Active Learning with Neural Networks: Insights from Nonparametric Statistics

Anonymous Author(s)
Affiliation
Address
email

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 minimax optimal label complexity guarantees for deep active learning. The key insight is to study deep active learning from the nonparametric classification perspective. Under common 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 a computationally efficient deep active learning algorithm that guarantees $\mathsf{polylog}(\frac{1}{\varepsilon})$ label complexity, without any low noise assumptions. We also provide extensions of 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 the natural function spaces associated with neural networks.

1 Introduction

We study active learning with hypothesis class being a set of neural networks, 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 labeling examples (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 success 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 that have made the attempts (Karzand and Nowak, 2020; Wang et al., 2021) at theoretically quantifying active learning gains in neural network settings. 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-seperated 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;

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

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 fill this gap by providing the first near minimax optimal label complexity guarantees for deep active learning. We obtain insights from the nonparametric setting where the conditional probability 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 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 (1) 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 (2) developing novel extensions of agnostic active learning algorithms (Balcan et al., 2009; Hanneke, 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 $\Theta(\mathrm{poly}(\frac{1}{\varepsilon}))$ (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 algorithm showing that deep active learning provably achieves $\mathrm{polylog}(\frac{1}{\varepsilon})$ label complexity when equipped with an abstention option (Chow, 1970). Such learning paradigm not only guarantees an exponential improvement 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 leave the decision-making process to a medical expert.

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 $h^\star$ can thus be expressed as $h^\star(x) := \mathrm{sign}(2\eta(x) - 1)$. For any classifier $h : \mathcal{X} \to \mathcal{Y}$, its (standard) error is calculated as $\mathrm{err}(h) := \mathbb{P}_{(x,y)\sim\mathcal{D}_{\mathcal{X}\mathcal{Y}}}(h(x) \neq y)$; and its (standard) excess error is defined as $\mathsf{excess}(h) := \mathrm{err}(h) - \mathrm{err}(h^\star)$. Our goal is to learn an accurate classifier with a small amount 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 $f : \mathcal{X} \to \mathbb{R}$ is defined as $\|f\|_{\mathcal{W}^{\alpha,\infty}} := \max_{\bar{\alpha},|\bar{\alpha}|\leq\alpha} \operatorname{ess\,sup}_{x\in\mathcal{X}} |\mathsf{D}^{\alpha} f(x)|$, where $\alpha = (\alpha_1, \ldots, \alpha_d)$, $|\alpha| = \sum_{i=1}^d \alpha_i$ and $\mathsf{D}^{\alpha} f$ 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 ReLU (Rectified Linear Unit) activation function, which is defined as $\mathsf{ReLU}(x) := \max\{x, 0\}$. Each neural network $f_{\mathsf{dnn}} : \mathcal{X} \to \mathbb{R}$ 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 $\mathsf{ReLU}(\sum_{i=1}^N w_i \bar{x}_i + b)$ 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 approximatability 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 coefficient in the active setting) in the rich nonparametric regimes.[1] We provide information 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 common Tsybakov noise condition with smoothness parameter $\beta \geq 0$ (Tsybakov, 2004). Let $\mathcal{H}_{\mathsf{dnn}}$ be an appropriately chosen set of neural network classifiers and denote $\theta_{\mathcal{H}_{\mathsf{dnn}}}(\varepsilon)$ as the disagreement coefficient (Hanneke, 2007, 2014) at level ε. We develop the following label complexity guarantee for deep active learning.

Theorem 1 (Informal)**.** *There exists an algorithm that returns a neural network classifier $\widehat{h} \in \mathcal{H}_{\mathsf{dnn}}$ with excess error $\widetilde{O}(\varepsilon)$ after querying $\widetilde{O}(\theta_{\mathcal{H}_{\mathsf{dnn}}}(\varepsilon^{\frac{\beta}{1+\beta}}) \cdot \varepsilon^{-\frac{d+2\alpha}{\alpha+\alpha\beta}})$ 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 $\theta_{\mathcal{H}_{\mathsf{dnn}}}(\varepsilon) \leq \frac{1}{\varepsilon}$ by definition, the label complexity presented in Theorem 1 is never worse than the passive learning rates $\widetilde{\Theta}(\varepsilon^{-\frac{d+2\alpha+\alpha\beta}{\alpha+\alpha\beta}})$ (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: $\theta_{\mathcal{H}_{\mathsf{dnn}}}(\varepsilon) = o(\varepsilon)$ (implying strict improvement over passive learning) and $\theta_{\mathcal{H}_{\mathsf{dnn}}}(\varepsilon) = o(1)$ (implying matching active learning lower bound).

In Section 3, we develop label complexity guarantee 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 predication, we have the following guarantee for learning with a set of neural networks.

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

The above $\mathrm{polylog}(\frac{1}{\varepsilon})$ label complexity bound is achieved *without any low noise assumptions*. Such exponential improvement theoretically justify the great empirical performance 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, our algorithm (Algorithm 2) developed for Theorem 2 can be efficiently implemented in $\widetilde{O}(\frac{1}{\varepsilon})$ time, given a least square regression oracle over a set of neural networks (which can be approximated by running stochastic gradient descent in practice).

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, our results extend the understanding of how approximation affects active learning to achieve fast rates (e.g., under Tsybakov noise in Theorem 1), and we identify conditions that greatly relax the (previous) approximation requirement in the learning with abstention setup (to develop Theorem 2). We also provide our active learning guarantees (under approximation) in generic forms (in Appendix B and Appendix D), which can be of independent interests. Other technical contributions including analysis of disagreement coefficient with a set of neural networks (Appendix C.2) and an important deep active learning lower bound (Theorem 8). We defer all proofs to the appendix due to lack of space.

Benefited from our generic approaches/analysis of deep active learning, in Section 4, we extend our results into learning smooth functions in the Radon BV^2 space (Ongie et al., 2020; Parhi and Nowak, 2021a,b; Unser, 2022), which is recently proposed as a natural space to analyze neural networks.

[1] As 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 estimation over exponentially my sub-regions (Castro and Nowak, 2008; Minsker, 2012; Locatelli et al., 2017, 2018; Shekhar et al., 2021; Kpotufe et al., 2021).

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 (Hanneke, 2007; Balcan et al., 2007, 2009; Zhang and Chaudhuri, 2014; Huang et al., 2015) 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; Friedman, 2009; Hanneke, 2014). 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 mean(or 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 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). More specifically, Puchkin and Zhivotovskiy (2021); Zhu and Nowak (2022) analyze active learning under Chow's excess error in the parametric regime, and show that one can alway $\mathrm{polylog}(\frac{1}{\varepsilon})$ label complexity. Shekhar et al. (2021) study nonparametric active learning under a slightly different version of the Chow's excess error, and propose algorithms with $\mathrm{poly}(\frac{1}{\varepsilon})$ label complexity; their algorithm follows similar logic as previous nonparametric active learning algorithms (e.g., Minsker (2012); Locatelli et al. (2017)) and is partition-based.

Inspired by the success of deep learning in the passive regime, active learning with deep 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 fill the gap between practice and theory by providing the first near minimax 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, 2021b) and VC/Pesudo dimension analyses of neural networks of a given structure (Bartlett et al., 2019).

2 Label complexity of deep active learning

We rigorously 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). Our main active learning guarantees are provided in Section 2.3, where we also discuss importance implications of our results.

2.1 Tsybakov noise condition

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

Definition 1 (Tsybakov noise)**.** *The marginal distribution $\mathcal{D}_{\mathcal{X}}$ satisfies the Tsybakov noise condition with parameter $\beta \geq 0$ and universal constants $c \geq 1$ if, $\forall t > 0$,*

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

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: (1) the smoothness conditions introduced in Section 1.1 with parameter $\alpha > 0$; and (2) the Tsybakov low noise condition (i.e., Definition 1) with parameter $\beta \geq 0$. We assume $\mathcal{D}_{\mathcal{XY}} \in \mathcal{P}(\alpha, \beta)$ 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 $g : \mathcal{X} \to \mathbb{R}$ and any error tolerance $\kappa > 0$, there exists a large enough neural network f_{dnn} such that $\|f_{\mathsf{dnn}} - g\|_\infty := \sup_{x \in \mathcal{X}} |f_{\mathsf{dnn}}(x) - g(x)| \leq \kappa$. Recently, *non-asympototic* approximation rates by ReLU neural networks have been developed for smooth functions in the Sobolev space, which we re-state in the following.[2]

Theorem 3 (Yarotsky (2017))**.** *Fix any $\kappa > 0$. For any $f^\star \equiv \eta \in \mathcal{W}_1^{\alpha,\infty}([-1,1]^d)$, 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 $\|f_{\mathsf{dnn}} - f^\star\|_\infty \leq \kappa$.*

The architecture of the neural network f_{dnn} appearing in the above theorem is independent of $f^\star$. Let $\mathcal{F}_{\mathsf{dnn}}$ 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., $\mathcal{H}_{\mathsf{dnn}} := \{h_f := \mathrm{sign}(2f(x) - 1) : f \in \mathcal{F}_{\mathsf{dnn}}\}$. The next results 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 $\mathcal{H}_{\mathsf{dnn}}$ be a set of neural network classifiers of the same architecture and with W parameters arranged in L layers. We then have*

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

With these recently established tools, we can construct a set of neural network classifiers $\mathcal{H}_{\mathsf{dnn}}$ such that (1) the best in-class classifier $\check{h} \in \mathcal{H}_{\mathsf{dnn}}$ has small excess error; and (2) $\mathcal{H}_{\mathsf{dnn}}$ has well-controlled VC dimension that is proportional to smooth/noise parameters. More specifically, we have the following proposition.

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

$$\inf_{h \in \mathcal{H}_{\mathsf{dnn}}} \mathrm{err}(h) - \mathrm{err}(h^\star) = O(\varepsilon) \quad \textit{and} \quad \mathrm{VCdim}(\mathcal{H}_{\mathsf{dnn}}) = \widetilde{O}(\varepsilon^{-\frac{d}{\alpha(1+\beta)}}).$$

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 $\mathcal{H}_{\mathsf{dnn}}$ with well-controlled VC dimension. As a warm-up, we first analyze the label complexity of such procedure in the passive regime (with fast rate).

Theorem 5. *Suppose $\mathcal{D}_{\mathcal{XY}} \in \mathcal{P}(\alpha, \beta)$. Fix any $\varepsilon, \delta > 0$. Let $\mathcal{H}_{\mathsf{dnn}}$ be the set of neural network classifiers constructed in Proposition 1. With $n = \widetilde{O}(\varepsilon^{-\frac{d+2\alpha+\alpha\beta}{\alpha(1+\beta)}})$ i.i.d. sampled points, with probability at least $1 - \delta$, the empirical risk minimizer $\widehat{h} \in \mathcal{H}_{\mathsf{dnn}}$ 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 our proposed learning procedure *with a set of neural networks* is minimax optimal (up to logarithmic factors).

2.3 Active learning and label complexity guarantees

The passive learning procedure presented in the previous section treats every label as equally important, 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 with the RobustCAL algorithm (Balcan et al., 2009; Hanneke, 2007, 2014), which is built upon the seminal CAL algorithm by Cohn et al. (1994).

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

We now introduce RobustCAL (Algorithm 1), which runs in epochs of geometrically increasing lengths. At the beginning of epoch m, based on previously *labeled* data points, RobustCAL updates a set of active classifier $\mathcal{H}_m$ such that, with high probability, the best classifier $\check{h} := \arg\min_{h \in \mathcal{H}} \mathrm{err}(h)$ remains *uneliminated*. Within each epoch m, RobustCAL 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., $\mathsf{DIS}(\mathcal{H}_m) := \{x \in \mathcal{X} : \exists h_1, h_2 \in \mathcal{H}_m \text{ s.t. } h_1(x) \neq h_2(x)\}$. RobustCAL returns any classifier $\widehat{h} \in \mathcal{H}_m$ that remains uneliminated after $M - 1$ epoch.

Algorithm 1 RobustCAL

Input: Accuracy level $\varepsilon \in (0, 1)$, confidence level $\delta \in (0, 1)$, and an appropriately chosen set of neural network classifiers $\mathcal{H} \equiv \mathcal{H}_{\mathsf{dnn}}$.

Set $T = \varepsilon^{-\frac{2+\beta}{1+\beta}} \cdot \mathrm{VCdim}(\mathcal{H})$, and $M = \lceil \log_2 T \rceil$.
Define $\tau_m := 2^m$ for $m \geq 1$ and $\tau_0 = 0$; $\mathcal{H}_0 = \mathcal{H}$.
Define $\rho_m := O\left(\left(\frac{\mathrm{VCdim}(\mathcal{H}) \cdot \log(\tau_{m-1}) \cdot \log(M/\delta)}{\tau_{m-1}}\right)^{\frac{1+\beta}{2+\beta}}\right)$ for $m \geq 2$ and $\rho_1 := 1$.
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 $\sum_{t=1}^{0} \ldots = 0$.
for epoch $m = 1, 2, \ldots, M$ **do**
  Update active set $\mathcal{H}_m := \left\{h \in \mathcal{H}_{m-1} : \widehat{R}_m(h) \leq \inf_{h \in \mathcal{H}_{m-1}} \widehat{R}_m(h) + \tau_{m-1} \cdot \rho_m\right\}$
  if epoch $m = M$ **then**
    Return any classifier $\widehat{h} \in \mathcal{H}_M$.
  for time $t = \tau_{m-1} + 1, \ldots, \tau_m$ **do**
    Observe $x_t \sim \mathcal{D}_{\mathcal{X}}$. Set $Q_t = \mathbb{1}(x_t \in \mathsf{DIS}(\mathcal{H}_m))$.
    if $Q_t = 1$ **then**
      Query the label y_t of x_t.

Since RobustCAL 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. The rate of convergence is captured by the *disagreement coefficient* (Hanneke, 2007, 2014), which we defined below.

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

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

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 $\theta_{\mathcal{H}}(\varepsilon_0) := \sup_{h \in \mathcal{H}} \theta_{\mathcal{H},h}(\varepsilon_0)$.

While using the well-known RobustCAL algorithm, our results are established on novel analysis. In particular, to achieve fast rate (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 (see Appendix B), which leads to the following guarantees.

Theorem 6. *Suppose $\mathcal{D}_{\mathcal{X}\mathcal{Y}} \in \mathcal{P}(\alpha, \beta)$. Fix any $\varepsilon, \delta > 0$. Let $\mathcal{H}_{\mathsf{dnn}}$ be the set of neural network classifiers constructed in Proposition 1. With probability at least $1 - \delta$, Algorithm 1 returns a classifier $\widehat{h} \in \mathcal{H}_{\mathsf{dnn}}$ with excess error $\widetilde{O}(\varepsilon)$ after querying $\widetilde{O}(\theta_{\mathcal{H}_{\mathsf{dnn}}}(\varepsilon^{\frac{\beta}{1+\beta}}) \cdot \varepsilon^{-\frac{d+2\alpha}{\alpha+\alpha\beta}})$ labels.*

We next discuss the label complexity of deep active learning achieved in Theorem 6.

- Ignoring the dependence on disagreement coefficient, the label complexity appearing in Theorem 6 matches, up to logarithmic factors, the lower bound $\Omega(\varepsilon^{-\frac{d+2\alpha}{\alpha+\alpha\beta}})$ 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., $\widetilde{\Theta}(\varepsilon^{-\frac{d+2\alpha+\alpha\beta}{\alpha(1+\beta)}})$ since $\theta_{\mathcal{H}_{\mathsf{dnn}}}(\varepsilon^{\frac{\beta}{1+\beta}}) \leq \varepsilon^{-\frac{\beta}{1+\beta}}$.
- We also identify cases when $\theta_{\mathcal{H}_{\mathsf{dnn}}}(\varepsilon^{\frac{\beta}{1+\beta}}) = o(\varepsilon^{-\frac{\beta}{1+\beta}})$, indicating *strict* improvement over passive learning (e.g., when $\mathcal{D}_{\mathcal{X}}$ is supported on countably many data points); and when $\theta_{\mathcal{H}_{\mathsf{dnn}}}(\varepsilon^{\frac{\beta}{1+\beta}}) =$

$O(1)$, indicating matching the minimax active lower bound (e.g., when $\mathcal{D}_{\mathcal{X}\mathcal{Y}}$ satisfies conditions such as *decomposibility* defined in Definition 4. See Appendix C.2 for detailed discussion).[3]

Our Theorem 6 has the following two important implications, 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 side, our guarantees are minimax optimal (up to disagreement coefficient), which, for the first time, close the gap between theory and practice in the field of 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 in fact learn with a set of neural network to achieve near minimax optimal label complexities.

3 Deep active learning with abstention: Exponential speedups

We analyze the label complexity of deep active learning under Chow's excess error (Chow, 1970). We first introduce the concepts of abstention and Chow's excess error in Section 3.1, and then provide our label complexity guarantees in Section 3.2.

3.1 Active learning without low noise conditions

The previous section analyze 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 such setting as follows.

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

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

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 $\mathsf{excess}_\gamma(\widehat{h}) := \mathrm{err}_\gamma(\widehat{h}) - \mathrm{err}(h^\star)$ (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 $\mathrm{polylog}(\frac{1}{\varepsilon})$ 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] We 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 discussion. We left a comprehensive analysis of the disagreement coefficient with a set of neural network classifiers for future work.

3.2 Exponential speedups with abstention

In this section, we work with a set of neural network *regression functions* $\mathcal{F}_{\mathsf{dnn}} : \mathcal{X} \to [0,1]$ (that approximates η) and then *construct* classifiers $h : \mathcal{X} \to \mathcal{Y} \cup \{\perp\}$ *with an additional abstention action*. We construct $\mathcal{F}_{\mathsf{dnn}}$ based on approximation results developed in Theorem 3, with additional preprocessing steps that: (1) clipping $f_{\mathsf{dnn}} : \mathcal{X} \to \mathbb{R}$ into $f'_{\mathsf{dnn}} : \mathcal{X} \to [0,1]$ (since we obviously have $\eta(x) \in [0,1]$); and (2) filtering out $f_{\mathsf{dnn}} \in \mathcal{F}_{\mathsf{dnn}}$ that are clearly not a good approximation of η.[4]

Algorithm 2 Efficient Deep Active Learning with Abstention

Input: Accuracy level $\varepsilon \in (0,1)$, abstention parameter $\gamma \in (0, 1/2)$, confidence level $\delta \in (0,1)$, and an appropriately chosen set of neural network regression functions $\mathcal{F} \equiv \mathcal{F}_{\mathsf{dnn}}$.

Set $T = \frac{\theta \cdot \mathrm{Pdim}(\mathcal{F})}{\varepsilon\,\gamma}$, $M = \lceil \log_2 T \rceil$, and $C_\delta = O(\mathrm{Pdim}(\mathcal{F}) \cdot \log(T/\delta))$.
Define $\mathcal{F}_0 := \mathcal{F}$; $\tau_m := 2^m$ for $m \geq 1$, $\tau_0 := 0$, and $\beta_m := 3(M - m + 1)C_\delta$.
Define $\widehat{R}_m(f) := \sum_{t=1}^{\tau_{m-1}} Q_t(\widehat{f}(x_t) - y_t)^2$ with the convention that $\sum_{t=1}^{0} \ldots = 0$.
for epoch $m = 1, 2, \ldots, M$ **do**
  Get $\widehat{f}_m = \arg\min_{f \in \mathcal{F}} \sum_{t=1}^{\tau_{m-1}} Q_t(f(x_t) - y_t)^2$.
  (Implicitely) Construct active set $\mathcal{F}_m = \left\{ f \in \mathcal{F} : \widehat{R}_m(f) \leq \widehat{R}_m(\widehat{f}_m) + \beta_m \right\}$.
  Construct classifier $\widehat{h}_m : \mathcal{X} \to \{+1, -1, \perp\}$ as

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

  and query function $g_m(x) := \mathbb{1}\left(\frac{1}{2} \in \left(\mathsf{lcb}(x; \mathcal{F}_m) - \frac{\gamma}{4}, \mathsf{ucb}(x; \mathcal{F}_m) + \frac{\gamma}{4}\right)\right) \cdot \mathbb{1}(\widehat{h}_m(x) \neq \perp)$.
  if epoch $m = M$ **then**
    Return classifier $\widehat{h}_M$.
  for time $t = \tau_{m-1} + 1, \ldots, \tau_m$ **do**
    Observe $x_t \sim \mathcal{D}_{\mathcal{X}}$. Set $Q_t = g_m(x_t)$.
    if $Q_t = 1$ **then**
      Query the label y_t of x_t.

Our algorithm is presented in Algorithm 2, which is adapted from the algorithm developed in Zhu and Nowak (2022). Algorithm 2 runs in epochs of geometrically increasing lengths. At the beginning of epoch $m \in [M]$, Algorithm 2 (implicitly) constructs an active set of regression functions $\mathcal{F}_m$ that are "close" to the true conditional probability η. For any $x \sim \mathcal{D}_{\mathcal{X}}$, Algorithm 2 constructs a lower bound $\mathsf{lcb}(x; \mathcal{F}_m) := \inf_{f \in \mathcal{F}_m} f(x)$ and an upper bound $\mathsf{ucb}(x; \mathcal{F}_m) := \sup_{f \in \mathcal{F}_m} f(x)$ as a confidence range of $\eta(x)$ (based on $\mathcal{F}_m$). An empirical classifier $\widehat{h}_m : \mathcal{X} \to \{+1, -1, \perp\}$ and a query function $g_m : \mathcal{X} \to \{0, 1\}$ are then constructed based on the confidence range (and the abstention parameter γ). For any time step t within epoch m, Algorithm 2 queries the label of the observed data point x_t if and only if $g_m(x_t) = 1$. Algorithm 2 returns $\widehat{h}_M$ as the learned classifier.

The algorithm provided in Zhu and Nowak (2022) requires the existence of a $\overline{f} \in \mathcal{F}$ such that $\|\overline{f} - \eta\|_\infty \leq \varepsilon$ (to achieve ε Chow's excess error). Such approximation requirement directly leads to $\mathrm{poly}(\frac{1}{\varepsilon})$ label complexity *in the nonparametric setting*, which is unacceptable. We identify conditions to bypass such requirement and provide sharper analysis of concentration results (see Appendix D). These novel extensions are important in the setting of active learning with neural networks, which help us derive the following deep active learning guarantees.

Theorem 7. *Fix any $\varepsilon, \delta, \gamma > 0$. Let $\mathcal{F}_{\mathsf{dnn}}$ be an appropriately chosen set of neural network regression function. With probability at least $1 - \delta$, Algorithm 2 returns a classifier $\widehat{h}$ with Chow's excess error $\widetilde{O}(\varepsilon)$ after querying $\mathrm{poly}(\frac{1}{\gamma}) \cdot \mathrm{polylog}(\frac{1}{\varepsilon\,\delta})$ labels.*

We discuss two important aspects of Algorithm 2/Theorem 7, i.e., exponential savings and computational efficiency, as follows (more detailed discussions deferred to Appendix E.2, due to lack of space).

[4] See Appendix D.1.3 for the detailed preprocessing steps.

- **Exponential speedups.** Theorem 7 shows that, equipped with an abstention option, deep active learning enjoys $\mathrm{polylog}(\frac{1}{\varepsilon})$ 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 property further implies $\mathrm{polylog}(\frac{1}{\varepsilon})$ 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 a set of neural networks, e.g., $\widehat{f} = \arg\min_{f \in \mathcal{F}_{\mathsf{dnn}}} \sum_{i=1}^{n} (f(x_i) - y_i)^2$.[5] Algorithm 2 can be then *efficiently* implemented with $\mathrm{poly}(\frac{1}{\gamma}) \cdot \frac{1}{\varepsilon}$ oracle calls.

While the label complexity obtained in Theorem 7 has desired dependence on $\mathrm{polylog}(\frac{1}{\varepsilon})$, its dependence on γ can be of order $\gamma^{-\mathrm{poly}(d)}$. 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 $\eta(x) := \mathsf{ReLU}(\langle w, x\rangle + a) + b$; 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, 2021a,b; Unser, 2022).[6]

The Radon BV^2 space. The Radon BV^2 unit ball over domain $\mathcal{X}$ is defined as $\mathscr{R}\mathsf{BV}_1^2(\mathcal{X}) := \{f : \|f\|_{\mathscr{R}\mathsf{BV}^2(\mathcal{X})} \leq 1\}$, where $\|f\|_{\mathscr{R}\mathsf{BV}^2(\mathcal{X})}$ denotes the Radon BV^2 norm of f over domain $\mathcal{X}$.[7] Following Parhi and Nowak (2021b), we assume $\mathcal{X} = \{x \in \mathbb{R}^d : \|x\|_2 \leq 1\}$ and $\eta \in \mathscr{R}\mathsf{BV}_1^2(\mathcal{X})$.

The Radon BV^2 space naturally contains neural networks of the form $f_{\mathsf{dnn}}(x) = \sum_{k=1}^{K} v_i \cdot \mathsf{ReLU}(w_i^\top x + b_i)$. 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., $\eta \equiv f_{\mathsf{dnn}}$), approximating η up to κ from a Sobolev perspective requires $\widetilde{O}(\kappa^{-d})$ 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 (2021b))**.** *Fix any $\kappa > 0$. For any $f^\star \in \mathscr{R}\mathsf{BV}_1^2(\mathcal{X})$, there exists a one-hidden layer neural network f_{dnn} of width $K = O(\kappa^{-\frac{2d}{d+3}})$ such that $\|f^\star - f_{\mathsf{dnn}}\|_\infty \leq \kappa$.*

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 $\eta \in \mathscr{R}\mathsf{BV}_1^2(\mathcal{X})$ and the Tsybakov noise condition is satisfied with parameter $\beta \geq 0$. Fix any $\varepsilon, \delta > 0$. Let $\mathcal{H}_{\mathsf{dnn}}$ be an appropriately chosen set of neural networks. With probability at least $1 - \delta$, Algorithm 1 returns a classifier $\widehat{h} \in \mathcal{H}_{\mathsf{dnn}}$ with excess error $\widetilde{O}(\varepsilon)$ after querying $\widetilde{O}(\theta_{\mathcal{H}_{\mathsf{dnn}}}(\varepsilon^{\frac{\beta}{1+\beta}}) \cdot \varepsilon^{-\frac{4d+6}{(1+\beta)(d+3)}})$ 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 is similar to results presented in Theorem 7, and is thus deferred to Appendix F.

[5] In practice, one can always approximate such oracle by running stochastic gradient descent.

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

[7] We provide more mathematical backgrounds and associated definitions in Appendix F.

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, Andrei Broder, and Tong Zhang. Margin based active learning. In *International Conference on Computational Learning Theory*, pages 35–50. Springer, 2007.

Maria-Florina Balcan, Alina Beygelzimer, and John Langford. Agnostic active learning. *Journal of Computer and System Sciences*, 75(1):78–89, 2009.

Peter L Bartlett, 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.

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.

George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.

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 vapnik-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.

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, Zuowei 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, 2021a.

Rahul Parhi and Robert D Nowak. Near-minimax optimal estimation with shallow relu neural networks. *arXiv preprint arXiv:2109.08844*, 2021b.

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, Zhihui 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, Pranjal 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.

Chicheng Zhang and Kamalika Chaudhuri. Beyond disagreement-based agnostic active learning. *Advances in Neural Information Processing Systems*, 27, 2014.

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] As discussed in Section 2.3, our active learning label complexity only matches the lower bound up to the disagreement coefficient. While we have provided examples (around Theorem 6 and in Appendix C.2) where the disagreement coefficient is well-bounded, we leave a more comprehensive analyses as 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]