows from the fact that $\widehat{f}_{\bar{m}}$ is the minimizer with respect to $\widehat{R}_{\bar{m}}$ and [Lemma 6](#); the last line follows from the construction of $\beta_{\bar{m}}$. □

We introduce more notations. Denote $(\mathcal{X}, \Sigma, \mathcal{D}_{\mathcal{X}})$ as the (marginal) probability space, and denote $\bar{\mathcal{X}}_m := \{x \in \mathcal{X} : g_m(x) = 1\} \in \Sigma$ be the region where query is requested within epoch m . Under the prerequisites of [Lemma 10](#) and [Lemma 11](#) (i.e., [Eq. \(19\)](#) holds true for epochs $m = 1, 2, \dots, \bar{m}$), we have $\mathcal{F}_m \subseteq \mathcal{F}_{m-1}$ for $m = 1, 2, \dots, \bar{m}$, which leads to $\bar{\mathcal{X}}_m \subseteq \bar{\mathcal{X}}_{m-1}$ for $m = 1, 2, \dots, \bar{m}$. We now define a sub probability measure $\bar{\mu}_m := (\mathcal{D}_{\mathcal{X}})|_{\bar{\mathcal{X}}_m}$ such that $\bar{\mu}_m(\omega) = \mathcal{D}_{\mathcal{X}}(\omega \cap \bar{\mathcal{X}}_m)$ for any $\omega \in \Sigma$. Fix any epoch $m \leq \bar{m}$ and consider any measurable function F (that is $\mathcal{D}_{\mathcal{X}}$ integrable), we have

$$\begin{aligned}
\mathbb{E}_{x \sim \mathcal{D}_{\mathcal{X}}}[\mathbb{1}(g_{\bar{m}}(x) = 1) \cdot F(x)] &= \int_{x \in \bar{\mathcal{X}}_{\bar{m}}} F(x) d\mathcal{D}_{\mathcal{X}}(x) \\
&\leq \int_{x \in \bar{\mathcal{X}}_m} F(x) d\mathcal{D}_{\mathcal{X}}(x) \\
&= \int_{x \in \mathcal{X}} F(x) d\bar{\mu}_m(x) \\
&=: \mathbb{E}_{x \sim \bar{\mu}_m}[F(x)], \tag{22}
\end{aligned}$$

where, by a slightly abuse of notations, we use $\mathbb{E}_{x \sim \mu}[\cdot]$ to denote the integration with any sub probability measure μ . In particular, [Eq. \(22\)](#) holds with equality when $m = \bar{m}$.

Lemma 10. Fix any epoch $\bar{m} \geq 2$. Suppose [Eq. \(19\)](#) holds true for epochs $m = 1, 2, \dots, \bar{m}$, we then have [Eq. \(20\)](#) holds true for epoch $m = \bar{m}$.

Proof. We prove Eq. (20) for epoch $m = \bar{m}$. We know that $\mathbb{1}(g_{\bar{m}}(x) = 1) = \mathbb{1}(g_{\bar{m}}(x) = 1) \cdot \mathbb{1}(w(x; \mathcal{F}_{\bar{m}}) > \gamma/2)$ from Lemma 7. Thus, for any $\check{m} \leq \bar{m}$, we have

$$\begin{aligned} \mathbb{E}_{x \sim \mathcal{D}_{\mathcal{X}}}[\mathbb{1}(g_{\bar{m}}(x) = 1)] &= \mathbb{E}_{x \sim \mathcal{D}_{\mathcal{X}}}[\mathbb{1}(g_{\bar{m}}(x) = 1) \cdot \mathbb{1}(w(x; \mathcal{F}_{\bar{m}}) > \gamma/2)] \\ &\leq \mathbb{E}_{x \sim \bar{\mu}_{\check{m}}}[\mathbb{1}(w(x; \mathcal{F}_{\bar{m}}) > \gamma/2)] \\ &\leq \mathbb{E}_{x \sim \bar{\mu}_{\check{m}}} \left(\mathbb{1} \left(\sup_{f \in \mathcal{F}_{\bar{m}}} |f(x) - \bar{f}(x)| > \gamma/4 \right) \right), \end{aligned} \quad (23)$$

where the second line uses Eq. (22) and the last line follows from the facts that $\bar{f} \in \mathcal{F}_{\bar{m}}$ (by Eq. (19)) and $w(x; \mathcal{F}_{\bar{m}}) > \gamma/2 \implies \exists f \in \mathcal{F}_{\bar{m}}, |f(x) - \bar{f}(x)| > \gamma/4$.

For any time step t , let $m(t)$ denote the epoch where t belongs to. From Eq. (19), we know that, $\forall f \in \mathcal{F}_{\bar{m}}$,

$$\begin{aligned} 9\beta_{\bar{m}} &\geq \sum_{t=1}^{\tau_{\bar{m}-1}} \mathbb{E}_t \left[Q_t (f(x_t) - \bar{f}(x_t))^2 \right] \\ &= \sum_{t=1}^{\tau_{\bar{m}-1}} \mathbb{E}_{x \sim \mathcal{D}_{\mathcal{X}}} \left[\mathbb{1}(g_{m(t)}(x) = 1) \cdot (f(x) - \bar{f}(x))^2 \right] \\ &= \sum_{\check{m}=1}^{\bar{m}-1} n_{\check{m}} \cdot \mathbb{E}_{x \sim \bar{\mu}_{\check{m}}} \left[(f(x) - \bar{f}(x))^2 \right] \\ &= \tau_{\bar{m}-1} \mathbb{E}_{x \sim \bar{\nu}_{\bar{m}}} \left[(f(x) - \bar{f}(x))^2 \right], \end{aligned} \quad (24)$$

where we use $Q_t = g_{m(t)}(x_t) = \mathbb{1}(g_{m(t)}(x) = 1)$ and Eq. (22) on the second line, and define a new sub probability measure

$$\bar{\nu}_{\bar{m}} := \frac{1}{\tau_{\bar{m}-1}} \sum_{\check{m}=1}^{\bar{m}-1} n_{\check{m}} \cdot \bar{\mu}_{\check{m}}$$

on the third line.

Plugging Eq. (24) into Eq. (23) leads to the bound

$$\begin{aligned} &\mathbb{E}_{x \sim \mathcal{D}_{\mathcal{X}}}[\mathbb{1}(g_{\bar{m}}(x) = 1)] \\ &\leq \mathbb{E}_{x \sim \bar{\nu}_{\bar{m}}} \left[\mathbb{1} \left(\exists f \in \mathcal{F}, |f(x) - \bar{f}(x)| > \gamma/4, \mathbb{E}_{x \sim \bar{\nu}_{\bar{m}}} \left[(f(x) - \bar{f}(x))^2 \right] \leq \frac{9\beta_{\bar{m}}}{\tau_{\bar{m}-1}} \right) \right], \end{aligned}$$

where we use the definition of $\bar{\nu}_{\bar{m}}$ again (note that Eq. (23) works with any $\check{m} \leq \bar{m}$). Based on the Definition 7,¹⁰ we then have

$$\begin{aligned} &\mathbb{E}_{x \sim \mathcal{D}_{\mathcal{X}}}[\mathbb{1}(g_{\bar{m}}(x) = 1)] \\ &\leq \frac{144\beta_{\bar{m}}}{\tau_{\bar{m}-1} \gamma^2} \cdot \theta_{\bar{f}}^{\text{val}} \left(\mathcal{F}, \gamma/4, \sqrt{9\beta_{\bar{m}}/2\tau_{\bar{m}-1}} \right) \\ &\leq \frac{144\beta_{\bar{m}}}{\tau_{\bar{m}-1} \gamma^2} \cdot \theta_{\bar{f}}^{\text{val}} \left(\mathcal{F}, \gamma/4, \sqrt{\beta_{\bar{m}}/\tau_{\bar{m}-1}} \right) \\ &\leq \frac{144\beta_{\bar{m}}}{\tau_{\bar{m}-1} \gamma^2} \cdot \bar{\theta}. \end{aligned}$$

□

Lemma 11. Fix any epoch $\bar{m} \geq 2$. Suppose Eq. (19) holds true for epochs $m = 1, 2, \dots, \bar{m}$, we then have Eq. (21) holds true for epoch $m = \bar{m}$.

Proof. We prove Eq. (21) for epoch $m = \bar{m}$. Similar to the proof of Lemma 10, we have

$$\begin{aligned} \mathbb{E}_{x \sim \mathcal{D}_{\mathcal{X}}}[\mathbb{1}(g_{\bar{m}}(x) = 1) \cdot w(x; \mathcal{F}_{\bar{m}})] &= \mathbb{E}_{x \sim \mathcal{D}_{\mathcal{X}}}[\mathbb{1}(g_{\bar{m}}(x) = 1) \cdot \mathbb{1}(w(x; \mathcal{F}_{\bar{m}}) > \gamma/2) \cdot w(x; \mathcal{F}_{\bar{m}})] \\ &\leq \mathbb{E}_{x \sim \bar{\mu}_{\bar{m}}}[\mathbb{1}(w(x; \mathcal{F}_{\bar{m}}) > \gamma/2) \cdot w(x; \mathcal{F}_{\bar{m}})] \end{aligned}$$

¹⁰Note that analyzing with a sub probability measure $\bar{\nu}$ does not cause any problem. See Zhu and Nowak (2022) for detailed discussion.

for any $\tilde{m} \leq \bar{m}$. With $\bar{\nu}_{\bar{m}} := \frac{1}{\tau_{\bar{m}-1}} \sum_{\tilde{m}=1}^{\bar{m}-1} n_{\tilde{m}} \cdot \bar{\mu}_{\tilde{m}}$, we then have

$$\begin{aligned}
& \mathbb{E}_{x \sim \mathcal{D}_{\mathcal{X}}} [\mathbb{1}(g_{\bar{m}}(x) = 1) \cdot w(x; \mathcal{F}_{\bar{m}})] \\
& \leq \mathbb{E}_{x \sim \bar{\nu}_{\bar{m}}} [\mathbb{1}(w(x; \mathcal{F}_{\bar{m}}) > \gamma/2) \cdot w(x; \mathcal{F}_{\bar{m}})] \\
& \leq \mathbb{E}_{x \sim \bar{\nu}_{\bar{m}}} \left[\mathbb{1} \left(\sup_{f \in \mathcal{F}_{\bar{m}}} |f(x) - \bar{f}(x)| > \gamma/4 \right) \cdot \left(\sup_{f, f' \in \mathcal{F}_{\bar{m}}} |f(x) - f'(x)| \right) \right] \\
& \leq 2 \mathbb{E}_{x \sim \bar{\nu}_{\bar{m}}} \left[\mathbb{1} \left(\sup_{f \in \mathcal{F}_{\bar{m}}} |f(x) - \bar{f}(x)| > \gamma/4 \right) \cdot \left(\sup_{f \in \mathcal{F}_{\bar{m}}} |f(x) - \bar{f}(x)| \right) \right] \\
& \leq 2 \int_{\gamma/4}^1 \mathbb{E}_{x \sim \bar{\nu}_{\bar{m}}} \left[\mathbb{1} \left(\sup_{f \in \mathcal{F}_{\bar{m}}} |f(x) - \bar{f}(x)| \geq \omega \right) \right] d\omega \\
& \leq 2 \int_{\gamma/4}^1 \frac{1}{\omega^2} d\omega \cdot \left(\frac{9\beta_{\bar{m}}}{\tau_{\bar{m}-1}} \cdot \theta_{\bar{f}}^{\text{val}}(\mathcal{F}, \gamma/4, \sqrt{9\beta_{\bar{m}}/2\tau_{\bar{m}-1}}) \right) \\
& \leq \frac{72\beta_{\bar{m}}}{\tau_{\bar{m}-1} \gamma} \cdot \theta_{\bar{f}}^{\text{val}}(\mathcal{F}, \gamma/4, \sqrt{\beta_{\bar{m}}/\tau_{\bar{m}-1}}) \\
& \leq \frac{72\beta_{\bar{m}}}{\tau_{\bar{m}-1} \gamma} \cdot \bar{\theta},
\end{aligned}$$

where we follow similar steps as in the proof of [Lemma 10](#) and use some basic arithmetic facts. \square

Lemma 12. *Eq. (19), Eq. (20) and Eq. (21) hold true for all $m \in [M]$.*

Proof. We first notice that, by [Lemma 9](#), [Eq. \(19\)](#) holds true for epochs $\bar{m} = 1, 2$ unconditionally. We also know that, by [Lemma 10](#) and [Lemma 11](#), once [Eq. \(19\)](#) holds true for epochs $m = 1, 2, \dots, \bar{m}$, [Eq. \(20\)](#) and [Eq. \(21\)](#) hold true for epochs $m = \bar{m}$ as well; at the same time, by [Lemma 9](#), once [Eq. \(20\)](#) holds true for epochs $m = 2, 3, \dots, \bar{m}$, [Eq. \(19\)](#) will hold true for epoch $m = \bar{m} + 1$.

We thus can start the induction procedure from $\bar{m} = 2$, and make sure that [Eq. \(19\)](#), [Eq. \(20\)](#) and [Eq. \(21\)](#) hold true for all $m \in [M]$. \square

D.3 Proof of [Theorem 14](#)

Theorem 14. *Suppose $\theta_{\mathcal{F}}^{\text{val}}(\gamma/4) \leq \bar{\theta}$ and the approximation level $\kappa \in (0, \gamma/4]$ satisfies*

$$\left(\frac{432\bar{\theta} \cdot M^2}{\gamma^2} \right) \cdot \kappa^2 \leq \frac{1}{10}. \tag{16}$$

With probability at least $1 - \delta$, [Algorithm 4](#) returns a classifier $\hat{h} : \mathcal{X} \rightarrow \{+1, -1, \perp\}$ with Chow's excess error

$$\text{excess}_{\gamma}(\hat{h}) = O \left(\varepsilon \cdot \log \left(\frac{\bar{\theta} \cdot \text{Pdim}(\mathcal{F})}{\varepsilon \gamma \delta} \right) \right),$$

after querying at most

$$O \left(\frac{M^2 \cdot \text{Pdim}(\mathcal{F}) \cdot \log(T/\delta) \cdot \bar{\theta}}{\gamma^2} \right)$$

labels.

Proof. We analyze under the good event \mathcal{E} defined in [Lemma 6](#), which holds with probability at least $1 - \frac{\delta}{2}$. Note that all supporting lemmas stated in [Appendix D.2.2](#) hold true under this event.

Fix any $m \in [M]$. We analyze the Chow's excess error of \hat{h}_m , which is measurable with respect to $\mathfrak{F}_{\tau_{m-1}}$. For any $x \in \mathcal{X}$, if $g_m(x) = 0$, [Lemma 8](#) implies that $\text{excess}_{\gamma}(\hat{h}_m; x) \leq 0$. If $g_m(x) = 1$, we know that $\hat{h}_m(x) \neq \perp$ and $\frac{1}{2} \in (\text{lcb}(x; \mathcal{F}_m) - \frac{\gamma}{4}, \text{ucb}(x; \mathcal{F}_m) + \frac{\gamma}{4})$. Since $\bar{f} \in \mathcal{F}_m$ by [Lemma 12](#)

(with Eq. (19)) and $\sup_{x \in \mathcal{X}} |\bar{f}(x) - f^*(x)| \leq \kappa \leq \gamma/4$ by construction. The error incurred in this case is upper bounded by

$$\begin{aligned} \text{excess}(\hat{h}_m; x) &\leq 2|f^*(x) - 1/2| \\ &\leq 2\kappa + 2|\bar{f}(x) - 1/2| \\ &\leq 2\kappa + 2w(x; \mathcal{F}_m) + \frac{\gamma}{2} \\ &\leq 4w(x; \mathcal{F}_m), \end{aligned}$$

where we use Lemma 7 in the last line.

Combining these two cases together, we have

$$\text{excess}(\hat{h}_m) \leq 4 \mathbb{E}_{x \sim \mathcal{D}_X} [\mathbb{1}(g_m(x) = 1) \cdot w(x; \mathcal{F}_m)].$$

Take $m = M$ and apply Lemma 12 (and Eq. (21)) leads to the following guarantee.

$$\begin{aligned} \text{excess}(\hat{h}_M) &\leq \frac{576\beta_M}{\tau_{M-1}\gamma} \cdot \theta_{\bar{f}}^{\text{val}}(\mathcal{F}, \gamma/4, \sqrt{\beta_M/\tau_{M-1}}) \\ &\leq O\left(\frac{\text{Pdim}(\mathcal{F}) \log(T/\delta)}{T\gamma} \cdot \bar{\theta}\right) \\ &= O\left(\varepsilon \cdot \log\left(\frac{\bar{\theta} \cdot \text{Pdim}(\mathcal{F})}{\varepsilon\gamma\delta}\right)\right), \end{aligned}$$

where we use the fact that $T = \frac{\bar{\theta} \cdot \text{Pdim}(\mathcal{F})}{\varepsilon\gamma}$.

We now analyze the label complexity (note that the sampling process of Algorithm 4 stops at time $t = \tau_{M-1}$). Note that $\mathbb{E}[\mathbb{1}(Q_t = 1) \mid \mathfrak{F}_{t-1}] = \mathbb{E}_{x \sim \mathcal{D}_X}[\mathbb{1}(g_m(x) = 1)]$ for any epoch $m \geq 2$ and time step t within epoch m . Combine Lemma 12 with Eq. (20) (and Lemma 12) leads to

$$\begin{aligned} \sum_{t=1}^{\tau_{M-1}} \mathbb{1}(Q_t = 1) &\leq \frac{3}{2} \sum_{t=1}^{\tau_{M-1}} \mathbb{E}[\mathbb{1}(Q_t = 1) \mid \mathfrak{F}_{t-1}] + 4 \log(2/\delta) \\ &\leq 3 + \frac{3}{2} \sum_{m=2}^{M-1} \frac{(\tau_m - \tau_{m-1}) \cdot 144\beta_m}{\tau_{m-1}\gamma^2} \cdot \bar{\theta} + 4 \log(2/\delta) \\ &\leq 3 + 4 \log(2/\delta) + O\left(\frac{M^2 \cdot \text{Pdim}(\mathcal{F}) \cdot \log(T/\delta) \cdot \bar{\theta}}{\gamma^2}\right) \\ &= O\left(\frac{M^2 \cdot \text{Pdim}(\mathcal{F}) \cdot \log(T/\delta) \cdot \bar{\theta}}{\gamma^2}\right) \end{aligned}$$

with probability at least $1 - \delta$ (due to another application of Lemma 5 with confidence level $\delta/2$); where we use the fact that $T = \frac{\bar{\theta} \cdot \text{Pdim}(\mathcal{F})}{\varepsilon\gamma}$. \square

E Proof of Theorem 7

We provide prerequisites in Appendix E.1 and the preprocessing procedures in Appendix E.2. We give the proof of Theorem 7 in Appendix E.3.

E.1 Prerequisites

E.1.1 Upper bound on the pseudo dimension

We present a result regarding the approximation and an upper bound on the pseudo dimension (i.e., Definition 6).

Proposition 6. *Suppose $\mathcal{D}_{\mathcal{X}Y} \in \mathcal{P}(\alpha, \beta)$. One can construct a set of neural network regression functions \mathcal{F}_{dnn} such that the following two properties hold simultaneously:*

$$\exists f \in \mathcal{F}_{\text{dnn}} \text{ s.t. } \|f - f^*\|_{\infty} \leq \kappa, \quad \text{and} \quad \text{Pdim}(\mathcal{F}_{\text{dnn}}) \leq c \cdot \kappa^{-\frac{d}{\alpha}} \log^2(\kappa^{-1}),$$

where $c > 0$ is a universal constant.

Proof. The result follows by combining Theorem 3 and Theorem 13. \square

E.1.2 Upper bounds on the value function disagreement coefficient

We derive upper bounds on the value function disagreement coefficient (i.e., Definition 7). We first introduce the (value function) eluder dimension, a complexity measure that is closely related to the value function disagreement coefficient Russo and Van Roy (2013); Foster et al. (2020).

Definition 8 (Value function eluder dimension). *For any $f^* \in \mathcal{F}$ and $\gamma_0 > 0$, let $\check{\epsilon}_{f^*}(\mathcal{F}, \gamma)$ be the length of the longest sequence of data points x^1, \dots, x^m such that for all i , there exists $f^i \in \mathcal{F}$ such that*

$$|f^i(x^i) - f^*(x^i)| > \gamma, \quad \text{and} \quad \sum_{j < i} (f^i(x^j) - f^*(x^j))^2 \leq \gamma^2.$$

The value function eluder dimension is defined as $\epsilon_{f^}(\mathcal{F}, \gamma_0) := \sup_{\gamma > \gamma_0} \check{\epsilon}_{f^*}(\mathcal{F}, \gamma)$.*

The next result shows that the value function disagreement coefficient can be upper bounded by eluder dimension.

Proposition 7 (Foster et al. (2020)). *Suppose \mathcal{F} is a uniform Glivenko-Cantelli class. For any $f^* : \mathcal{X} \rightarrow [0, 1]$ and $\gamma, \varepsilon > 0$, we have $\theta_{f^*}^{\text{val}}(\mathcal{F}, \gamma, \varepsilon) \leq 4 \epsilon_{f^*}(\mathcal{F}, \gamma)$.*

We remark here that the requirement that \mathcal{F} is a uniform Glivenko-Cantelli class is rather weak: It is satisfied as long as \mathcal{F} has finite pseudo dimension (Anthony, 2002).

In the following, we only need to derive upper bounds on the value function eluder dimension, which upper bounds on the value function disagreement coefficient.¹¹ We first define two definitions: (i) the standard definition of covering number (e.g., see Wainwright (2019)), and (ii) a newly-proposed definition of approximate Lipschitzness.

Definition 9. *An ι -covering of a set \mathcal{X} with respect to a metric ρ is a set $\{x_1, \dots, x_N\} \subseteq \mathcal{X}$ such that for each $x \in \mathcal{X}$, there exists some $i \in [N]$ such that $\rho(x, x_i) \leq \iota$. The ι -covering number $\mathcal{N}(\iota; \mathcal{X}, \rho)$ is the cardinality of the smallest ι -cover.*

Definition 10. *We call a function $f : \mathcal{X} \rightarrow \mathbb{R}$ (L, κ) -approximate Lipschitz if*

$$|f(x) - f(x')| \leq L \cdot \|x - x'\|_2 + \kappa$$

for any $x, x' \in \mathcal{X}$.

We next provide upper bounds on value function eluder dimension and value function disagreement coefficient.

Theorem 15. *Suppose \mathcal{F} is a set of $(L, \kappa/4)$ -approximate Lipschitz functions. For any $\kappa' \geq \kappa$, we have $\sup_{f \in \mathcal{F}} \epsilon_f(\mathcal{F}, \kappa') \leq 17 \cdot \mathcal{N}(\frac{\kappa'}{8L}; \mathcal{X}, \|\cdot\|_2)$.*

Proof. Fix any $f \in \mathcal{F}$ and $\bar{\kappa} \geq \kappa'$. We first give upper bounds on $\check{\epsilon}_f(\mathcal{F}, \bar{\kappa})$.

We construct $\mathcal{G} := \mathcal{F} - f$, which is a set of $(2L, \kappa/2)$ -Lipschitz functions. Fix any eluder sequence x^1, \dots, x^m at scale $\bar{\kappa}$ and any $\check{x} \in \mathcal{X}$. We claim that $|\{x_j\}_{j \leq m} \cap \mathcal{S}| \leq 17$ where $\mathcal{S} := \{x \in \mathcal{X} : \|x - \check{x}\|_2 \leq \frac{\bar{\kappa}}{8L}\}$. Suppose $\{x_j\}_{j \leq m} \cap \mathcal{S} = x_{j_1}, \dots, x_{j_k}$ (j_i is ordered based on the ordering of $\{x_j\}_{j \leq m}$). Since x^{j_k} is added into the eluder sequence, there must exist a $g^{j_k} \in \mathcal{G}$ such that

$$|g^{j_k}(x^{j_k})| > \bar{\kappa}, \quad \text{and} \quad \sum_{j < j_k} (g^{j_k}(x^j))^2 \leq \bar{\kappa}^2. \quad (25)$$

Since g^{j_k} is $(2L, \kappa/2)$ -Lipschitz, $\bar{\kappa} \geq \kappa' \geq \kappa$ and $x^{j_k} \in \mathcal{S}$, we must have $g^{j_k}(x) \geq \frac{\bar{\kappa}}{4}$ for any $x \in \mathcal{S}$. As a result, we must have $|\{x_j\}_{j < j_k} \cap \mathcal{S}^i| \leq 16$ as otherwise the second constraint in Eq. (25) will be violated. We cover the space \mathcal{X} with $\mathcal{N}(\frac{\bar{\kappa}}{8L}; \mathcal{X}, \|\cdot\|_2)$ balls of radius $\frac{\bar{\kappa}}{8L}$. Since the eluder sequence contains at most 17 data points within each ball, we know that $\check{\epsilon}_f(\mathcal{F}, \bar{\kappa}) \leq 17 \cdot \mathcal{N}(\frac{\bar{\kappa}}{8L}; \mathcal{X}, \|\cdot\|_2)$.

The desired result follows by noticing that $17 \cdot \mathcal{N}(\frac{\bar{\kappa}}{8L}; \mathcal{X}, \|\cdot\|_2)$ is non-increasing in $\bar{\kappa}$. \square

¹¹We focus on Euclidean geometry on \mathcal{X} (i.e., using $\|\cdot\|_2$ norm) in deriving the upper bound. Slightly tighter bounds might be possible with other norms.

Corollary 1. Suppose $\mathcal{X} \subseteq \mathbb{B}_r^d := \{x \in \mathbb{R}^d : \|x\|_2 \leq r\}$ and \mathcal{F} is a set of $(L, \kappa/4)$ -approximate Lipschitz functions. For any $\kappa' \geq \kappa$, we have $\theta_{\mathcal{F}}^{\text{val}}(\kappa') := \sup_{f \in \mathcal{F}, \iota > 0} \theta_f^{\text{val}}(\mathcal{F}, \kappa', \iota) \leq c \cdot (\frac{Lr}{\kappa'})^d$ with a universal constant $c > 0$.

Proof. It is well-known that $\mathcal{N}(\iota; \mathbb{B}_r^d, \|\cdot\|_2) \leq (1 + 2r/\iota)^d$ (Wainwright, 2019). The desired result thus follows from combining Theorem 15 with Proposition 7. \square

E.2 The preprocessing step: Clipping and filtering

Let $\eta : \mathcal{X} \rightarrow [0, 1]$ denote the true conditional probability and \mathcal{F}_{dnn} denote a set of neural network regression functions (e.g., constructed based on Theorem 3). We assume that (i) η is L -Lipschitz, and (ii) there exists a $f \in \mathcal{F}$ such that $\|f - \eta\|_\infty \leq \kappa$ for some approximation factor $\kappa > 0$. We present the preprocessing step below in Algorithm 5.

Algorithm 5 The Preprocessing Step: Clipping and Filtering

Input: A set of regression functions \mathcal{F} , Lipschitz parameter $L > 0$, approximation factor $\kappa > 0$.

1: **Clipping.** Set $\check{\mathcal{F}} := \{\check{f} : f \in \mathcal{F}\}$, where, for any $f \in \mathcal{F}$, we denote

$$\check{f}(x) := \begin{cases} 1, & \text{if } f(x) \geq 1; \\ 0, & \text{if } f(x) \leq 0; \\ f(x), & \text{o.w.} \end{cases}$$

2: **Filtering.** Set $\tilde{\mathcal{F}} := \{\check{f} \in \check{\mathcal{F}} : \check{f} \text{ is } (L, 2\kappa)\text{-approximate Lipschitz}\}$

3: **Return** $\tilde{\mathcal{F}}$.

Proposition 8. Suppose η is L -Lipschitz and \mathcal{F}_{dnn} is a set of neural networks (of the same architecture) with W parameters arranged in L layers such that there exists a $f \in \mathcal{F}_{\text{dnn}}$ with $\|f - \eta\|_\infty \leq \kappa$. Let $\tilde{\mathcal{F}}_{\text{dnn}}$ be the set of functions obtained by applying Algorithm 5 on \mathcal{F}_{dnn} , we then have (i) $\text{Pdim}(\tilde{\mathcal{F}}_{\text{dnn}}) = O(WL \log(W))$, and (ii) there exists a $\check{f} \in \tilde{\mathcal{F}}_{\text{dnn}}$ such that $\|\check{f} - \eta\|_\infty \leq \kappa$.

Proof. Suppose f is a neural network function, we first notice that the “clipping” step can be implemented by adding one additional layer with $O(1)$ additional parameters for each neural network function. More specifically, fix any $f : \mathcal{X} \rightarrow \mathbb{R}$, we can set $\check{f}(x) := \text{ReLU}(f(x)) - \text{ReLU}(f(x) - 1)$. Set $\check{\mathcal{F}}_{\text{dnn}} := \{\check{f} : f \in \mathcal{F}_{\text{dnn}}\}$, we then have $\text{Pdim}(\check{\mathcal{F}}_{\text{dnn}}) = O(WL \log(W))$ based on Theorem 13. Let $\tilde{\mathcal{F}}_{\text{dnn}}$ be the filtered version of $\check{\mathcal{F}}_{\text{dnn}}$. Since $\tilde{\mathcal{F}}_{\text{dnn}} \subseteq \check{\mathcal{F}}_{\text{dnn}}$, we have $\text{Pdim}(\tilde{\mathcal{F}}_{\text{dnn}}) = O(WL \log(W))$.

Since $\eta : \mathcal{X} \rightarrow [0, 1]$, we have $\|\check{f} - \eta\|_\infty \leq \|f - \eta\|_\infty$, which implies that there must exists a $\check{f} \in \check{\mathcal{F}}_{\text{dnn}}$ such that $\|\check{f} - \eta\|_\infty \leq \kappa$. To prove the second statement, it suffices to show that the $\check{f} \in \check{\mathcal{F}}$ that achieves κ approximation error is not removed in the “filtering” step, i.e., \check{f} is $(L, 2\kappa)$ -approximate Lipschitz. For any $x, x' \in \mathcal{X}$, we have

$$\begin{aligned} |\check{f}(x) - \check{f}(x')| &= |\check{f}(x) - \eta(x) + \eta(x) - \eta(x') + \eta(x') - \check{f}(x')| \\ &\leq L\|x - x'\|_2 + 2\kappa, \end{aligned}$$

where we use the L -Lipschitzness of η and the fact that $\|\check{f} - \eta\|_\infty \leq \kappa$. \square

Proposition 9. Suppose η is L -Lipschitz and $\mathcal{X} \subseteq \mathbb{B}_r^d$. Fix any $\kappa \in (0, \gamma/32]$. There exists a set of neural network regression functions \mathcal{F}_{dnn} such that the followings hold simultaneously.

1. $\text{Pdim}(\mathcal{F}_{\text{dnn}}) \leq c \cdot \kappa^{-\frac{d}{\alpha}} \log^2(\kappa^{-1})$ with a universal constant $c > 0$.
2. There exists a $\bar{f} \in \mathcal{F}_{\text{dnn}}$ such that $\|\bar{f} - \eta\|_\infty \leq \kappa$.
3. $\theta_{\mathcal{F}_{\text{dnn}}}^{\text{val}}(\gamma/4) := \sup_{f \in \mathcal{F}_{\text{dnn}}, \iota > 0} \theta_f^{\text{val}}(\mathcal{F}_{\text{dnn}}, \gamma/4, \iota) \leq c' \cdot (\frac{Lr}{\gamma})^d$ with a universal constant $c' > 0$.

Proof. Let \mathcal{F}_{dnn} be obtained by (i) invoking Theorem 3 with approximation level κ , and (ii) invoking Algorithm 5 on the set of functions obtained in step (i). The first two statements follow from

[Proposition 8](#), and the third statement follows from [Corollary 1](#) (note that to achieve guarantees for disagreement coefficient at level $\gamma/4$, we need to have $\kappa \leq \gamma/32$ when invoking [Theorem 3](#)). \square

E.3 Proof of [Theorem 7](#)

Theorem 7. Fix any $\varepsilon, \delta, \gamma > 0$. With probability at least $1 - \delta$, [Algorithm 2](#) (with an appropriate initialization at line 1) returns a classifier \hat{h} with Chow's excess error $\tilde{O}(\varepsilon)$ after querying $\text{poly}(\frac{1}{\gamma}) \cdot \text{polylog}(\frac{1}{\varepsilon\delta})$ labels.

Proof. Let line 1 of [Algorithm 2](#) be the set of neural networks \mathcal{F}_{dnn} generated from [Proposition 9](#) with approximation level $\kappa \in (0, \gamma/32]$ (and constants c, c' specified therein). To apply results derived in [Theorem 14](#), we need to satisfy [Eq. \(16\)](#), i.e., specifying an approximation level $\kappa \in (0, \gamma/32]$ such that the following holds true

$$\frac{1}{\kappa^2} \geq \frac{4320 \cdot c' \cdot (\frac{Lr}{\gamma})^d \cdot \left(\left[\log_2 \left(\frac{c' \cdot (\frac{Lr}{\gamma})^d \cdot c \cdot (\kappa^{-\frac{d}{\alpha}} \log^2(\kappa^{-1}))}{\varepsilon \gamma} \right) \right] \right)^2}{\gamma^2}$$

For the setting we considered, i.e., $\mathcal{X} = [0, 1]^d$ and $\eta \in \mathcal{W}_1^{\alpha, \infty}(\mathcal{X})$, we have $r = \sqrt{d} = O(1)$ and $L \leq \sqrt{d} = O(1)$ (e.g., see [Theorem 4.1](#) in [Heinonen \(2005\)](#)).¹² We thus only need to select a $\kappa \in (0, \gamma/32]$ such that

$$\frac{1}{\kappa} \geq \bar{c} \cdot \left(\frac{1}{\gamma} \right)^{\frac{d}{2}+1} \cdot \left(\log \frac{1}{\varepsilon \gamma} + \log \frac{1}{\kappa} \right),$$

with a universal constant $\bar{c} > 0$ (that is possibly d -dependent and α -dependent). Since $x \geq 2a \log a \implies x \geq a \log x$ for any $a > 0$, we can select a $\kappa > 0$ such that

$$\frac{1}{\kappa} = \check{c} \cdot \left(\frac{1}{\gamma} \right)^{\frac{d}{2}+1} \cdot \log \frac{1}{\varepsilon \gamma}$$

with a universal constant $\check{c} > 0$. With such choice of κ , from [Proposition 9](#), we have

$$\text{Pdim}(\mathcal{F}_{\text{dnn}}) = O \left(\left(\frac{1}{\gamma} \right)^{\frac{d^2+d}{2\alpha}} \cdot \text{polylog} \left(\frac{1}{\varepsilon \gamma} \right) \right).$$

Plugging this bound on $\text{Pdim}(\mathcal{F}_{\text{dnn}})$ and the upper bound on $\theta_{\mathcal{F}_{\text{dnn}}}^{\text{val}}(\gamma/4)$ from [Proposition 9](#) into the guarantee of [Theorem 14](#) leads to $\text{excess}_\gamma(\hat{h}) = O(\varepsilon \cdot \log(\frac{1}{\varepsilon \gamma \delta}))$ after querying

$$O \left(\left(\frac{1}{\gamma} \right)^{d+2+\frac{d^2+d}{2\alpha}} \cdot \text{polylog} \left(\frac{1}{\varepsilon \gamma \delta} \right) \right)$$

labels. \square

F Other omitted details for [Section 3](#)

We discuss the proper abstention property of classifier learned in [Algorithm 2](#) the computational efficiency of [Algorithm 2](#) in [Appendix F.1](#). We provide the proof of [Theorem 8](#) in [Appendix F.2](#).

F.1 Proper abstention and computational efficiency

F.1.1 Proper abstention

We first recall the definition of *proper abstention* proposed in [Zhu and Nowak \(2022\)](#).

¹²Recall that we ignore constants that can be potentially α -dependent and d -dependent.

Definition 11 (Proper abstention). *A classifier $\hat{h} : \mathcal{X} \rightarrow \mathcal{Y} \cup \{\perp\}$ enjoys proper abstention if and only if it abstains in regions where abstention is indeed the optimal choice, i.e., $\{x \in \mathcal{X} : \hat{h}(x) = \perp\} \subseteq \{x \in \mathcal{X} : \eta(x) \in [\frac{1}{2} - \gamma, \frac{1}{2} + \gamma]\} =: \mathcal{X}_\gamma$.*

We next show that the classifier \hat{h} returned by [Algorithm 4](#) enjoys the proper abstention property. We also convert the abstaining classifier $\hat{h} : \mathcal{X} \rightarrow \mathcal{Y} \cup \{\perp\}$ into a standard classifier $\check{h} : \mathcal{X} \rightarrow \mathcal{Y}$ and quantify its standard excess error. The conversion is through randomizing the prediction of \hat{h} over its abstention region, i.e., if $\hat{h}(x) = \perp$, then its randomized version $\check{h}(x)$ predicts $+1/-1$ with equal probability ([Puchkin and Zhivotovskiy, 2021](#)).

Proposition 10. *The classifier \hat{h} returned by [Algorithm 4](#) enjoys proper abstention. With randomization over the abstention region, we have the following upper bound on its standard excess error*

$$\text{err}(\check{h}) - \text{err}(h^*) = \text{err}_\gamma(\hat{h}) - \text{err}(h^*) + \gamma \cdot \mathbb{P}_{x \sim \mathcal{D}_X}(x \in \mathcal{X}_\gamma). \quad (26)$$

Proof. The proper abstention property of \hat{h} returned by [Algorithm 4](#) is achieved via conservation: \hat{h} will avoid abstention unless it is absolutely sure that abstention is the optimal choice (also see the proof of [Lemma 8](#)).

Let $\check{h} : \mathcal{X} \rightarrow \mathcal{Y}$ be the randomized version of $\bar{h} : \mathcal{X} \rightarrow \{+1, -1, \perp\}$ (over the abstention region $\{x \in \mathcal{X} : \hat{h}(x) = \perp\} \subseteq \mathcal{X}_\gamma$). We can see that, compared to the Chow’s abstention error $1/2 - \gamma$, the additional error incurred over the abstention region is exactly $\gamma \cdot \mathbb{P}_{x \sim \mathcal{D}_X}(x \in \mathcal{X}_\gamma)$. We thus have

$$\text{err}(\hat{h}) - \text{err}(h^*) \leq \text{err}_\gamma(\hat{h}) - \text{err}(h^*) + \gamma \cdot \mathbb{P}_{x \sim \mathcal{D}_X}(x \in \mathcal{X}_\gamma). \quad \square$$

To characterize the standard excess error of classifier with proper abstention, we only need to upper bound the term $\mathbb{P}_{x \sim \mathcal{D}_X}(x \in \mathcal{X}_\gamma)$, which does *not* depend on the (random) classifier \hat{h} . Instead, it only depends on the marginal distribution.

We next introduce the Massart ([Massart and Nédélec, 2006](#)), which can be viewed as the extreme version of the Tsybakov noise by sending $\beta \rightarrow \infty$.

Definition 12 (Massart noise). *The marginal distribution \mathcal{D}_X satisfies the Massart noise condition with parameter $\tau_0 > 0$ if $\mathbb{P}_{x \sim \mathcal{D}_X}(|\eta(x) - 1/2| \leq \tau_0) = 0$.*

Proposition 11. *Suppose Massart noise holds. By setting the abstention parameter $\gamma = \tau_0$ in [Algorithm 4](#) (and randomization over the abstention region), with probability at least $1 - \delta$, we obtain a classifier with standard excess error $\tilde{O}(\varepsilon)$ after querying $\text{poly}(\frac{1}{\tau_0}) \cdot \text{polylog}(\frac{1}{\varepsilon\delta})$ labels.*

Proof. This is a direct consequence of [Theorem 7](#) and [Proposition 10](#). □

F.1.2 Computational efficiency

We discuss the efficient implementation of [Algorithm 4](#) and its computational complexity in the section. The computational efficiency of [Algorithm 4](#) mainly follows from the analysis in [Zhu and Nowak \(2022\)](#). We provide the discussion here for completeness.

Regression oracle. We introduce the regression oracle over the set of initialized neural networks \mathcal{F}_{dnn} (line 1 at [Algorithm 2](#)). Given any set \mathcal{S} of weighted examples $(w, x, y) \in \mathbb{R}_+ \times \mathcal{X} \times \mathcal{Y}$ as input, the regression oracle outputs

$$\hat{f}_{\text{dnn}} := \arg \min_{f \in \mathcal{F}_{\text{dnn}}} \sum_{(w, x, y) \in \mathcal{S}} w(f(x) - y)^2.$$

While the exact computational complexity of such oracle with a set of neural networks remains elusive, in practice, running stochastic gradient descent often leads to great approximations. We quantify the computational complexity in terms of the number of calls to the regression oracle. Any future analysis on such oracle can be incorporated into our guarantees.

We first state some known results in computing the confidence intervals with respect to a general set of regression functions \mathcal{F} .

Proposition 12 (Krishnamurthy et al. (2017); Foster et al. (2018, 2020)). *Consider the setting studied in Algorithm 4. Fix any epoch $m \in [M]$ and denote $\mathcal{B}_m := \{(x_t, Q_t, y_t)\}_{t=1}^{\tau_m-1}$. Fix any $\iota > 0$. For any data point $x \in \mathcal{X}$, there exists algorithms $\mathbf{Alg}_{\text{lcb}}$ and $\mathbf{Alg}_{\text{ucb}}$ that certify*

$$\begin{aligned} \text{lcb}(x; \mathcal{F}_m) - \iota &\leq \mathbf{Alg}_{\text{lcb}}(x; \mathcal{B}_m, \beta_m, \iota) \leq \text{lcb}(x; \mathcal{F}_m) \quad \text{and} \\ \text{ucb}(x; \mathcal{F}_m) &\leq \mathbf{Alg}_{\text{ucb}}(x; \mathcal{B}_m, \beta_m, \iota) \leq \text{ucb}(x; \mathcal{F}_m) + \iota. \end{aligned}$$

The algorithms take $O(\frac{1}{\iota^2} \log \frac{1}{\iota})$ calls of the regression oracle for general \mathcal{F} and take $O(\log \frac{1}{\iota})$ calls of the regression oracle if \mathcal{F} is convex and closed under pointwise convergence.

Proof. See Algorithm 2 in Krishnamurthy et al. (2017) for the general case; and Algorithm 3 in Foster et al. (2018) for the case when \mathcal{F} is convex and closed under pointwise convergence. \square

We now state the computational guarantee of Algorithm 4, given the regression oracle introduced above.

Theorem 16. *Algorithm 4 can be efficiently implemented via the regression oracle and enjoys the same theoretical guarantees stated in Theorem 7. The number of oracle calls needed is $\text{poly}(\frac{1}{\gamma}) \cdot \frac{1}{\epsilon}$; the per-example inference time of the learned \hat{h}_M is $\tilde{O}(\frac{1}{\gamma^2} \cdot \text{polylog}(\frac{1}{\epsilon\gamma}))$ for general \mathcal{F} , and $\tilde{O}(\text{polylog}(\frac{1}{\epsilon\gamma}))$ when \mathcal{F} is convex.*

Proof. Fix any epoch $m \in [M]$. Denote $\bar{\iota} := \frac{\gamma}{8M}$ and $\iota_m := \frac{(M-m)\gamma}{8M}$. With any observed $x \in \mathcal{X}$, we construct the approximated confidence intervals $\widehat{\text{lcb}}(x; \mathcal{F}_m)$ and $\widehat{\text{ucb}}(x; \mathcal{F}_m)$ as follows.

$$\widehat{\text{lcb}}(x; \mathcal{F}_m) := \mathbf{Alg}_{\text{lcb}}(x; \mathcal{B}_m, \beta_m, \bar{\iota}) - \iota_m \quad \text{and} \quad \widehat{\text{ucb}}(x; \mathcal{F}_m) := \mathbf{Alg}_{\text{ucb}}(x; \mathcal{B}_m, \beta_m, \bar{\iota}) + \iota_m.$$

For efficient implementation of Algorithm 4, we replace $\text{lcb}(x; \mathcal{F}_m)$ and $\text{ucb}(x; \mathcal{F}_m)$ with $\widehat{\text{lcb}}(x; \mathcal{F}_m)$ and $\widehat{\text{ucb}}(x; \mathcal{F}_m)$ in the construction of \hat{h}_m and g_m .

Based on Proposition 12, we know that

$$\begin{aligned} \text{lcb}(x; \mathcal{F}_m) - \iota_m - \bar{\iota} &\leq \widehat{\text{lcb}}(x; \mathcal{F}_m) \leq \text{lcb}(x; \mathcal{F}_m) - \iota_m \quad \text{and} \\ \text{ucb}(x; \mathcal{F}_m) + \iota_m &\leq \widehat{\text{ucb}}(x; \mathcal{F}_m) \leq \text{ucb}(x; \mathcal{F}_m) + \iota_m + \bar{\iota}. \end{aligned}$$

Since $\iota_m + \bar{\iota} \leq \frac{\gamma}{8}$ for any $m \in [M]$, the guarantee stated in Lemma 7 can be modified as $g_m(x) = 1 \implies w(x; \mathcal{F}_m) \geq \frac{\gamma}{4}$. The guarantee stated in Lemma 8 also holds true since we have $\widehat{\text{lcb}}(x; \mathcal{F}_m) \leq \text{lcb}(x; \mathcal{F}_m)$ and $\widehat{\text{ucb}}(x; \mathcal{F}_m) \geq \text{ucb}(x; \mathcal{F}_m)$ by construction. Suppose $\mathcal{F}_m \subseteq \mathcal{F}_{m-1}$ (as in Lemma 9), we have

$$\begin{aligned} \widehat{\text{lcb}}(x; \mathcal{F}_m) &\geq \text{lcb}(x; \mathcal{F}_m) - \iota_m - \bar{\iota} \geq \text{lcb}(x; \mathcal{F}_{m-1}) - \iota_{m-1} \geq \widehat{\text{lcb}}(x; \mathcal{F}_{m-1}) \quad \text{and} \\ \widehat{\text{ucb}}(x; \mathcal{F}_m) &\leq \text{ucb}(x; \mathcal{F}_m) + \iota_m + \bar{\iota} \leq \text{ucb}(x; \mathcal{F}_{m-1}) + \iota_{m-1} \leq \widehat{\text{ucb}}(x; \mathcal{F}_{m-1}), \end{aligned}$$

which ensures that $\mathbb{1}(g_m(x) = 1) \leq \mathbb{1}(g_{m-1}(x) = 1)$. Thus, the inductive lemmas appearing in Appendix D.2.2 can be proved similarly with changes only in constant terms (also change the constant terms in the definition of θ and in Eq. (16), since $\frac{\gamma}{2}$ is replaced by $\frac{\gamma}{4}$ in Lemma 7). As a result, the guarantees stated in Theorem 14 (and Theorem 7) hold true with changes only in constant terms.

We now discuss the computational complexity of the efficient implementation. At the beginning of each epoch m . We use one oracle call to compute $\hat{f}_m := \arg \min_{f \in \mathcal{F}} \sum_{t=1}^{\tau_m-1} Q_t(f(x_t) - y_t)^2$. The main computational cost comes from computing $\widehat{\text{lcb}}$ and $\widehat{\text{ucb}}$ at each time step. We take $\iota = \bar{\iota} := \frac{\gamma}{8M}$ into Proposition 12, which leads to $O(\frac{(\log T)^2}{\gamma^2} \cdot \log(\frac{\log T}{\gamma}))$ calls of the regression oracle for general \mathcal{F} and $O(\log(\frac{\log T}{\gamma}))$ calls of the regression oracle for any convex \mathcal{F} that is closed under pointwise convergence. This also serves as the per-example inference time for \hat{h}_M . The total computational cost of Algorithm 4 is then derived by multiplying the per-round cost by T and plugging $T = \frac{\theta \text{Pdim}(\mathcal{F})}{\epsilon\gamma} = \tilde{O}(\text{poly}(\frac{1}{\gamma}) \cdot \frac{1}{\epsilon})$ into the bound. \square

E.2 Proof of Theorem 8

For ease of construction, we suppose the instance space is $\mathcal{X} = \mathbb{B}_1^d := \{x \in \mathbb{R}^d : \|x\|_2 \leq 1\}$. Part of our construction is inspired by Li et al. (2021).

Theorem 8. Fix any $\gamma \in (0, 1/8)$. For any accuracy level ε sufficiently small, there exists a problem instance such that (1) $\eta \in \mathcal{W}_1^{1,\infty}(\mathcal{X})$ and is of the form $\eta(x) := \text{ReLU}(\langle w, x \rangle + a) + b$; and (2) for any active learning algorithm, it takes at least $\gamma^{-\Omega(d)}$ labels to identify an ε -optimal classifier, for either standard excess error or Chow's excess error (with parameter γ).

Proof. Fix any $\gamma \in (0, 1/8)$. We first claim that we can find a discrete subset $\bar{\mathcal{X}} \subseteq \mathcal{X}$ with cardinality $|\bar{\mathcal{X}}| \geq (1/8\gamma)^{d/2}$ such that $\|x_i\|_2 = 1$ and $\langle x_1, x_2 \rangle \leq 1 - 4\gamma$ for any $x_i \in \bar{\mathcal{X}}$. To prove this, we first notice that $\|x_1 - x_2\|_2 \geq \tau \iff \langle x_1, x_2 \rangle \leq 1 - \tau^2/2$. Since the τ -packing number on the unit sphere is at least $(1/\tau)^d$, setting $\tau = \sqrt{8\gamma}$ leads to the desired claim.

We set $\mathcal{D}_{\mathcal{X}} := \text{unif}(\bar{\mathcal{X}})$ and $\mathcal{F}_{\text{dnn}} := \{\text{ReLU}(\langle w, \cdot \rangle - (1 - 4\gamma)) + (1/2 - 2\gamma) : w \in \bar{\mathcal{X}}\}$. We have $\mathcal{F}_{\text{dnn}} \subseteq \mathcal{W}_1^{1,\infty}(\mathcal{X})$ since $\|w\|_2 \leq 1$ for any $w \in \bar{\mathcal{X}}$. We randomly select a $w^* \in \bar{\mathcal{X}}$ and set $f^*(\cdot) = \eta(\cdot) = \text{ReLU}(\langle w^*, \cdot \rangle - (1 - 4\gamma)) + (1/2 - 2\gamma)$. We assume that the labeling feedback is the conditional expectation, i.e., $\eta(x)$ is provided if x is queried. We see that $f^*(x) = 1/2 - 2\gamma$ for any $x \in \mathcal{X}$ but $x \neq w^*$, and $f^*(w^*) = 1/2 + 2\gamma$. We can see that mistakenly select the wrong $\hat{f} \neq f^*$ leads to $\frac{\gamma}{4} \cdot \frac{2}{|\bar{\mathcal{X}}|} = \frac{\gamma}{2|\bar{\mathcal{X}}|}$ excess error. Note that the excess error holds true in both standard excess error and Chow's excess error (with parameter γ) since $\mathcal{D}_{\mathcal{X}}(x \in \mathcal{X} : \eta(x) \in [1/2 - \gamma, 1/2 + \gamma]) = 0$ by construction.

We suppose the desired excess error ε is sufficiently small (e.g., $\varepsilon \leq \frac{\gamma}{8|\bar{\mathcal{X}}|}$). We now show that, with label complexity at most $K := \lfloor |\bar{\mathcal{X}}|/2 \rfloor = \Omega(\gamma^{-d/2})$, any active learning algorithm will, in expectation, pick a classifier that has $\Omega(\varepsilon)$ excess error. Since the worst case error of any randomized algorithm is lower bounded by the expected error of the best deterministic algorithm against a input distribution (Yao, 1977), we only need to analyze a deterministic learner. We set the input distribution as the uniform distribution over instances with parameter $w^* \in \bar{\mathcal{X}}$. For any deterministic algorithm, we use $s := (x_{i_1}, \dots, x_{i_K})$ to denote the data points queried under the constraint that at most K labels can be queried. We denote $\hat{f} \in \mathcal{F}$ as the learned classifier conditioned on s . Since $w^* \sim \text{unif}(\bar{\mathcal{X}})$, we know that, with probability at least $\frac{1}{2}$, $w^* \notin s$. Conditioned on that event, we know that, with probability at least $\frac{1}{2}$, the learner will output $\hat{f} \neq f^*$ since more than half of the data points remains unqueried. The deterministic algorithm thus outputs the wrong $\hat{f} \neq f^*$ with probability at least $\frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4}$, which has $\frac{\gamma}{2|\bar{\mathcal{X}}|}$ excess error as previously discussed. When $\varepsilon \leq \frac{\gamma}{8|\bar{\mathcal{X}}|}$, this leads to $\Omega(\varepsilon)$ excess error in expectation. \square

G Omitted details for Section 4

We provide mathematical backgrounds for the Radon BV^2 space in Appendix G.1, derive approximation results and passive learning results in Appendix G.2, and derive active learning results in Appendix G.3.

G.1 The Radon BV^2 space

We provide explicit definition of the $\|f\|_{\mathcal{R}\text{BV}^2(\mathcal{X})}$ and associated mathematical backgrounds in this section. Also see Ongie et al. (2020); Parhi and Nowak (2021, 2022a,b); Unser (2022) for more discussions.

We first introduce the *Radon transform* of a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ as

$$\mathcal{R}\{f\}(\gamma, t) := \int_{\{x: \gamma^\top x = t\}} f(x) \, ds(x), \quad (\gamma, t) \in \mathbb{S}^{d-1} \times \mathbb{R},$$

where s denotes the surface measure on the hyperplane $\{x : \gamma^\top x = t\}$. The Radon domain is parameterized by a *direction* $\gamma \in \mathbb{S}^{d-1}$ and an *offset* $t \in \mathbb{R}$. We also introduce the *ramp filter* as

$$\Lambda^{d-1} := (-\partial_t^2)^{\frac{d-1}{2}},$$

where ∂_t denotes the partial derivative with respect to the offset variable, t , of the Radon domain, and the fractional powers are defined in terms of Riesz potentials.

With the above preparations, we can define the $\mathcal{R} \text{TV}^2$ -seminorm as

$$\mathcal{R} \text{TV}^2(f) := c_d \|\partial_t^2 \Lambda^{d-1} \mathcal{R}f\|_{\mathcal{M}(\mathbb{S}^{d-1} \times \mathbb{R})},$$

where $c_d = 1/(2(2\pi)^{d-1})$ is a dimension-dependent constant, and $\|\cdot\|_{\mathcal{M}(\mathbb{S}^{d-1} \times \mathbb{R})}$ denotes the *total variation norm* (in terms of measures) over the bounded domain $\mathbb{S}^{d-1} \times \mathbb{R}$. The $\mathcal{R} \text{BV}^2$ norm of f over \mathbb{R}^d is defined as

$$\|f\|_{\mathcal{R} \text{BV}^2(\mathbb{R}^d)} := \mathcal{R} \text{TV}^2(f) + |f(0)| + \sum_{k=1}^d |f(e_k) - f(0)|,$$

where $\{e_k\}_{k=1}^d$ denotes the canonical basis of \mathbb{R}^d . The $\mathcal{R} \text{BV}^2(\mathbb{R}^d)$ space is then defined as

$$\mathcal{R} \text{BV}^2(\mathbb{R}^d) := \{f \in L^{\infty,1}(\mathbb{R}^d) : \mathcal{R} \text{BV}^2(f) < \infty\},$$

where $L^{\infty,1}(\mathbb{R}^d)$ is the Banach space of functions mapping $\mathbb{R}^d \rightarrow \mathbb{R}$ of at most linear growth. To define the $\mathcal{R} \text{BV}^2$ norm of f over a bounded domain $\mathcal{X} \subseteq \mathbb{R}^d$, we use the standard approach of considering restrictions of functions in $\mathcal{R} \text{BV}^2(\mathbb{R}^d)$, i.e.,

$$\|f\|_{\mathcal{R} \text{BV}^2(\mathcal{X})} := \inf_{g \in \mathcal{R} \text{BV}^2(\mathbb{R}^d)} \|g\|_{\mathcal{R} \text{BV}^2(\mathbb{R}^d)} \quad \text{s.t.} \quad g|_{\mathcal{X}} = f.$$

In the rest of [Appendix G](#), we use $\mathcal{P}(\beta)$ to denote the set of distributions that satisfy (1) Tsybakov noise condition with parameter $\beta \geq 0$; and (2) $\eta \in \mathcal{R} \text{BV}_1^2(\mathcal{X})$.

G.2 Approximation and passive learning results

Proposition 13. *Suppose $\mathcal{D}_{\mathcal{X}\mathcal{Y}} \in \mathcal{P}(\beta)$. One can construct a set of neural network classifier \mathcal{H}_{dnn} such that the following two properties hold simultaneously:*

$$\min_{h \in \mathcal{H}_{\text{dnn}}} \text{err}(h) - \text{err}(h^*) = O(\varepsilon) \quad \text{and} \quad \text{VCdim}(\mathcal{H}_{\text{dnn}}) = \tilde{O}(\varepsilon^{-\frac{2d}{(1+\beta)(d+3)}}).$$

Proof. We take $\kappa = \varepsilon^{\frac{1}{1+\beta}}$ in [Theorem 9](#) to construct a set of neural network classifiers \mathcal{H}_{dnn} with $W = O(\varepsilon^{-\frac{2d}{(1+\beta)(d+3)}})$ total parameters arranged in $L = O(1)$ layers. According to [Theorem 4](#), we know

$$\text{VCdim}(\mathcal{H}_{\text{dnn}}) = O(\varepsilon^{-\frac{2d}{(1+\beta)(d+3)}} \cdot \log(\varepsilon^{-1})) = \tilde{O}(\varepsilon^{-\frac{2d}{(1+\beta)(d+3)}}).$$

We now show that there exists a classifier $\bar{h} \in \mathcal{H}_{\text{dnn}}$ with small excess error. Let $\bar{h} = h_{\bar{f}}$ be the classifier such that $\|\bar{f} - \eta\|_{\infty} \leq \kappa$. We can see that

$$\begin{aligned} \text{excess}(\bar{h}) &= \mathbb{E}[\mathbb{1}(\bar{h}(x) \neq y) - \mathbb{1}(h^*(x) \neq y)] \\ &= \mathbb{E}[|2\eta(x) - 1| \cdot \mathbb{1}(\bar{h}(x) \neq h^*(x))] \\ &\leq 2\kappa \cdot \mathbb{P}_{x \sim \mathcal{D}_{\mathcal{X}}}(x \in \mathcal{X} : |\eta(x) - 1/2| \leq \kappa) \\ &= O(\kappa^{1+\beta}) \\ &= O(\varepsilon), \end{aligned}$$

where the third line follows from the fact that \bar{h} and h^* disagrees only within region $\{x \in \mathcal{X} : |\eta(x) - 1/2| \leq \kappa\}$ and the incurred error is at most 2κ on each disagreed data point. The fourth line follows from the Tsybakov noise condition and the last line follows from the selection of κ . \square

Theorem 17. *Suppose $\mathcal{D}_{\mathcal{X}\mathcal{Y}} \in \mathcal{P}(\beta)$. Fix any $\varepsilon, \delta > 0$. Let \mathcal{H}_{dnn} be the set of neural network classifiers constructed in [Proposition 13](#). With $n = \tilde{O}(\varepsilon^{-\frac{4d+6+\beta(d+3)}{(1+\beta)(d+3)}})$ i.i.d. sampled data points, with probability at least $1 - \delta$, the empirical risk minimizer $\hat{h} \in \mathcal{H}_{\text{dnn}}$ achieves excess error $O(\varepsilon)$.*

Proof. [Proposition 13](#) certifies $\min_{h \in \mathcal{H}_{\text{dnn}}} \text{err}(h) - \text{err}(h^*) = O(\varepsilon)$ and $\text{VCdim}(\mathcal{H}_{\text{dnn}}) = O\left(\varepsilon^{-\frac{2d}{(1+\beta)(d+3)}} \cdot \log(\varepsilon^{-1})\right)$. Take $\rho = 1$ in [Theorem 11](#), leads to

$$\text{err}(\hat{h}) - \text{err}(h^*) \leq O\left(\varepsilon + \left(\varepsilon^{-\frac{2d}{(1+\beta)(d+3)}} \cdot \log(\varepsilon^{-1}) \cdot \frac{\log n}{n}\right)^{\frac{1+\beta}{2+\beta}} + \frac{\log \delta^{-1}}{n}\right),$$

Taking $n = O\left(\varepsilon^{-\frac{4d+6+\beta(d+3)}{(1+\beta)(d+3)}} \cdot \log(\varepsilon^{-1}) + \varepsilon^{-1} \cdot \log(\delta^{-1})\right) = \tilde{O}\left(\varepsilon^{-\frac{4d+6+\beta(d+3)}{(1+\beta)(d+3)}}\right)$ thus ensures that $\text{err}(\hat{h}) - \text{err}(h^*) = O(\varepsilon)$. \square

G.3 Active learning results

Theorem 10. *Suppose $\eta \in \mathcal{R} \text{BV}_1^2(\mathcal{X})$ and the Tsybakov noise condition is satisfied with parameter $\beta \geq 0$. Fix any $\varepsilon, \delta > 0$. There exists an algorithm such that, with probability at least $1 - \delta$, it learns a classifier $\hat{h} \in \mathcal{H}_{\text{dnn}}$ with excess error $\tilde{O}(\varepsilon)$ after querying $\tilde{O}(\theta_{\mathcal{H}_{\text{dnn}}}(\varepsilon^{\frac{\beta}{1+\beta}}) \cdot \varepsilon^{-\frac{4d+6}{(1+\beta)(d+3)}})$ labels.*

Proof. Construct \mathcal{H}_{dnn} based on [Proposition 13](#) such that $\min_{h \in \mathcal{H}_{\text{dnn}}} \text{err}(h) - \text{err}(h^*) = O(\varepsilon)$ and $\text{VCdim}(\mathcal{H}_{\text{dnn}}) = \tilde{O}\left(\varepsilon^{-\frac{2d}{(1+\beta)(d+3)}}\right)$. Taking such \mathcal{H}_{dnn} as the initialization of [Algorithm 3](#) (line 1) and applying [Theorem 12](#) leads to the desired result. \square

To derive deep active learning guarantee with abstention in the Radon BV^2 space, we first present two supporting results below.

Proposition 14. *Suppose $\mathcal{D}_{\mathcal{X}\mathcal{Y}} \in \mathcal{P}(\beta)$. One can construct a set of neural network regression functions \mathcal{F}_{dnn} such that the following two properties hold simultaneously:*

$$\exists f \in \mathcal{F}_{\text{dnn}} \text{ s.t. } \|f - f^*\|_\infty \leq \kappa, \quad \text{and} \quad \text{Pdim}(\mathcal{F}_{\text{dnn}}) \leq c \cdot \kappa^{-\frac{2d}{d+3}} \log^2(\kappa^{-1}),$$

where $c > 0$ is a universal constant.

Proof. The result follows by combining [Theorem 9](#) and [Theorem 13](#). \square

Proposition 15. *Suppose η is L -Lipschitz and $\mathcal{X} \subseteq \mathbb{B}_r^d$. Fix any $\kappa \in (0, \gamma/32]$. There exists a set of neural network regression functions \mathcal{F}_{dnn} such that the followings hold simultaneously.*

1. $\text{Pdim}(\mathcal{F}_{\text{dnn}}) \leq c \cdot \kappa^{-\frac{2d}{d+3}} \log^2(\kappa^{-1})$ with a universal constant $c > 0$.
2. There exists a $\bar{f} \in \mathcal{F}_{\text{dnn}}$ such that $\|\bar{f} - \eta\|_\infty \leq \kappa$.
3. $\theta_{\mathcal{F}_{\text{dnn}}}^{\text{val}}(\gamma/4) := \sup_{f \in \mathcal{F}_{\text{dnn}}, \iota > 0} \theta_f^{\text{val}}(\mathcal{F}_{\text{dnn}}, \gamma/4, \iota) \leq c' \cdot \left(\frac{Lr}{\gamma}\right)^d$ with a universal constant $c' > 0$.

Proof. The implementation and proof are similar to those in [Proposition 9](#), except we use [Proposition 14](#) instead of [Proposition 6](#). \square

We now state and prove deep active learning guarantees in the Radon BV^2 space.

Theorem 18. *Suppose $\eta \in \mathcal{R} \text{BV}_1^2(\mathcal{X})$. Fix any $\varepsilon, \delta, \gamma > 0$. There exists an algorithm such that, with probability at least $1 - \delta$, it learns a classifier \hat{h} with Chow's excess error $\tilde{O}(\varepsilon)$ after querying $\text{poly}(\frac{1}{\gamma}) \cdot \text{polylog}(\frac{1}{\varepsilon\delta})$ labels.*

Proof. The result is obtained by applying [Algorithm 4](#) with line 1 be the set of neural networks \mathcal{F}_{dnn} generated from [Proposition 15](#) with approximation level $\kappa \in (0, \gamma/32]$ (and constants c, c' specified therein). The rest of the proof proceeds in a similar way as the proof [Theorem 7](#). Since we have $r = 1$ and $L \leq 1$ ([Parhi and Nowak, 2022b](#)), we only need to choose a $\kappa > 0$ such that

$$\frac{1}{\kappa} = \check{c} \cdot \left(\frac{1}{\gamma}\right)^{\frac{d}{2}+1} \cdot \log \frac{1}{\varepsilon\gamma}$$

with a universal constant $\check{c} > 0$. With such choice of κ , we have

$$\text{Pdim}(\mathcal{F}_{\text{dnn}}) = O\left(\left(\frac{1}{\gamma}\right)^{\frac{d^2+2d}{d+3}} \text{polylog}\left(\frac{1}{\varepsilon\gamma}\right)\right).$$

Plugging this bound on $\text{Pdim}(\mathcal{F}_{\text{dnn}})$ and the upper bound on $\theta_{\mathcal{F}_{\text{dnn}}}^{\text{val}}(\gamma/4)$ from [Proposition 15](#) into the guarantee of [Theorem 14](#) leads to $\text{excess}_\gamma(\hat{h}) = O(\varepsilon \cdot \log(\frac{1}{\varepsilon\gamma\delta}))$ after querying

$$O\left(\left(\frac{1}{\gamma}\right)^{d+2+\frac{d^2+2d}{d+3}} \cdot \text{polylog}\left(\frac{1}{\varepsilon\gamma\delta}\right)\right)$$

labels. □