
\mathcal{H} -Consistency Bounds for Surrogate Loss Minimizers

Pranjal Awasthi¹ Anqi Mao² Mehryar Mohri^{1,2} Yutao Zhong²

Abstract

We present a detailed study of estimation errors in terms of surrogate loss estimation errors. We refer to such guarantees as \mathcal{H} -consistency bounds, since they account for the hypothesis set \mathcal{H} adopted. These guarantees are significantly stronger than \mathcal{H} -calibration or \mathcal{H} -consistency. They are also more informative than similar *excess error bounds* derived in the literature, when \mathcal{H} is the family of all measurable functions. We prove general theorems providing such guarantees, for both the distribution-dependent and distribution-independent settings. We show that our bounds are tight, modulo a convexity assumption. We also show that previous excess error bounds can be recovered as special cases of our general results. We then present a series of explicit bounds in the case of the zero-one loss, with multiple choices of the surrogate loss and for both the family of linear functions and neural networks with one hidden-layer. We further prove more favorable distribution-dependent guarantees in that case. We also present a series of explicit bounds in the case of the adversarial loss, with surrogate losses based on the supremum of the ρ -margin, hinge or sigmoid loss and for the same two general hypothesis sets. Here too, we prove several enhancements of these guarantees under natural distributional assumptions. Finally, we report the results of simulations illustrating our bounds and their tightness.

1. Introduction

Most learning algorithms rely on optimizing a surrogate loss function distinct from the *target loss function* tailored to the task considered. This is typically because the target loss

¹Google Research, New York, NY; ²Courant Institute of Mathematical Sciences, New York, NY. Correspondence to: Anqi Mao <aqmao@cims.nyu.edu>, Yutao Zhong <yutao@cims.nyu.edu>.

function is computationally hard to optimize or because it does not admit favorable properties, such as differentiability or smoothness, crucial to the convergence of optimization algorithms. But, what guarantees can we count on for the target loss estimation error, when minimizing a surrogate loss estimation error?

A desirable property of a surrogate loss function, often referred to in that context is *Bayes-consistency*. It requires that asymptotically, nearly optimal minimizers of the surrogate excess error also nearly optimally minimize the target excess error (Steinwart, 2007). This property holds for a broad family of convex surrogate losses of the standard binary and multi-class classification losses (Zhang, 2004a; Bartlett et al., 2006; Tewari & Bartlett, 2007; Steinwart, 2007). But, Bayes-consistency is not relevant when learning with a hypothesis set \mathcal{H} distinct from the family of all measurable functions. Instead, the hypothesis set-dependent notion of \mathcal{H} -consistency should be adopted, as argued by Long & Servedio (2013) (see also (Kuznetsov et al., 2014) and (Zhang & Agarwal, 2020)). More recently, Awasthi et al. (2021a) further studied \mathcal{H} -consistency guarantees for the *adversarial loss* (Goodfellow et al., 2014; Madry et al., 2017; Tsipras et al., 2018; Carlini & Wagner, 2017). Nevertheless, consistency and \mathcal{H} -consistency are both asymptotic properties and thus do not provide any guarantee for approximate minimizers learned from finite samples.

Instead, we will consider upper bounds on the target estimation error expressed in terms of the surrogate estimation error, which we refer to as \mathcal{H} -consistency bounds, since they account for the hypothesis set \mathcal{H} adopted. These guarantees are significantly stronger than \mathcal{H} -calibration or \mathcal{H} -consistency (Section 6) or some margin-based properties of convex surrogate losses for linear predictors studied by Ben-David et al. (2012) and Long & Servedio (2011). They are also more informative than similar *excess error bounds* derived in the literature, which correspond to the special case where \mathcal{H} is the family of all measurable functions (Zhang, 2004a; Bartlett et al., 2006) (see also (Mohri et al., 2018)[section 4.7]). We prove general theorems providing such guarantees, which could be used in both distribution-dependent and distribution-independent settings (Section 4). We show that our bounds are tight, modulo a convexity assumption (Section 5.2 and 6.1). We also show that previous excess error bounds can be recovered as special cases of our

general results (Section 5.1).

We then present a series of explicit bounds in the case of the 0/1 loss (Section 5), with multiple choices of the surrogate loss and for both the family of linear functions (Section 5.3) and that of neural networks with one hidden-layer (Section 5.4). We further prove more favorable distribution-dependent guarantees in that case (Section 5.5).

We also present a detailed analysis of the *adversarial loss* (Section 6). We show that there can be no non-trivial adversarial \mathcal{H} -consistency bound for supremum-based convex loss functions and supremum-based sigmoid loss function, under mild assumptions that hold for most hypothesis sets used in practice (Section 6.2). These results imply that the loss functions commonly used in practice for optimizing the adversarial loss cannot benefit from any useful \mathcal{H} -consistency bound guarantee! These are novel results that go beyond the negative ones given for convex surrogates by Awasthi et al. (2021a).

We present new \mathcal{H} -consistency bounds for the adversarial loss with surrogate losses based on the supremum of the ρ -margin loss, for linear hypothesis sets (Section 6.3) and the family of neural networks with one hidden-layer (Section 6.4). Here too, we prove several enhancements of these guarantees under some natural distributional assumptions (Section 6.5).

Our results help compare different surrogate loss functions of the zero-one loss or adversarial loss, given the specific hypothesis set used, based on the functional form of their \mathcal{H} -consistency bounds. These results, combined with approximation error properties of surrogate losses, can help select the most suitable surrogate loss in practice. In addition to several general theorems, our study required a careful inspection of the properties of various surrogate loss functions and hypothesis sets. Our proofs and techniques could be adopted for the analysis of many other surrogate loss functions and hypothesis sets.

In Section 7, we report the results of simulations illustrating our bounds and their tightness. We give a detailed discussion of related work in Appendix A. We start with some preliminary definitions and notation.

2. Preliminaries

Let \mathcal{X} denote the input space and $\mathcal{Y} = \{-1, +1\}$ the binary label space. We will denote by \mathcal{D} a distribution over $\mathcal{X} \times \mathcal{Y}$, by \mathcal{P} a set of such distributions and by \mathcal{H} a hypothesis set of functions mapping from \mathcal{X} to \mathbb{R} . The *generalization error* and *minimal generalization error* for a loss function $\ell(h, x, y)$ are defined as $\mathcal{R}_\ell(h) = \mathbb{E}_{(x,y) \sim \mathcal{D}}[\ell(h, x, y)]$ and $\mathcal{R}_{\ell, \mathcal{H}}^* = \inf_{h \in \mathcal{H}} \mathcal{R}_\ell(h)$. Let \mathcal{H}_{all} denote the hypothesis set of all measurable functions. The *excess error* of a hypothesis

h is defined as the difference $\mathcal{R}_\ell(h) - \mathcal{R}_{\ell, \mathcal{H}_{\text{all}}}^*$, which can be decomposed into the sum of two terms, the *estimation error* and *approximation error*:

$$\mathcal{R}_\ell(h) - \mathcal{R}_{\ell, \mathcal{H}_{\text{all}}}^* = (\mathcal{R}_\ell(h) - \mathcal{R}_{\ell, \mathcal{H}}^*) + (\mathcal{R}_{\ell, \mathcal{H}}^* - \mathcal{R}_{\ell, \mathcal{H}_{\text{all}}}^*). \quad (1)$$

Given two loss functions ℓ_1 and ℓ_2 , a fundamental question is whether ℓ_1 is *consistent* with respect to ℓ_2 for a hypothesis set \mathcal{H} and a set of distributions \mathcal{P} (Bartlett et al., 2006; Steinwart, 2007; Long & Servedio, 2013; Bao et al., 2021; Awasthi et al., 2021a).

Definition 1 ((\mathcal{P}, \mathcal{H})-consistency). *We say that ℓ_1 is (\mathcal{P}, \mathcal{H})-consistent with respect to ℓ_2 , if, for all distributions $\mathcal{D} \in \mathcal{P}$ and sequences $\{h_n\}_{n \in \mathbb{N}} \subset \mathcal{H}$, we have*

$$\lim_{n \rightarrow +\infty} \mathcal{R}_{\ell_1}(h_n) - \mathcal{R}_{\ell_1, \mathcal{H}}^* = 0 \Rightarrow \lim_{n \rightarrow +\infty} \mathcal{R}_{\ell_2}(h_n) - \mathcal{R}_{\ell_2, \mathcal{H}}^* = 0. \quad (2)$$

We will denote by Φ a margin-based loss if a loss function ℓ can be represented as $\ell(h, x, y) = \Phi(yh(x))$ and by $\tilde{\Phi} := \sup_{x': \|x-x'\|_p \leq \gamma} \Phi(yh(x'))$, $p \in [1, +\infty]$, the supremum-based counterpart. In the standard binary classification, ℓ_2 is the 0/1 loss $\ell_{0-1} := \mathbb{1}_{\text{sign}(h(x)) \neq y}$, where $\text{sign}(\alpha) = \mathbb{1}_{\alpha \geq 0} - \mathbb{1}_{\alpha < 0}$ and ℓ_1 is the margin-based loss for some function $\Phi: \mathbb{R} \rightarrow \mathbb{R}_+$, typically convex. In the adversarial binary classification, ℓ_2 is the adversarial 0/1 loss $\ell_\gamma := \sup_{x': \|x-x'\|_p \leq \gamma} \mathbb{1}_{yh(x') \leq 0}$, for some $\gamma \in (0, 1)$ and ℓ_1 is the supremum-based margin loss $\tilde{\Phi}$.

Let $B_p^d(r)$ denote the d -dimensional ℓ_p -ball with radius r : $B_p^d(r) = \{z \in \mathbb{R}^d \mid \|z\|_p \leq r\}$. Without loss of generality, we consider $\mathcal{X} = B_p^d(1)$. Let $p, q \in [1, +\infty]$ be conjugate numbers, that is $\frac{1}{p} + \frac{1}{q} = 1$. We will specifically study the family of linear hypotheses $\mathcal{H}_{\text{lin}} = \{x \mapsto w \cdot x + b \mid \|w\|_q \leq W, |b| \leq B\}$ and one-hidden-layer ReLU networks $\mathcal{H}_{\text{NN}} = \{x \mapsto \sum_{j=1}^n u_j (w_j \cdot x + b)_+ \mid \|u\|_1 \leq \Lambda, \|w_j\|_q \leq W, |b| \leq B\}$, where $(\cdot)_+ = \max(\cdot, 0)$. Finally, for any $\epsilon > 0$, we will denote by $\langle t \rangle_\epsilon$ the ϵ -truncation of $t \in \mathbb{R}$ defined by $t \mathbb{1}_{t > \epsilon}$.

3. \mathcal{H} -consistency bound definitions

(\mathcal{P}, \mathcal{H})-Consistency is an asymptotic relation between two loss functions. However, we are interested in a more quantitative relation in many applications. This motivates the study of *\mathcal{H} -consistency bound*.

Definition 2 (\mathcal{H} -consistency bound). *If for some non-decreasing function $f: \mathbb{R}_+ \rightarrow \mathbb{R}_+$, a bound of the following form holds for all $h \in \mathcal{H}$ and $\mathcal{D} \in \mathcal{P}$:*

$$\mathcal{R}_{\ell_2}(h) - \mathcal{R}_{\ell_2, \mathcal{H}}^* \leq f(\mathcal{R}_{\ell_1}(h) - \mathcal{R}_{\ell_1, \mathcal{H}}^*), \quad (3)$$

then, we call it an \mathcal{H} -consistency bound. Furthermore, if \mathcal{P} consists of all distributions over $\mathcal{X} \times \mathcal{Y}$, we say that the bound is distribution-independent.

When $\mathcal{H} = \mathcal{H}_{\text{all}}$ and \mathcal{P} is the set of all distributions, a bound of the form (3) is also called a *consistency excess error bound*. Note when $f(0) = 0$ and f is continuous at 0, the \mathcal{H} -consistency bound (3) implies \mathcal{H} -consistency (2). Thus, \mathcal{H} -consistency bounds provide stronger quantitative results than consistency and calibration. Furthermore, there is a fundamental reason to study such bounds from the statistical learning point of view: they can be turned into more favorable generalization bounds for the target loss ℓ_2 than the excess error bound. For example, when \mathcal{P} is the set of all distributions, by (1), relation (3) implies that, for all $h \in \mathcal{H}$, the following inequality holds:

$$\mathcal{R}_{\ell_2}(h) - \mathcal{R}_{\ell_2, \mathcal{H}_{\text{all}}}^* \leq f(\mathcal{R}_{\ell_1}(h) - \mathcal{R}_{\ell_1, \mathcal{H}}^*) + \mathcal{R}_{\ell_2, \mathcal{H}}^* - \mathcal{R}_{\ell_2, \mathcal{H}_{\text{all}}}^*. \quad (4)$$

Similarly, the excess error bound can be written as follows:

$$\mathcal{R}_{\ell_2}(h) - \mathcal{R}_{\ell_2, \mathcal{H}_{\text{all}}}^* \leq f(\mathcal{R}_{\ell_1}(h) - \mathcal{R}_{\ell_1, \mathcal{H}}^* + \mathcal{R}_{\ell_1, \mathcal{H}}^* - \mathcal{R}_{\ell_1, \mathcal{H}_{\text{all}}}^*). \quad (5)$$

If we further bound the estimation error $[\mathcal{R}_{\ell_1}(h) - \mathcal{R}_{\ell_1, \mathcal{H}}^*]$ by the empirical error plus a complexity term, (4) and (5) both turn into generalization bounds. However, the generalization bound obtained by (4) is linearly dependent on the approximation error of target loss ℓ_2 , while the one obtained by (5) depends on the approximation error of the surrogate loss ℓ_1 and can potentially be worse than linear dependence. Moreover, (4) can be easily used to compare different surrogates by directly comparing the corresponding mapping f . However, only comparing the mapping f for different surrogates in (5) is not sufficient since the approximation errors of surrogates may differ as well.

Minimizability gap. We will adopt the standard notation for the conditional distribution of Y given $X = x$: $\eta(x) = \mathcal{D}(Y = 1 \mid X = x)$ and will also use the shorthand $\Delta\eta(x) = \eta(x) - \frac{1}{2}$. It is useful to write the generalization error as $\mathcal{R}_\ell(h) = \mathbb{E}_X[\mathcal{C}_\ell(h, x)]$, where $\mathcal{C}_\ell(h, x)$ is the *conditional ℓ -risk* defined by $\mathcal{C}_\ell(h, x) = \eta(x)\ell(h, x, +1) + (1 - \eta(x))\ell(h, x, -1)$. The *minimal conditional ℓ -risk* is denoted by $\mathcal{C}_{\ell, \mathcal{H}}^*(x) = \inf_{h \in \mathcal{H}} \mathcal{C}_\ell(h, x)$. We also use the following shorthand for the gap $\Delta\mathcal{C}_{\ell, \mathcal{H}}(h, x) = \mathcal{C}_\ell(h, x) - \mathcal{C}_{\ell, \mathcal{H}}^*(x)$. We call $\langle \Delta\mathcal{C}_{\ell, \mathcal{H}}(h, x) \rangle_\epsilon = \Delta\mathcal{C}_{\ell, \mathcal{H}}(h, x) \mathbb{1}_{\Delta\mathcal{C}_{\ell, \mathcal{H}}(h, x) > \epsilon}$ the *conditional ϵ -regret* for ℓ . To simplify the notation, we also define for any $t \in [0, 1]$, $\mathcal{C}_\ell(h, x, t) = t\ell(h, x, +1) + (1 - t)\ell(h, x, -1)$ and $\Delta\mathcal{C}_{\ell, \mathcal{H}}(h, x, t) = \mathcal{C}_\ell(h, x, t) - \inf_{h \in \mathcal{H}} \mathcal{C}_\ell(h, x, t)$. Thus, $\Delta\mathcal{C}_{\ell, \mathcal{H}}(h, x, \eta(x)) = \Delta\mathcal{C}_{\ell, \mathcal{H}}(h, x)$.

A key quantity that appears in our bounds is the (ℓ, \mathcal{H}) -*minimizability gap* $\mathcal{M}_{\ell, \mathcal{H}}$, which is the difference of the best-in-class error and the expectation of the minimal conditional ℓ -risk:

$$\mathcal{M}_{\ell, \mathcal{H}} = \mathcal{R}_{\ell, \mathcal{H}}^* - \mathbb{E}_X[\mathcal{C}_{\ell, \mathcal{H}}^*(x)].$$

This is an inherent property of the hypothesis set \mathcal{H} and distribution \mathcal{D} that we cannot hope to estimate or minimize.

As an example, the minimizability gap for the 0/1 loss and adversarial 0/1 loss with \mathcal{H}_{all} can be expressed as follows:

$$\begin{aligned} \mathcal{M}_{\ell_{0-1}, \mathcal{H}_{\text{all}}} &= \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^* - \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}] = 0, \\ \mathcal{M}_{\ell_\gamma, \mathcal{H}_{\text{all}}} &= \mathcal{R}_{\ell_\gamma, \mathcal{H}_{\text{all}}}^* - \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}]. \end{aligned}$$

Steinwart (2007, Lemma 2.5) shows that the minimizability gap vanishes when the loss ℓ is *minimizable*. Awasthi et al. (2021a) point out that the minimizability condition does not hold for adversarial loss functions, and therefore that, in general, $\mathcal{M}_{\ell_\gamma, \mathcal{H}_{\text{all}}}$ is strictly positive, thereby presenting additional challenges for adversarial robust classification. Thus, the minimizability gap is critical in the study of adversarial surrogate loss functions. The minimizability gaps for some common loss functions and hypothesis sets are given in Table 1 in Section 5.2 for completeness.

4. General theorems

We first introduce two main theorems that provide a general \mathcal{H} -consistency bound between any target loss and surrogate loss. These bounds are \mathcal{H} -dependent, taking into consideration the specific hypothesis set used by a learning algorithm. To the best of our knowledge, no such guarantee has appeared in the past. For both theoretical and practical computational reasons, learning algorithms typically seek a good hypothesis within a restricted subset of \mathcal{H}_{all} . Thus, in general, \mathcal{H} -dependent bounds can provide more relevant guarantees than excess error bounds. Our proposed bounds are also more general in the sense that \mathcal{H}_{all} can be used as a special case. Theorems 1 and 2 are counterparts of each other, while the latter may provide a more explicit form of bounds as in (3).

Theorem 1 (Distribution-dependent Ψ -bound). *Assume that there exists a convex function $\Psi: \mathbb{R}_+ \rightarrow \mathbb{R}$ with $\Psi(0) \geq 0$ and $\epsilon \geq 0$ such that the following holds for all $h \in \mathcal{H}$ and $x \in \mathcal{X}$:*

$$\Psi(\langle \Delta\mathcal{C}_{\ell_2, \mathcal{H}}(h, x) \rangle_\epsilon) \leq \Delta\mathcal{C}_{\ell_1, \mathcal{H}}(h, x). \quad (6)$$

Then, the following inequality holds for any $h \in \mathcal{H}$:

$$\begin{aligned} &\Psi(\mathcal{R}_{\ell_2}(h) - \mathcal{R}_{\ell_2, \mathcal{H}}^* + \mathcal{M}_{\ell_2, \mathcal{H}}) \\ &\leq \mathcal{R}_{\ell_1}(h) - \mathcal{R}_{\ell_1, \mathcal{H}}^* + \mathcal{M}_{\ell_1, \mathcal{H}} + \max\{\Psi(0), \Psi(\epsilon)\}. \end{aligned} \quad (7)$$

Theorem 2 (Distribution-dependent Γ -bound). *Assume that there exists a concave function $\Gamma: \mathbb{R}_+ \rightarrow \mathbb{R}$ and $\epsilon \geq 0$ such that the following holds for all $h \in \mathcal{H}$ and $x \in \mathcal{X}$:*

$$\langle \Delta\mathcal{C}_{\ell_2, \mathcal{H}}(h, x) \rangle_\epsilon \leq \Gamma(\Delta\mathcal{C}_{\ell_1, \mathcal{H}}(h, x)). \quad (8)$$

Then, the following inequality holds for any $h \in \mathcal{H}$:

$$\mathcal{R}_{\ell_2}(h) - \mathcal{R}_{\ell_2, \mathcal{H}}^* \leq \Gamma(\mathcal{R}_{\ell_1}(h) - \mathcal{R}_{\ell_1, \mathcal{H}}^* + \mathcal{M}_{\ell_1, \mathcal{H}}) - \mathcal{M}_{\ell_2, \mathcal{H}} + \epsilon. \quad (9)$$

The proofs of Theorems 1 and 2 are included in Appendix D, where we make use of the convexity of Ψ and concavity of Γ . Below, we will mainly focus on the case where $\Psi(0) = 0$ and $\epsilon = 0$. Note that if ℓ_2 is upper bounded by ℓ_1 and $\mathcal{R}_{\ell_1, \mathcal{H}}^* - \mathcal{M}_{\ell_1, \mathcal{H}} = \mathcal{R}_{\ell_2, \mathcal{H}}^* - \mathcal{M}_{\ell_2, \mathcal{H}}$, then, the following inequality automatically holds for any $h \in \mathcal{H}$:

$$\mathcal{R}_{\ell_2}(h) - \mathcal{R}_{\ell_2, \mathcal{H}}^* + \mathcal{M}_{\ell_2, \mathcal{H}} \leq \mathcal{R}_{\ell_1}(h) - \mathcal{R}_{\ell_1, \mathcal{H}}^* + \mathcal{M}_{\ell_1, \mathcal{H}}.$$

This is a special case of Theorems 1 and 2. Indeed, since $\mathcal{R}_{\ell_1, \mathcal{H}}^* - \mathcal{M}_{\ell_1, \mathcal{H}} = \mathcal{R}_{\ell_2, \mathcal{H}}^* - \mathcal{M}_{\ell_2, \mathcal{H}}$, we have $\mathcal{C}_{\ell_2, \mathcal{H}}^*(x) \equiv \mathcal{C}_{\ell_1, \mathcal{H}}^*(x)$ and thus $\Delta \mathcal{C}_{\ell_2, \mathcal{H}}(h, x) \leq \Delta \mathcal{C}_{\ell_1, \mathcal{H}}(h, x)$. Therefore, Φ and Γ can be the identity function. We refer to such cases as “trivial cases”. They occur when $\mathcal{M}_{\ell_1, \mathcal{H}}$ and $\mathcal{M}_{\ell_2, \mathcal{H}}$ respectively coincide with the corresponding approximation errors and $\mathcal{R}_{\ell_1, \mathcal{H}_{\text{all}}}^* = \mathcal{R}_{\ell_2, \mathcal{H}_{\text{all}}}^*$. We will later see such cases for specific loss functions and hypothesis sets (See (38) in Appendix K.1.6 and (56) in Appendix L.1.1). Let us point out, however, that the corresponding \mathcal{H} -consistency bounds are still valid and worth studying since they can be shown to be the tightest (Theorems 4 and 6).

Theorem 1 is distribution-dependent, in the sense that, for a fixed distribution, if we find a Ψ that satisfies condition (6), then the bound (7) only gives guarantee for that same distribution. Since the distribution \mathcal{D} of interest is typically unknown, to obtain guarantees for \mathcal{D} , if the only information given is that \mathcal{D} belongs to a set of distributions \mathcal{P} , we need to find a Ψ that satisfies condition (6) for all the distributions in \mathcal{P} . The choice of Ψ is critical, since it determines the form of the bound obtained.

We say that Ψ is *optimal* if any function that makes the bound (7) hold for all distributions in \mathcal{P} is everywhere no larger than Ψ . The optimal Ψ leads to the tightest \mathcal{H} -consistency bound (7) uniform over \mathcal{P} . Specifically, when \mathcal{P} consists of all distributions, we say that the bound is distribution-independent. The above also applies to Theorem 2, except that Γ is *optimal* if any function that makes the bound (9) hold for all distributions in \mathcal{P} is everywhere no less than Γ .

When ℓ_2 is the 0/1 loss or the adversarial 0/1 loss, the conditional ϵ -regret that appears in condition (6) has explicit forms for common hypothesis sets as characterized later in Lemma 1 and 2, establishing the basis for introducing non-adversarial and adversarial \mathcal{H} -estimation error transformation in Section 5.2 and 6.1. We will see later in these sections that the transformations introduced are often the optimal Ψ we are seeking for, which respectively leads to tight non-adversarial and adversarial distribution-independent guarantees. In Section 5 and 6, we also apply our general theorems and tools to loss functions and hypothesis sets widely used in practice. Each case requires a careful analysis that we present in detail.

5. Guarantees for the zero-one loss $\ell_2 = \ell_{0-1}$

In this section, we discuss guarantees in the non-adversarial scenario where ℓ_2 is the zero-one loss, ℓ_{0-1} . The lemma stated next characterizes the minimal conditional ℓ_{0-1} -risk and the conditional ϵ -regret, which will be helpful for introducing the general tools in Section 5.2. The proof is given in Appendix E. For convenience, we will adopt the following notation: $\overline{\mathcal{H}}(x) = \{h \in \mathcal{H} : \text{sign}(h(x))\Delta\eta(x) \leq 0\}$.

Lemma 1. *Assume that \mathcal{H} satisfies the following condition for any $x \in \mathcal{X}$: $\{\text{sign}(h(x)) : h \in \mathcal{H}\} = \{-1, +1\}$. Then, the minimal conditional ℓ_{0-1} -risk is*

$$\mathcal{C}_{\ell_{0-1}, \mathcal{H}}^*(x) = \mathcal{C}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^*(x) = \min\{\eta(x), 1 - \eta(x)\}.$$

The conditional ϵ -regret for ℓ_{0-1} can be characterized as

$$\langle \Delta \mathcal{C}_{\ell_{0-1}, \mathcal{H}}(h, x) \rangle_{\epsilon} = \langle 2|\Delta\eta(x)| \rangle_{\epsilon} \mathbb{1}_{h \in \overline{\mathcal{H}}(x)}.$$

5.1. Hypothesis set of all measurable functions

Before introducing our general tools, we will consider the case where $\mathcal{H} = \mathcal{H}_{\text{all}}$ and will show that previous excess error bounds can be recovered as special cases of our results. As shown in (Steinwart, 2007), both $\mathcal{M}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}$ and $\mathcal{M}_{\Phi, \mathcal{H}_{\text{all}}}$ vanish. Thus by Lemma 1, we obtain the following corollary of Theorem 1 by taking $\epsilon = 0$.

Corollary 1. *Assume that there exists a convex function $\Psi: \mathbb{R}_+ \rightarrow \mathbb{R}$ with $\Psi(0) = 0$ such that for any $x \in \mathcal{X}$, $\Psi(2|\Delta\eta(x)|) \leq \inf_{h \in \overline{\mathcal{H}}_{\text{all}}(x)} \Delta \mathcal{C}_{\Phi, \mathcal{H}_{\text{all}}}(h, x)$. Then, for any hypothesis $h \in \mathcal{H}_{\text{all}}$, the following inequality holds:*

$$\Psi(\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^*) \leq \mathcal{R}_{\Phi}(h) - \mathcal{R}_{\Phi, \mathcal{H}_{\text{all}}}^*.$$

Furthermore, Corollary 2 follows from Corollary 1 by taking the convex function $\Psi(t) = (t/(2c))^s$.

Corollary 2. *Assume there exist $s \geq 1$ and $c > 0$ such that for any $x \in \mathcal{X}$, $|\Delta\eta(x)| \leq c \inf_{h \in \overline{\mathcal{H}}_{\text{all}}(x)} (\Delta \mathcal{C}_{\Phi, \mathcal{H}_{\text{all}}}(h, x))^{\frac{1}{s}}$. Then, for any hypothesis $h \in \mathcal{H}_{\text{all}}$,*

$$\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^* \leq 2c \left(\mathcal{R}_{\Phi}(h) - \mathcal{R}_{\Phi, \mathcal{H}_{\text{all}}}^* \right)^{\frac{1}{s}}.$$

The excess error bound results in the literature are all covered by the above corollaries. As shown in Appendix F, Theorem 4.7 in (Mohri et al., 2018) is a special case of Corollary 2 and Theorem 1.1 in (Bartlett et al., 2006) is a special case of Corollary 1.

5.2. General hypothesis sets \mathcal{H}

In this section, we provide general tools to study \mathcal{H} -consistency bounds when the target loss is the 0/1 loss. We will then apply them to study specific hypothesis sets and surrogates in Section 5.3 and 5.4. Lemma 1 characterizes




Figure 1: Left: surrogates. Right: \mathcal{H}_{lin} -est. error trans. inv.

the conditional ϵ -regret for ℓ_{0-1} with common hypothesis sets. Thus, Theorems 1 and 2 can be instantiated as Theorems 8 and 9 in these cases (see Appendix C). They are powerful distribution-dependent bounds and, as discussed in Section 4, the bounds become distribution-independent if the corresponding conditions can be verified for all the distributions with some Ψ , which is equivalent to verifying the condition in the following theorem.

Theorem 3 (Distribution-independent Ψ -bound). *Assume that \mathcal{H} satisfies the condition of Lemma 1. Assume that there exists a convex function $\Psi: \mathbb{R}_+ \rightarrow \mathbb{R}$ with $\Psi(0) = 0$ and $\epsilon \geq 0$ such that for any $t \in [1/2, 1]$,*

$$\Psi(\langle 2t - 1 \rangle_\epsilon) \leq \inf_{x \in \mathcal{X}, h \in \mathcal{H}: h(x) < 0} \Delta \mathcal{C}_{\Phi, \mathcal{H}}(h, x, t).$$

Then, for any hypothesis $h \in \mathcal{H}$ and any distribution,

$$\begin{aligned} & \Psi(\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}}^* + \mathcal{M}_{\ell_{0-1}, \mathcal{H}}) \\ & \leq \mathcal{R}_{\Phi}(h) - \mathcal{R}_{\Phi, \mathcal{H}}^* + \mathcal{M}_{\Phi, \mathcal{H}} + \max\{0, \Psi(\epsilon)\}. \end{aligned} \quad (10)$$

The counterpart of Theorem 3 is Theorem 12 (distribution-independent Γ -bound), deferred to Appendix C due to space limitations. The proofs for both theorems are included in Appendix G. Theorem 3 provides the general tool to derive distribution-independent \mathcal{H} -consistency bounds. They are in fact tight if we choose Ψ to be the \mathcal{H} -estimation error transformation defined as follows.

Definition 3 (\mathcal{H} -estimation error transformation). *The \mathcal{H} -estimation error transformation of Φ is defined on $t \in [0, 1]$ by $\mathcal{T}_{\Phi}(t) = \mathcal{T}(t)\mathbb{1}_{t \in [\epsilon, 1]} + (\mathcal{T}(\epsilon)/\epsilon)t\mathbb{1}_{t \in [0, \epsilon]}$, where $\mathcal{T}(t) := \inf_{x \in \mathcal{X}, h \in \mathcal{H}: h(x) < 0} \Delta \mathcal{C}_{\Phi, \mathcal{H}}(h, x, \frac{t+1}{2})$.*

When $\epsilon = 0$, $\mathcal{T}_{\Phi}(t)$ coincides with $\mathcal{T}(t)$. Observe that for any $t \in [(1 + \epsilon)/2, 1]$, the following equality holds:

$$\mathcal{T}_{\Phi}(2t - 1) = \inf_{x \in \mathcal{X}, h \in \mathcal{H}: h(x) < 0} \Delta \mathcal{C}_{\Phi, \mathcal{H}}(h, x, t).$$

Taking $\Psi = \mathcal{T}_{\Phi}$ satisfies the condition in Theorem 3 if \mathcal{T}_{Φ} is convex with $\mathcal{T}_{\Phi}(0) = 0$. Moreover, as mentioned earlier, it actually leads to the tightest \mathcal{H} -consistency bound (10) when $\epsilon = 0$.

Theorem 4 (Tightness). *Suppose that \mathcal{H} satisfies the condition of Lemma 1 and that $\epsilon = 0$. If \mathcal{T}_{Φ} is convex with $\mathcal{T}_{\Phi}(0) = 0$, then, for any $t \in [0, 1]$ and $\delta > 0$, there exist a distribution \mathcal{D} and a hypothesis $h \in \mathcal{H}$ such that*

Table 1: Loss functions and their minimizability gaps.

Loss Functions	Definitions	$\mathcal{M}_{\ell, \mathcal{H}_{\text{lin}}}$	$\mathcal{M}_{\ell, \mathcal{H}_{\text{NN}}}$
Hinge	$\Phi_{\text{hinge}}(t) = \max\{0, 1 - t\}$	(25)	(40)
Logistic	$\Phi_{\text{log}}(t) = \log_2(1 + e^{-t})$	(27)	(42)
Exponential	$\Phi_{\text{exp}}(t) = e^{-t}$	(29)	(44)
Quadratic	$\Phi_{\text{quad}}(t) = (1 - t)^2 \mathbb{1}_{t \leq 1}$	(31)	(31)
Sigmoid	$\Phi_{\text{sig}}(t) = 1 - \tanh(kt), k > 0$	(33)	(48)
ρ -Margin	$\Phi_{\rho}(t) = \min\{1, \max\{0, 1 - \frac{t}{\rho}\}\}, \rho > 0$	(36)	(40)
Sup- ρ -Margin	$\bar{\Phi}_{\rho} = \sup_{x': \ x - x'\ _{\rho} \leq \gamma} \Phi_{\rho}(yh(x'))$	(53)	(58)
Zero-One	$\ell_{0-1} = \mathbb{1}_{\text{sign}(h(x)) \neq y}$	(24)	(39)
Adversarial Zero-One	$\ell_{\gamma} = \sup_{x': \ x - x'\ _{\rho} \leq \gamma} \mathbb{1}_{yh(x') \leq 0}$	(52)	(57)

Table 2: \mathcal{H}_{lin} -estimation error transformation and \mathcal{H}_{lin} -consistency bounds with $\epsilon = 0$.

Surrogates	$\mathcal{T}_{\Phi}(t), t \in [0, 1]$	Bound
Hinge	$\min\{B, 1\}t$	(26)
Logistic	$\begin{cases} \frac{t+1}{2} \log_2(t+1) + \frac{1-t}{2} \log_2(1-t), & t \leq \frac{e^B - 1}{e^B + 1}, \\ 1 - \frac{t+1}{2} \log_2(1 + e^{-B}) - \frac{1-t}{2} \log_2(1 + e^B), & t > \frac{e^B - 1}{e^B + 1}. \end{cases}$	(28)
Exponential	$\begin{cases} 1 - \sqrt{1 - t^2}, & t \leq \frac{e^{2B} - 1}{e^{2B} + 1}, \\ 1 - \frac{t+1}{2} e^{-B} - \frac{1-t}{2} e^B, & t > \frac{e^{2B} - 1}{e^{2B} + 1}. \end{cases}$	(30)
Quadratic	$\begin{cases} t^2, & t \leq B, \\ 2Bt - B^2, & t > B. \end{cases}$	(32)
Sigmoid	$\tanh(kB)t$	(34)
ρ -Margin	$\frac{\min\{B, \rho\}}{\rho}t$	(37)

$$\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}}^* + \mathcal{M}_{\ell_{0-1}, \mathcal{H}} = t \text{ and } \mathcal{T}_{\Phi}(t) \leq \mathcal{R}_{\Phi}(h) - \mathcal{R}_{\Phi, \mathcal{H}}^* + \mathcal{M}_{\Phi, \mathcal{H}} \leq \mathcal{T}_{\Phi}(t) + \delta.$$

The proof is included in Appendix I. In other words, when $\epsilon = 0$, if \mathcal{T}_{Φ} is convex with $\mathcal{T}_{\Phi}(0) = 0$, it is optimal for the distribution-independent bound (10). Moreover, if \mathcal{T}_{Φ} is additionally invertible and non-increasing, \mathcal{T}_{Φ}^{-1} is the optimal function for the distribution-independent bound in Theorem 12 (Appendix C) and the two bounds are equivalent.

In the following sections, we will see that all these assumptions hold for common loss functions with linear and neural network hypothesis sets. Next, we will apply Theorems 3 and 4 to the linear models (Section 5.3) and neural networks (Section 5.4). Each case requires a detailed analysis (See Appendix K.1 and K.2).

The loss functions considered below and their minimizability gaps are defined in Table 1. In some cases, the minimizability gap coincides with the approximation error. For example, $\mathcal{M}_{\Phi_{\text{sig}}, \mathcal{H}_{\text{lin}}} = \mathcal{R}_{\Phi_{\text{sig}}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X[1 - |1 - 2\eta(x)| \tanh(k(W\|x\|_p + B))]$ coincides with the $(\Phi_{\text{sig}}, \mathcal{H}_{\text{lin}})$ -approximation error $\mathcal{R}_{\Phi_{\text{sig}}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X[1 - |1 - 2\eta(x)|]$ for $B = +\infty$; $\mathcal{M}_{\Phi_{\text{hinge}}, \mathcal{H}_{\text{NN}}} = \mathcal{R}_{\Phi_{\text{hinge}}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X[1 - |2\eta(x) - 1| \min\{\Lambda W\|x\|_p + \Lambda B, 1\}]$ coincides with the $(\Phi_{\text{hinge}}, \mathcal{H}_{\text{NN}})$ -approximation error $\mathcal{R}_{\Phi_{\text{hinge}}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X[1 - |1 - 2\eta(x)|]$ for $\Lambda B \geq 1$. The detailed derivation is included in Appendix K, L.

5.3. Linear hypotheses

By applying Theorems 3 and 4, we can derive \mathcal{H}_{lin} -consistency bounds for common loss functions defined in Table 1. Table 2 supplies the \mathcal{H}_{lin} -estimation error transformation \mathcal{T}_{Φ} and the corresponding bounds for those

Table 3: \mathcal{H}_{NN} -estimation error transformation and \mathcal{H}_{NN} -consistency bounds with $\epsilon = 0$.

Surrogates	$\mathcal{T}_\Phi(t), t \in [0, 1]$	Bound
Hinge	$\min\{\Lambda B, 1\}t$	(41)
Logistic	$\begin{cases} \frac{t+1}{2} \log_2(t+1) + \frac{1-t}{2} \log_2(1-t), & t \leq \frac{e^{\Lambda B} - 1}{e^{\Lambda B} + 1}, \\ 1 - \frac{t+1}{2} \log_2(1 + e^{-\Lambda B}) - \frac{1-t}{2} \log_2(1 + e^{\Lambda B}), & t > \frac{e^{\Lambda B} - 1}{e^{\Lambda B} + 1}. \end{cases}$	(43)
Exponential	$\begin{cases} 1 - \sqrt{1 - t^2}, & t \leq \frac{e^{2\Lambda B} - 1}{e^{2\Lambda B} + 1}, \\ 1 - \frac{t+1}{2} e^{-\Lambda B} - \frac{1-t}{2} e^{\Lambda B}, & t > \frac{e^{2\Lambda B} - 1}{e^{2\Lambda B} + 1}. \end{cases}$	(45)
Quadratic	$\begin{cases} t^2, & t \leq \Lambda B, \\ 2\Lambda B t - (\Lambda B)^2, & t > \Lambda B. \end{cases}$	(47)
Sigmoid	$\tanh(k\Lambda B)t$	(49)
ρ -Margin	$\frac{\min\{\Lambda B, \rho\}}{\rho} t$	(51)

loss functions. The inverse \mathcal{T}_Φ^{-1} is given in Table 5 of Appendix B. Surrogates Φ and their corresponding \mathcal{T}_Φ^{-1} ($B = 0.8$) are visualized in Figure 1. Theorems 3 and 4 apply to all these cases since \mathcal{T}_Φ is convex, increasing, invertible and satisfies that $\mathcal{T}_\Phi(0) = 0$. More precisely, taking $\Psi = \mathcal{T}_\Phi$ and $\epsilon = 0$ in (10) and using the inverse function \mathcal{T}_Φ^{-1} directly give the tightest bound. As an example, for the sigmoid loss, $\mathcal{T}_{\Phi_{\text{sig}}}^{-1}(t) = \frac{t}{\tanh(kB)}$. Then the bound (10) becomes $\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{lin}}}^* \leq (\mathcal{R}_{\Phi_{\text{sig}}}(h) - \mathcal{R}_{\Phi_{\text{sig}}, \mathcal{H}_{\text{lin}}}^* + \mathcal{M}_{\Phi_{\text{sig}}, \mathcal{H}_{\text{lin}}}) / \tanh(kB) - \mathcal{M}_{\ell_{0-1}, \mathcal{H}_{\text{lin}}}$, which is (34) in Table 2. Furthermore, after plugging in the minimizability gaps concluded in Table 1, we will obtain the novel bound $\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^* \leq (\mathcal{R}_{\Phi_{\text{sig}}}(h) - \mathbb{E}_X[1 - |1 - 2\eta(x)| \tanh(k(W\|x\|_p + B))]) / \tanh(kB)$ ((35) in Appendix K.1.5). The bounds for other surrogates are similarly derived in Appendix K.1. For the logistic loss and exponential loss, to simplify the expression, the bounds are obtained by plugging in an upper bound of \mathcal{T}_Φ^{-1} .

Let us emphasize that these \mathcal{H} -consistency bounds are novel in the sense that they are all hypothesis set-dependent and, to our knowledge, no such guarantee has been presented before. More precisely, the bounds of Table 2 depend directly on the parameter B in the linear models and parameters of the loss function (e.g., k in sigmoid loss). Thus, for a fixed hypothesis $h \in \mathcal{H}_{\text{lin}}$, we may give the tightest bound by choosing the best parameter B . As an example, Appendix K.1.5 shows that the bound (35) with $B = +\infty$ coincides with the excess error bound known for the sigmoid loss (Bartlett et al., 2006). However, for a fixed hypothesis h , by varying B (hypothesis set) and k (loss function), we may obtain a finer bound! Thus studying hypothesis set-dependent bounds can guide us to select the most suitable hypothesis set and loss function. Moreover, as shown by Theorem 4, all the bounds obtained by directly using \mathcal{T}_Φ^{-1} are tight and cannot be further improved.

5.4. One-hidden-layer ReLU neural networks

In this section, we give \mathcal{H} -consistency bounds for one-hidden-layer ReLU neural networks \mathcal{H}_{NN} . Table 3 is the counterpart of Table 2 for \mathcal{H}_{NN} . Different from the bounds in the linear case, all the bounds in Table 3 not only depend on B , but also depend on Λ , which is a new parameter in

\mathcal{H}_{NN} . This further illustrates that our bounds are hypothesis set-dependent and that, as with the linear case, adequately choosing the parameters Λ and B in \mathcal{H}_{NN} would give us better hypothesis set-dependent guarantees than standard excess error bounds. The inverse \mathcal{T}_Φ^{-1} is given in Table 6 of Appendix B. Our proofs and techniques could also be adopted for the analysis of multi-layer neural networks.

5.5. Guarantees under Massart’s noise condition

The distribution-independent \mathcal{H} -consistency bound (10) cannot be improved, since they are tight as shown in Theorem 4. However, the bounds can be further improved in the distribution-dependent setting. Indeed, we will study how \mathcal{H} -consistency bounds can be improved under low noise conditions, which impose the restrictions on the conditional distribution $\eta(x)$. We consider Massart’s noise condition (Massart & Nédélec, 2006) which is defined as follows.

Definition 4 (Massart’s noise). *The distribution \mathcal{D} over $\mathcal{X} \times \mathcal{Y}$ satisfies Massart’s noise condition if $|\Delta\eta(x)| \geq \beta$ for almost all $x \in \mathcal{X}$, for some constant $\beta \in (0, 1/2]$.*

When it is known that the distribution \mathcal{D} satisfies Massart’s noise condition with β , in contrast with the distribution-independent bounds, we can require the bounds (7) and (9) to hold uniformly only for such distributions. With Massart’s noise condition, we introduce a modified \mathcal{H} -estimation error transformation in Proposition 1 (Appendix M), which verifies condition (13) of Theorem 8 (the finer distribution dependent guarantee mentioned before, deferred to Appendix C) for all distributions under the noise condition. Then, using this transformation, we can obtain more favorable distribution-dependent bounds. As an example, we consider the quadratic loss Φ_{quad} , the logistic loss Φ_{log} and the exponential loss Φ_{exp} with \mathcal{H}_{all} . For all distributions and $h \in \mathcal{H}_{\text{all}}$, as shown in (Zhang, 2004a; Bartlett et al., 2006; Mohri et al., 2018), the following holds:

$$\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^* \leq \sqrt{2}(\mathcal{R}_\Phi(h) - \mathcal{R}_{\Phi, \mathcal{H}_{\text{all}}}^*)^{1/2},$$

when the surrogate loss Φ is Φ_{log} or Φ_{exp} . If $\Phi = \Phi_{\text{quad}}$, then the constant multiplier $\sqrt{2}$ can be removed. For distributions that satisfy Massart’s noise condition with β , as proven in Appendix M, for any $h \in \mathcal{H}_{\text{all}}$ such that $\mathcal{R}_\Phi(h) \leq \mathcal{R}_{\Phi, \mathcal{H}_{\text{all}}}^* + \mathcal{J}(2\beta)$, the consistency excess error bound is improved from the square-root dependency to a linear dependency:

$$\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^* \leq 2\beta(\mathcal{R}_\Phi(h) - \mathcal{R}_{\Phi, \mathcal{H}_{\text{all}}}^*) / \mathcal{J}(2\beta), \quad (11)$$

where $\mathcal{J}(t)$ equals to $t^2, \frac{t+1}{2} \log_2(t+1) + \frac{1-t}{2} \log_2(1-t)$ and $1 - \sqrt{1 - t^2}$ for $\Phi_{\text{quad}}, \Phi_{\text{log}}$ and Φ_{exp} respectively. These linear dependent bounds are tight, as illustrated in Section 7.

6. Guarantees for the adversarial loss $\ell_2 = \ell_\gamma$

In this section, we discuss the adversarial scenario where ℓ_2 is the adversarial 0/1 loss ℓ_γ . We consider *symmetric* hypothesis sets, which satisfy: $h \in \mathcal{H}$ if and only if $-h \in \mathcal{H}$. For convenience, we will adopt the following definitions:

$$\underline{h}_\gamma(x) = \inf_{x': \|x-x'\|_p \leq \gamma} h(x') \quad \bar{h}_\gamma(x) = \sup_{x': \|x-x'\|_p \leq \gamma} h(x').$$

We also define $\bar{\mathcal{H}}_\gamma(x) = \{h \in \mathcal{H} : \underline{h}_\gamma(x) \leq 0 \leq \bar{h}_\gamma(x)\}$. The following characterization of the minimal conditional ℓ_γ -risk and conditional ϵ -regret is based on (Awasthi et al., 2021a, Lemma 27) and will be helpful in introducing the general tools in Section 6.1. The proof is similar and is included in Appendix E for completeness.

Lemma 2. *Assume that \mathcal{H} is symmetric. Then, the minimal conditional ℓ_γ -risk is*

$$\mathcal{C}_{\ell_\gamma, \mathcal{H}}^*(x) = \min\{\eta(x), 1 - \eta(x)\} \mathbb{1}_{\bar{\mathcal{H}}_\gamma(x) \neq \mathcal{H}} + \mathbb{1}_{\bar{\mathcal{H}}_\gamma(x) = \mathcal{H}}.$$

The conditional ϵ -regret for ℓ_γ can be characterized as

$$\langle \Delta \mathcal{C}_{\ell_\gamma, \mathcal{H}}(h, x) \rangle_\epsilon = \begin{cases} \langle |\Delta \eta(x)| + \frac{1}{2} \rangle_\epsilon & h \in \bar{\mathcal{H}}_\gamma(x) \not\subseteq \mathcal{H} \\ \langle 2\Delta \eta(x) \rangle_\epsilon & \bar{h}_\gamma(x) < 0 \\ \langle -2\Delta \eta(x) \rangle_\epsilon & \underline{h}_\gamma(x) > 0 \\ 0 & \text{otherwise} \end{cases}$$

6.1. General hypothesis sets \mathcal{H}

As with the non-adversarial case, we begin by providing general theoretical tools to study \mathcal{H} -consistency bounds when the target loss is the adversarial 0/1 loss. Lemma 2 characterizes the conditional ϵ -regret for ℓ_γ with symmetric hypothesis sets. Thus, Theorems 1 and 2 can be instantiated as Theorems 10 and 11 (See Appendix C) in these cases. These results are distribution-dependent and can serve as general tools. For example, we can use these tools to derive more favorable guarantees under noise conditions (Section 6.5). As in the previous section, we present their distribution-independent version in the following theorem.

Theorem 5 (Adversarial distribution-independent Ψ -bound). *Suppose that \mathcal{H} is symmetric. Assume there exist a convex function $\Psi: \mathbb{R}_+ \rightarrow \mathbb{R}$ with $\Psi(0) = 0$ and $\epsilon \geq 0$ such that the following holds for any $t \in [1/2, 1]$:*

$$\begin{aligned} \Psi(\langle t \rangle_\epsilon) &\leq \inf_{x \in \mathcal{X}, h \in \bar{\mathcal{H}}_\gamma(x) \not\subseteq \mathcal{H}} \Delta \mathcal{C}_{\bar{\Phi}, \mathcal{H}}(h, x, t), \\ \Psi(\langle 2t - 1 \rangle_\epsilon) &\leq \inf_{x \in \mathcal{X}, h \in \mathcal{H}: \bar{h}_\gamma(x) < 0} \Delta \mathcal{C}_{\bar{\Phi}, \mathcal{H}}(h, x, t). \end{aligned}$$

Then, for any hypothesis $h \in \mathcal{H}$ and any distribution,

$$\begin{aligned} &\Psi(\mathcal{R}_{\ell_\gamma}(h) - \mathcal{R}_{\ell_\gamma, \mathcal{H}}^* + \mathcal{M}_{\ell_\gamma, \mathcal{H}}) \\ &\leq \mathcal{R}_{\bar{\Phi}}(h) - \mathcal{R}_{\bar{\Phi}, \mathcal{H}}^* + \mathcal{M}_{\bar{\Phi}, \mathcal{H}} + \max\{0, \Psi(\epsilon)\}. \quad (12) \end{aligned}$$

The counterpart of Theorem 5 is Theorem 13 (adversarial distribution-independent Γ -bound), deferred to Appendix C due to space limitations. The proofs for both theorems are included in Appendix H. As with the non-adversarial scenario, the tightest distribution-independent \mathcal{H} -consistency bounds obtained by Theorem 5 can be achieved by the optimal Ψ , which is the *adversarial \mathcal{H} -estimation error transformation* defined as follows.

Definition 5 (Adversarial \mathcal{H} -estimation error transformation). *The adversarial \mathcal{H} -estimation error transformation of $\tilde{\Phi}$ is defined on $t \in [0, 1]$ by $\mathcal{T}_{\bar{\Phi}}(t) = \min\{\mathcal{T}_1(t), \mathcal{T}_2(t)\}$,*

$$\begin{aligned} \text{where } \mathcal{T}_1(t) &:= \widehat{\mathcal{T}}_1(t) \mathbb{1}_{t \in [1/2, 1]} + 2\widehat{\mathcal{T}}_1(1/2) t \mathbb{1}_{t \in [0, 1/2)}, \\ \mathcal{T}_2(t) &:= \widehat{\mathcal{T}}_2(t) \mathbb{1}_{t \in [\epsilon, 1]} + (\widehat{\mathcal{T}}_2(\epsilon)/\epsilon) t \mathbb{1}_{t \in [0, \epsilon)}, \\ \text{with } \widehat{\mathcal{T}}_1(t) &:= \inf_{x \in \mathcal{X}, h \in \bar{\mathcal{H}}_\gamma(x) \not\subseteq \mathcal{H}} \Delta \mathcal{C}_{\bar{\Phi}, \mathcal{H}}(h, x, t), \\ \widehat{\mathcal{T}}_2(t) &:= \inf_{x \in \mathcal{X}, h \in \mathcal{H}: \bar{h}_\gamma(x) < 0} \Delta \mathcal{C}_{\bar{\Phi}, \mathcal{H}}\left(h, x, \frac{t+1}{2}\right). \end{aligned}$$

It is clear that $\mathcal{T}_{\bar{\Phi}}$ satisfies assumptions in Theorem 5. The next theorem shows that it gives the tightest \mathcal{H} -consistency bound (12) under certain conditions.

Theorem 6 (Adversarial tightness). *Suppose that \mathcal{H} is symmetric and that $\epsilon = 0$. If $\mathcal{T}_{\bar{\Phi}} = \min\{\mathcal{T}_1, \mathcal{T}_2\}$ is convex with $\mathcal{T}_{\bar{\Phi}}(0) = 0$ and $\mathcal{T}_2 \leq \mathcal{T}_1$, then, for any $t \in [0, 1]$ and $\delta > 0$, there exist a distribution \mathcal{D} and a hypothesis $h \in \mathcal{H}$ such that $\mathcal{R}_{\ell_\gamma}(h) - \mathcal{R}_{\ell_\gamma, \mathcal{H}}^* + \mathcal{M}_{\ell_\gamma, \mathcal{H}} = t$ and $\mathcal{T}_{\bar{\Phi}}(t) \leq \mathcal{R}_{\bar{\Phi}}(h) - \mathcal{R}_{\bar{\Phi}, \mathcal{H}}^* + \mathcal{M}_{\bar{\Phi}, \mathcal{H}} \leq \mathcal{T}_{\bar{\Phi}}(t) + \delta$.*

The proof is included in Appendix I. In other words, when $\epsilon = 0$, if $\mathcal{T}_2 \leq \mathcal{T}_1$ and $\mathcal{T}_{\bar{\Phi}}$ is convex with $\mathcal{T}_{\bar{\Phi}}(0) = 0$, $\mathcal{T}_{\bar{\Phi}}$ is the optimal function for the distribution-independent bound (12). Moreover, if $\mathcal{T}_{\bar{\Phi}}$ is additionally invertible and non-increasing, $\mathcal{T}_{\bar{\Phi}}^{-1}$ is the optimal function for the distribution-independent bound in Theorem 13 (Appendix C) and the two bounds will be equivalent.

We will see that all these assumptions hold for cases considered in Section 6.3 and 6.4. Next, we will apply Theorem 5 along with the tightness guarantee Theorem 6 to study specific hypothesis sets and adversarial surrogate loss functions in Section 6.2 for negative results and Section 6.3 and 6.4 for positive results. A careful analysis is presented in each case (See Appendix L).

6.2. Negative results for adversarial robustness

Awasthi et al. (2021a) show that supremum-based convex loss functions of the type $\tilde{\Phi} = \sup_{x': \|x-x'\|_p \leq \gamma} \Phi(yh(x'))$, where Φ is convex and non-increasing, are *not \mathcal{H} -calibrated* with respect to ℓ_γ for \mathcal{H} containing 0, that is *regular for adversarial calibration*, e.g., \mathcal{H}_{lin} and \mathcal{H}_{NN} .

Definition 6 (Regularity for adversarial calibration). [Definition 5 in (Awasthi et al., 2021a)] *We say that a hypothesis set \mathcal{H} is regular for adversarial calibration if there exists*

Table 4: Adversarial \mathcal{H} -consistency bounds. They are completely new consistency bounds in the adversarial setting and can turn into more significant ϵ -consistency results. The minimizability gaps appearing in the bounds for the surrogates are concluded in Table 1. The detailed derivation is included in Appendix L, N.

Surrogates	Bound (\mathcal{H}_{lin})	Bound (\mathcal{H}_{NN})	Distribution set
$\tilde{\Phi}_\rho$	(54)	(59)	All distributions
$\tilde{\Phi}_{\text{hinge}}$	(67)	(71)	Massart's noise
$\tilde{\Phi}_{\text{sig}}$	(69)	(73)	Massart's noise

a distinguishing x in \mathcal{X} , that is if there exist $f, g \in \mathcal{H}$ such that $\inf_{\|x'-x\|_p \leq \gamma} f(x') > 0$ and $\sup_{\|x'-x\|_p \leq \gamma} g(x') < 0$.

Similarly, we show that there are no non-trivial adversarial \mathcal{H} -consistency bounds with respect to ℓ_γ for supremum-based convex loss functions and supremum-based symmetric loss functions (see Definition 7 below) including sigmoid loss with such hypothesis sets.

Definition 7 (Symmetric loss). We say that a margin-based loss Φ is symmetric if there exists a constant $C \geq 0$ such that $\Phi(t) + \Phi(-t) = C$ for any $t \in \mathbb{R}$, and denote it by Φ_{sym} . We also define its supremum-based counterpart as $\tilde{\Phi}_{\text{sym}} := \sup_{x': \|x-x'\|_p \leq \gamma} \Phi_{\text{sym}}(yh(x'))$ and call $\tilde{\Phi}_{\text{sym}}$ the supremum-based symmetric loss.

For the sigmoid loss $\Phi_{\text{sig}}(t) = 1 - \tanh(kt)$, $k > 0$, we have $\Phi_{\text{sig}}(t) + \Phi_{\text{sig}}(-t) = 2$, which implies that Φ_{sig} is symmetric. Note that Awasthi et al. (2021a) do not study the sigmoid loss, which is non-convex. Thus, our results below go beyond their results for convex adversarial surrogates.

Theorem 7 (Negative results for robustness). Suppose that \mathcal{H} contains 0 and is regular for adversarial calibration. Let ℓ_1 be supremum-based convex loss or supremum-based symmetric loss and $\ell_2 = \ell_\gamma$. Then, $f(t) \geq 1/2$ for any $t \geq 0$ are the only non-decreasing functions f such that (3) holds.

The proof is given in Appendix J. In other words, the function f in bound (3) must be lower bounded by 1/2 for such adversarial surrogates. Theorem 7 implies that the loss functions commonly used in practice for optimizing the adversarial loss cannot benefit from any useful \mathcal{H} -consistency bound guarantee. Instead, we show in Section 6.3 and 6.4 that the supremum-based ρ -margin loss $\tilde{\Phi}_\rho = \sup_{x': \|x-x'\|_p \leq \gamma} \Phi_\rho(yh(x'))$ proposed by Awasthi et al. (2021a) admits favorable adversarial \mathcal{H} -consistency bounds. These bounds would also imply significantly stronger results than the asymptotic \mathcal{H} -consistency guarantee in (Awasthi et al., 2021a).

6.3. Linear hypotheses

In this section, by applying Theorems 10 and 11, we derive the adversarial \mathcal{H}_{lin} -consistency bound (54) in Table 4 for

supremum-based ρ -margin loss. This is a completely new consistency bound in the adversarial setting. As with the non-adversarial case, the bound is dependent on the parameter B in linear hypothesis set and ρ in the loss function. This helps guide the choice of loss functions once the hypothesis set is fixed. More precisely, if $B > 0$ is known, we can always choose $\rho < B$ such that the bound is the tightest. Moreover, the bound can turn into more significant ϵ -consistency results in adversarial setting than the \mathcal{H} -consistency result in (Awasthi et al., 2021a).

Corollary 3. Let \mathcal{D} be a distribution over $\mathcal{X} \times \mathcal{Y}$ such that $\mathcal{M}_{\tilde{\Phi}_\rho, \mathcal{H}_{\text{lin}}} \leq \epsilon$ for some $\epsilon \geq 0$. Then, the following holds:

$$\mathcal{R}_{\ell_\gamma}(h) - \mathcal{R}_{\ell_\gamma, \mathcal{H}_{\text{lin}}}^* \leq \rho \left(\mathcal{R}_{\tilde{\Phi}_\rho}(h) - \mathcal{R}_{\tilde{\Phi}_\rho, \mathcal{H}_{\text{lin}}}^* + \epsilon \right) / \min\{B, \rho\}.$$

Awasthi et al. (2021a) show that $\tilde{\Phi}_\rho$ is \mathcal{H}_{lin} -consistent with respect to ℓ_γ when $\mathcal{M}_{\tilde{\Phi}_\rho, \mathcal{H}_{\text{lin}}} = 0$. This result can be immediately implied by Corollary 3. Moreover, Corollary 3 provides guarantees for more general cases where $\mathcal{M}_{\tilde{\Phi}_\rho, \mathcal{H}_{\text{lin}}}$ can be nonzero.

6.4. One-hidden-layer ReLU neural networks

For the one-hidden-layer ReLU neural networks \mathcal{H}_{NN} and $\tilde{\Phi}_\rho$, we have the \mathcal{H}_{NN} -consistency bound (59) in Table 4. Note $\inf_{x \in \mathcal{X}} \sup_{h \in \mathcal{H}_{\text{NN}}} h_\gamma(x)$ does not have an explicit expression. However, (59) can be further relaxed to be (60) in Appendix L.2, which is identical to the bound in the linear case modulo the replacement of B by ΛB . As in the linear case, the bound is new and also implies stronger ϵ -consistency results as follows:

Corollary 4. Let \mathcal{D} be a distribution over $\mathcal{X} \times \mathcal{Y}$ such that $\mathcal{M}_{\tilde{\Phi}_\rho, \mathcal{H}_{\text{NN}}} \leq \epsilon$ for some $\epsilon \geq 0$. Then,

$$\mathcal{R}_{\ell_\gamma}(h) - \mathcal{R}_{\ell_\gamma, \mathcal{H}_{\text{NN}}}^* \leq \rho \left(\mathcal{R}_{\tilde{\Phi}_\rho}(h) - \mathcal{R}_{\tilde{\Phi}_\rho, \mathcal{H}_{\text{NN}}}^* + \epsilon \right) / \min\{\Lambda B, \rho\}.$$

Besides the bounds for $\tilde{\Phi}_\rho$, Table 4 gives a series of results that are all new in the adversarial setting. Like the bounds in Table 2 and 3, they are all hypothesis set dependent and very useful. For example, the improved bounds for $\tilde{\Phi}_{\text{hinge}}$ and $\tilde{\Phi}_{\text{sig}}$ under noise conditions in the table can also turn into meaningful consistency results under Massart's noise condition, as shown in Section 6.5.

6.5. Guarantees under Massart's noise condition

Section 6.2 shows that non-trivial distribution-independent bounds for supremum-based hinge loss and supremum-based sigmoid loss do not exist. However, under Massart's noise condition (Definition 4), we will show that there exist non-trivial adversarial \mathcal{H} -consistency bounds for the two loss functions. Furthermore, we will see that the bounds are linear dependent as those in Section 5.5.




Figure 2: Left: tightness of bound (11) in Section 5.5. Right: tightness of bounds (54), (67) and (69) in Section 6.3 and 6.5.

As with the non-adversarial scenario, we introduce a modified adversarial \mathcal{H} -estimation error transformation in Proposition 2 (Appendix N). Using this tool, we derive adversarial \mathcal{H} -consistency bounds for $\tilde{\Phi}_{\text{hinge}}$ and $\tilde{\Phi}_{\text{sig}}$ under Massart’s noise condition in Table 4. From the bounds (67), (69), (71), and (73), we can also obtain novel ϵ -consistency results for $\tilde{\Phi}_{\text{hinge}}$ and $\tilde{\Phi}_{\text{sig}}$ with linear models and neural networks under Massart’s noise condition.

Corollary 5. *Let \mathcal{H} be \mathcal{H}_{lin} or \mathcal{H}_{NN} . Let \mathcal{D} be a distribution over $\mathcal{X} \times \mathcal{Y}$ which satisfies Massart’s noise condition with β such that $\mathcal{M}_{\tilde{\Phi}, \mathcal{H}} \leq \epsilon$ for some $\epsilon \geq 0$. Then,*

$$\mathcal{R}_{\ell_\gamma}(h) - \mathcal{R}_{\ell_\gamma, \mathcal{H}}^* \leq \frac{1 + 2\beta}{4\beta} (\mathcal{R}_{\tilde{\Phi}}(h) - \mathcal{R}_{\tilde{\Phi}, \mathcal{H}}^* + \epsilon) / \mathcal{T}(B),$$

where $\mathcal{T}(t)$ equals to $\min\{t, 1\}$ and $\tanh(kt)$ for $\tilde{\Phi}_{\text{hinge}}$ and $\tilde{\Phi}_{\text{sig}}$ respectively, B is replaced by ΛB for $\mathcal{H} = \mathcal{H}_{\text{NN}}$.

In Section 7, we will further show that these linear dependency bounds in adversarial setting are tight, along with the non-adversarial bounds we discussed earlier in Section 5.5.

7. Simulations

Here, we present experiments on simulated data to illustrate our bounds and their tightness. We generate data points $x \in \mathbb{R}$ on $[-1, +1]$. All risks are approximated by their empirical counterparts computed over 10^7 i.i.d. samples.

Non-adversarial. To demonstrate the tightness of our non-adversarial bounds, we consider a scenario where the marginal distribution is symmetric about $x = 0$ with labels flipped. With probability $\frac{1}{16}$, $(x, y) = (1, -1)$; with probability $\frac{7}{16}$, the label is $+1$ and the data follows the truncated normal distribution on $[\sigma, 1]$ with both mean and standard deviation σ . We consider Φ_{quad} , Φ_{log} and Φ_{exp} defined in Table 1. The distribution considered satisfies Massart’s noise condition with $\beta = \frac{1}{2}$. Thus, our bound (11) in Section 5.5 becomes $\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^* \leq \mathcal{R}_{\Phi}(h) - \mathcal{R}_{\Phi, \mathcal{H}_{\text{all}}}^*$, for any $h \in \mathcal{H}_{\text{all}}$ such that $\mathcal{R}_{\Phi}(h) \leq \mathcal{R}_{\Phi, \mathcal{H}_{\text{all}}}^* + 1$. All the minimal generalization errors vanish in this case. As shown in Figure 2, for $h(x) = -5x$, the bounds corresponding to Φ_{quad} , Φ_{log} and Φ_{exp} are all tight as $\sigma \rightarrow 0$.

Adversarial. To demonstrate the tightness of our adversarial bounds, the distribution is modified as follows:

with probability $\frac{1}{16}$, $(x, y) = (1, -1)$; with probability $\frac{1}{16}$, $(x, y) = (-1, +1)$; with probability $\frac{7}{8}$, the label is -1 and the data follows the truncated normal distribution on $[-1, \gamma - \sigma]$ with mean $\gamma - \sigma$ and standard deviation σ . We set $\gamma = 0.1$ and consider $\tilde{\Phi}_{\rho}$ with $\rho = 1$, $\tilde{\Phi}_{\text{hinge}}$ and $\tilde{\Phi}_{\text{sig}}$ with $k = 1$. The distribution considered satisfies Massart’s noise condition with $\beta = \frac{1}{2}$. Thus, our bounds (54), (67) and (69) in Table 4 become $\mathcal{R}_{\ell_\gamma}(h) \leq \mathcal{R}_{\tilde{\Phi}}(h)$, for any $h \in \mathcal{H}_{\text{lin}}$. As shown in Figure 2, for $h(x) = -5x$, the bounds corresponding to $\tilde{\Phi}_{\rho}$, $\tilde{\Phi}_{\text{hinge}}$ and $\tilde{\Phi}_{\text{sig}}$ are all tight as $\sigma \rightarrow 0$.

8. Conclusion

We presented an exhaustive study of \mathcal{H} -consistency bounds, including a series of new guarantees for both the non-adversarial zero-one loss function and the adversarial zero-one loss function. Our hypothesis-dependent guarantees are significantly stronger than the consistency or calibration ones. Our results include a series of theoretical and conceptual tools helpful for the analysis of other loss functions and other hypothesis sets, including multi-class classification or ranking losses. They can be further extended to the analysis of non-i.i.d. settings such as that of drifting distributions (Helmbold & Long, 1994; Long, 1999; Barve & Long, 1997; Bartlett et al., 2000; Mohri & Medina, 2012; Gama et al., 2014) or, more generally, time series prediction (Engle, 1982; Bollerslev, 1986; Brockwell & Davis, 1986; Box & Jenkins, 1990; Hamilton, 1994; Meir, 2000; Kuznetsov & Mohri, 2015; 2017; 2020). Our results can also be extended to many other loss functions, using our general proof techniques or a similar analysis.

Acknowledgements

Part of the work of A. Mao and Y. Zhong was done during their internship at Google Research.

References

Attias, I., Kontorovich, A., and Mansour, Y. Improved generalization bounds for robust learning. *arXiv preprint arXiv:1810.02180*, 2018.

Awasthi, P., Dutta, A., and Vijayaraghavan, A. On robustness to adversarial examples and polynomial optimization. In *Advances in Neural Information Processing Systems*, pp. 13737–13747, 2019.

Awasthi, P., Frank, N., and Mohri, M. Adversarial learning guarantees for linear hypotheses and neural networks. In *International Conference on Machine Learning*, pp. 431–441, 2020.

Awasthi, P., Frank, N., Mao, A., Mohri, M., and Zhong, Y. Calibration and consistency of adversarial surrogate

- losses. In *Advances in Neural Information Processing Systems*, pp. 9804–9815, 2021a.
- Awasthi, P., Frank, N., and Mohri, M. On the existence of the adversarial bayes classifier. In *Advances in Neural Information Processing Systems*, pp. 2978–2990, 2021b.
- Bao, H., Scott, C., and Sugiyama, M. Corrigendum to: Calibrated surrogate losses for adversarially robust classification. *arXiv preprint arXiv:2005.13748*, 2021.
- Bartlett, P. L., Ben-David, S., and Kulkarni, S. Learning changing concepts by exploiting the structure of change. *Machine Learning*, 41:153–174, 2000.
- Bartlett, P. L., Jordan, M. I., and McAuliffe, J. D. Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473):138–156, 2006.
- Bartlett, P. L., Bubeck, S., and Cherapanamjeri, Y. Adversarial examples in multi-layer random relu networks. *arXiv preprint arXiv:2106.12611*, 2021.
- Barve, R. D. and Long, P. M. On the complexity of learning from drifting distributions. *Information and Computation*, 138(2):101–123, 1997.
- Ben-David, S., Loker, D., Srebro, N., and Sridharan, K. Minimizing the misclassification error rate using a surrogate convex loss. *arXiv preprint arXiv:1206.6442*, 2012.
- Bollerslev, T. Generalized autoregressive conditional heteroskedasticity. *J Econometrics*, 1986.
- Box, G. E. P. and Jenkins, G. *Time Series Analysis, Forecasting and Control*. Holden-Day, Incorporated, 1990.
- Brockwell, P. J. and Davis, R. A. *Time Series: Theory and Methods*. Springer-Verlag, New York, 1986.
- Bubeck, S. and Sellke, M. A universal law of robustness via isoperimetry. *arXiv preprint arXiv:2105.12806*, 2021.
- Bubeck, S., Lee, Y. T., Price, E., and Razenshteyn, I. Adversarial examples from cryptographic pseudo-random generators. *arXiv preprint arXiv:1811.06418*, 2018a.
- Bubeck, S., Price, E., and Razenshteyn, I. Adversarial examples from computational constraints. *arXiv preprint arXiv:1805.10204*, 2018b.
- Bubeck, S., Cherapanamjeri, Y., Gidel, G., and Combes, R. T. d. A single gradient step finds adversarial examples on random two-layers neural networks. *arXiv preprint arXiv:2104.03863*, 2021.
- Carlini, N. and Wagner, D. Towards evaluating the robustness of neural networks. In *IEEE Symposium on Security and Privacy (SP)*, pp. 39–57, 2017.
- Cullina, D., Bhagoji, A. N., and Mittal, P. PAC-learning in the presence of evasion adversaries. *arXiv preprint arXiv:1806.01471*, 2018.
- Diakonikolas, I., Kane, D. M., and Manurangsi, P. The complexity of adversarially robust proper learning of halfspaces with agnostic noise. *arXiv preprint arXiv:2007.15220*, 2020.
- Engle, R. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50(4):987–1007, 1982.
- Feige, U., Mansour, Y., and Schapire, R. Learning and inference in the presence of corrupted inputs. In *Conference on Learning Theory*, pp. 637–657, 2015.
- Feige, U., Mansour, Y., and Schapire, R. E. Robust inference for multiclass classification. In *Algorithmic Learning Theory*, pp. 368–386, 2018.
- Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., and Bouchachia, H. A survey on concept drift adaptation. *ACM Computing Surveys (CSUR)*, 46, 04 2014.
- Gao, W. and Zhou, Z.-H. On the consistency of auc pairwise optimization. In *International Joint Conference on Artificial Intelligence*, pp. 939–945, 2015.
- Goodfellow, I. J., Shlens, J., and Szegedy, C. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- Hamilton, J. D. *Time series analysis*. Princeton, 1994.
- Helmhold, D. P. and Long, P. M. Tracking drifting concepts by minimizing disagreements. *Machine Learning*, 14(1): 27–46, 1994.
- Khim, J. and Loh, P.-L. Adversarial risk bounds for binary classification via function transformation. *arXiv preprint arXiv:1810.09519*, 2018.
- Kuznetsov, V. and Mohri, M. Learning theory and algorithms for forecasting non-stationary time series. In *Proceedings of NIPS*, pp. 541–549, 2015.
- Kuznetsov, V. and Mohri, M. Generalization bounds for non-stationary mixing processes. *Mach. Learn.*, 106(1): 93–117, 2017.
- Kuznetsov, V. and Mohri, M. Discrepancy-based theory and algorithms for forecasting non-stationary time series. *Ann. Math. Artif. Intell.*, 88(4):367–399, 2020.
- Kuznetsov, V., Mohri, M., and Syed, U. Multi-class deep boosting. In *Advances in Neural Information Processing Systems*, pp. 2501–2509, 2014.

- Long, P. and Servedio, R. Learning large-margin half-spaces with more malicious noise. In *Advances in Neural Information Processing Systems*, pp. 91–99, 2011.
- Long, P. and Servedio, R. Consistency versus realizable \mathcal{H} -consistency for multiclass classification. In *International Conference on Machine Learning*, pp. 801–809, 2013.
- Long, P. M. The complexity of learning according to two models of a drifting environment. *Machine Learning*, 37: 337–354, 1999.
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
- Massart, P. and Nédélec, É. Risk bounds for statistical learning. *The Annals of Statistics*, 34(5):2326–2366, 2006.
- Meir, R. Nonparametric time series prediction through adaptive model selection. *Machine Learning*, pp. 5–34, 2000.
- Mohri, M. and Medina, A. M. New analysis and algorithm for learning with drifting distributions. In *Algorithmic Learning Theory - 23rd International Conference, ALT 2012, Lyon, France, October 29-31, 2012. Proceedings*, volume 7568 of *Lecture Notes in Computer Science*, pp. 124–138. Springer, 2012.
- Mohri, M., Rostamizadeh, A., and Talwalkar, A. *Foundations of Machine Learning*. MIT Press, second edition, 2018.
- Montasser, O., Hanneke, S., and Srebro, N. \mathcal{V}_c classes are adversarially robustly learnable, but only improperly. *arXiv preprint arXiv:1902.04217*, 2019.
- Montasser, O., Hanneke, S., and Srebro, N. Reducing adversarially robust learning to non-robust pac learning. *arXiv preprint arXiv:2010.12039*, 2020.
- Robey, A., Chamon, L., Pappas, G., Hassani, H., and Ribeiro, A. Adversarial robustness with semi-infinite constrained learning. In *Advances in Neural Information Processing Systems*, pp. 6198–6215, 2021.
- Shafahi, A., Najibi, M., Ghiasi, M. A., Xu, Z., Dickerson, J., Studer, C., Davis, L. S., Taylor, G., and Goldstein, T. Adversarial training for free! In *Advances in Neural Information Processing Systems*, pp. 3353–3364, 2019.
- Steinwart, I. How to compare different loss functions and their risks. *Constructive Approximation*, 26(2):225–287, 2007.
- Tewari, A. and Bartlett, P. L. On the consistency of multi-class classification methods. *Journal of Machine Learning Research*, 8(36):1007–1025, 2007.
- Tsipras, D., Santurkar, S., Engstrom, L., Turner, A., and Madry, A. Robustness may be at odds with accuracy. *arXiv preprint arXiv:1805.12152*, 2018.
- Uematsu, K. and Lee, Y. On theoretically optimal ranking functions in bipartite ranking. *Department of Statistics, The Ohio State University, Tech. Rep*, 863, 2011.
- Viallard, P., VIDOT, E. G., Habrard, A., and Morvant, E. A pac-bayes analysis of adversarial robustness. In *Advances in Neural Information Processing Systems*, pp. 14421–14433, 2021.
- Wong, E., Rice, L., and Kolter, J. Z. Fast is better than free: Revisiting adversarial training. *arXiv preprint arXiv:2001.03994*, 2020.
- Yin, D., Ramchandran, K., and Bartlett, P. L. Rademacher complexity for adversarially robust generalization. In *International Conference of Machine Learning*, pp. 7085–7094, 2019.
- Zhang, H., Yu, Y., Jiao, J., Xing, E. P., Ghaoui, L. E., and Jordan, M. I. Theoretically principled trade-off between robustness and accuracy. *arXiv preprint arXiv:1901.08573*, 2019.
- Zhang, M. and Agarwal, S. Bayes consistency vs. \mathcal{H} -consistency: The interplay between surrogate loss functions and the scoring function class. In *Advances in Neural Information Processing Systems*, pp. 16927–16936, 2020.
- Zhang, T. Statistical behavior and consistency of classification methods based on convex risk minimization. *The Annals of Statistics*, 32(1):56–85, 2004a.
- Zhang, T. Statistical analysis of some multi-category large margin classification methods. *Journal of Machine Learning Research*, 5(Oct):1225–1251, 2004b.

Contents of Appendix

A	Related Work	14
B	Deferred Tables	15
C	Deferred Theorems	16
D	Proof of Theorem 1 and Theorem 2	17
E	Proof of Lemma 1 and Lemma 2	18
F	Comparison with Previous Results when $\mathcal{H} = \mathcal{H}_{\text{all}}$	19
F.1	Comparison with (Mohri et al., 2018, Theorem 4.7)	19
F.2	Comparison with (Bartlett et al., 2006, Theorem 1.1)	20
G	Proof of Theorem 3 and Theorem 12	20
H	Proof of Theorem 5 and Theorem 13	21
I	Proof of Theorem 4 and Theorem 6	21
J	Proof of Theorem 7	23
K	Derivation of Non-Adversarial \mathcal{H}-Consistency Bounds	23
K.1	Linear Hypotheses	23
K.1.1	Hinge Loss	24
K.1.2	Logistic Loss	25
K.1.3	Exponential Loss	27
K.1.4	Quadratic Loss	29
K.1.5	Sigmoid Loss	30
K.1.6	ρ -Margin Loss	31
K.2	One-Hidden-Layer ReLU Neural Network	32
K.2.1	Hinge Loss	33
K.2.2	Logistic Loss	34
K.2.3	Exponential Loss	36
K.2.4	Quadratic Loss	38
K.2.5	Sigmoid Loss	39
K.2.6	ρ -Margin Loss	40
L	Derivation of Adversarial \mathcal{H}-Consistency Bounds	42

L.1	Linear Hypotheses	42
L.1.1	Supremum-Based ρ -Margin Loss	42
L.2	One-Hidden-Layer ReLU Neural Networks	44
L.2.1	Supremum-Based ρ -Margin Loss	44
M	Derivation of Non-Adversarial \mathcal{H}_{all}-Consistency Bounds under Massart's Noise Condition	47
M.1	Quadratic Loss	47
M.2	Logistic Loss	48
M.3	Exponential Loss	49
N	Derivation of Adversarial \mathcal{H}-Consistency Bounds under Massart's Noise Condition	49
N.1	Linear Hypotheses	50
N.1.1	Supremum-Based Hinge Loss	50
N.1.2	Supremum-Based Sigmoid Loss	52
N.2	One-Hidden-Layer ReLU Neural Networks	54
N.2.1	Supremum-Based Hinge Loss	54
N.2.2	Supremum-Based Sigmoid Loss	57

A. Related Work

Bayes-consistency (also known as consistency) and excess error bounds between margin-based loss functions and the zero-one loss have been widely studied in the literature (Zhang, 2004a; Bartlett et al., 2006; Steinwart, 2007; Mohri et al., 2018). Consistency studies the asymptotic relation between the surrogate excess error and the target excess error while excess error bounds study the quantitative relation between them and thus is stronger. They both consider the hypothesis set of all measurable functions. Zhang (2004a), Bartlett et al. (2006), and Steinwart (2007) studied consistency via the lens of calibration and showed that calibration and consistency are equivalent in the standard binary classification when considering the hypothesis set of all measurable functions.

Zhang (2004a) studied the closeness to the optimal excess error of the zero-one loss minimizers of convex surrogates. Bartlett et al. (2006) extended the results of Zhang (2004a) and developed a general methodology for coming up with quantitative bounds between the excess error corresponding to the zero-one loss and that of margin-based surrogate loss functions for all distributions. In a more recent work, Mohri et al. (2018) simplified these results and provided different proofs for the excess error bounds of various loss functions widely used in practice. Calibration and consistency analysis have also been extended to the multi-class settings (Zhang, 2004b; Tewari & Bartlett, 2007) and to ranking problems (Uematsu & Lee, 2011; Gao & Zhou, 2015).

Bayes-consistency is not an appropriate notion when studying learning with a hypothesis set \mathcal{H} that is distinct from the family of all measurable functions. Therefore, a new hypothesis set-dependent notion namely, \mathcal{H} -consistency, has been proposed and explored in the more recent literature (Long & Servedio, 2013; Kuznetsov et al., 2014; Zhang & Agarwal, 2020). In particular, Long & Servedio (2013) argued that \mathcal{H} -consistency is a more useful notion than consistency by empirically showing that certain loss functions that are \mathcal{H} -consistent but not Bayes consistent can perform significantly better than a loss function known to be Bayes consistent. The work of Kuznetsov et al. (2014) extended the \mathcal{H} -consistency results in (Long & Servedio, 2013) to the case of structured prediction and provided positive results for \mathcal{H} -consistency of several multi-class ensemble algorithms.

In a recent work Zhang & Agarwal (2020) investigated the empirical phenomenon in (Long & Servedio, 2013) and designed a class of piecewise linear scoring functions such that minimizing a surrogate that is not \mathcal{H} -consistent over this larger class yields \mathcal{H} -consistency of linear models. For linear predictors, more general margin-based properties of convex surrogate losses are also studied in (Long & Servedio, 2011; Ben-David et al., 2012). Aiming for such margin-based error guarantees, Ben-David et al. (2012) argued that the hinge loss is optimal among convex losses.

Most recently, the notion of \mathcal{H} -consistency along with \mathcal{H} -calibration have also been studied in the context of adversarially robust classification (Bao et al., 2021; Awasthi et al., 2021a). In the adversarial scenario, in contrast to standard classification, the target loss is the adversarial zero-one loss (Goodfellow et al., 2014; Madry et al., 2017; Carlini & Wagner, 2017; Tsipras et al., 2018; Shafahi et al., 2019; Wong et al., 2020). This corresponds to the worst zero-one loss incurred over an adversarial perturbation of x within a γ -ball as measured in a norm, typically ℓ_p for $p \in [1, +\infty]$. The adversarial loss presents new challenges and makes the consistency analysis significantly more complex.

The work of Bao et al. (2021) initiated the study of \mathcal{H} -calibration with respect to the adversarial zero-one loss for the linear models. They showed that convex surrogates are not calibrated and introduced a class of nonconvex margin-based surrogate losses. They then provided sufficient conditions for such nonconvex losses to be calibrated in the linear case. The work of Awasthi et al. (2021a) extended the results in (Bao et al., 2021) to the general nonlinear hypothesis sets and pointed out that although \mathcal{H} -calibration is a necessary condition of \mathcal{H} -consistency, it is not sufficient in the adversarial scenario. They then proposed sufficient conditions which guarantee calibrated losses to be consistent in the setting of adversarially robust classification.

All the above mentioned publications either studied asymptotic properties (Bayes-consistency or \mathcal{H} -consistency) or studied quantitative relations when \mathcal{H} is the family of all measurable functions (excess error bounds). Instead, our work considers a hypothesis set-dependent quantitative relation between the surrogate estimation error and the target estimation error. This is significantly stronger than \mathcal{H} -calibration or \mathcal{H} -consistency and is also more informative than excess error bounds which correspond to a special case of our results with $\mathcal{H} = \mathcal{H}_{\text{all}}$. As a by-product, our theory contributes more significant consistency results for the poorly understood setting of adversarial robustness. There have also been recent works on different theoretical aspects of adversarial robustness such as tension between the zero-one loss and the adversarial zero-one loss (Tsipras et al., 2018; Zhang et al., 2019), computational bottlenecks for adversarial loss (Bubeck et al., 2018a;b; Awasthi et al., 2019), adversarial examples (Bartlett et al., 2021; Bubeck et al., 2021), sample complexity of adversarial surrogate

losses (Khim & Loh, 2018; Cullina et al., 2018; Yin et al., 2019; Montasser et al., 2019; Awasthi et al., 2020), computational complexity of adversarially robust linear classifiers (Diakonikolas et al., 2020), connections with PAC learning (Montasser et al., 2020; Viillard et al., 2021), perturbations beyond ℓ_p norm (Feige et al., 2015; 2018; Attias et al., 2018), adversarial robustness optimization (Robey et al., 2021), overparametrization (Bubeck & Sellke, 2021) and Bayes optimality (Awasthi et al., 2021b).

B. Deferred Tables

Table 5: Non-adversarial \mathcal{H}_{lin} -estimation error transformation ($\epsilon = 0$) and \mathcal{H}_{lin} -consistency bounds. All the bounds are hypothesis set-dependent (parameter B in \mathcal{H}_{lin}) and provide novel guarantees as discussed in Section 5.3. The minimizability gaps appearing in the bounds for the surrogates are concluded in Table 1. The detailed derivation is included in Appendix K.1.

Surrogates	$\mathcal{T}_{\Phi}(t), t \in [0, 1]$	$\mathcal{T}_{\Phi}^{-1}(t), t \in \mathbb{R}_+$	Bound
Hinge	$\min\{B, 1\} t$	$\frac{t}{\min\{B, 1\}}$	(26)
Logistic	$\begin{cases} \frac{t+1}{2} \log_2(t+1) + \frac{1-t}{2} \log_2(1-t), & t \leq \frac{e^{B-1}}{e^B+1}, \\ 1 - \frac{t+1}{2} \log_2(1+e^{-B}) - \frac{1-t}{2} \log_2(1+e^B), & t > \frac{e^{B-1}}{e^B+1}. \end{cases}$	upper bounded by $\begin{cases} \sqrt{2t}, & t \leq \frac{1}{2} \left(\frac{e^{B-1}}{e^B+1} \right)^2, \\ 2 \left(\frac{e^{B+1}}{e^B-1} \right) t, & t > \frac{1}{2} \left(\frac{e^{B-1}}{e^B+1} \right)^2. \end{cases}$	(28)
Exponential	$\begin{cases} 1 - \sqrt{1-t^2}, & t \leq \frac{e^{2B-1}}{e^{2B}+1}, \\ 1 - \frac{t+1}{2} e^{-B} - \frac{1-t}{2} e^B, & t > \frac{e^{2B-1}}{e^{2B}+1}. \end{cases}$	upper bounded by $\begin{cases} \sqrt{2t}, & t \leq \frac{1}{2} \left(\frac{e^{2B-1}}{e^{2B}+1} \right)^2, \\ 2 \left(\frac{e^{2B+1}}{e^{2B}-1} \right) t, & t > \frac{1}{2} \left(\frac{e^{2B-1}}{e^{2B}+1} \right)^2. \end{cases}$	(30)
Quadratic	$\begin{cases} t^2, & t \leq B, \\ 2Bt - B^2, & t > B. \end{cases}$	$\begin{cases} \sqrt{t}, & t \leq B^2, \\ \frac{t}{2B} + \frac{B}{2}, & t > B^2. \end{cases}$	(32)
Sigmoid	$\tanh(kB) t$	$\frac{\tanh(kB)}{t}$	(34)
ρ -Margin	$\frac{\min\{B, \rho\}}{\rho} t$	$\frac{\rho}{\min\{B, \rho\}} t$	(37)

Table 6: Non-adversarial \mathcal{H}_{NN} -estimation error transformation ($\epsilon = 0$) and \mathcal{H}_{NN} -consistency bounds. All the bounds are hypothesis set-dependent (parameter Λ and B in \mathcal{H}_{NN}) and provide novel guarantees as discussed in Section 5.4. The minimizability gaps appearing in the bounds for the surrogates are concluded in Table 1. The detailed derivation is included in Appendix K.2.

Surrogates	$\mathcal{T}_{\Phi}(t), t \in [0, 1]$	$\mathcal{T}_{\Phi}^{-1}(t), t \in \mathbb{R}_+$	Bound
Hinge	$\min\{\Lambda B, 1\} t$	$\frac{t}{\min\{\Lambda B, 1\}}$	(41)
Logistic	$\begin{cases} \frac{t+1}{2} \log_2(t+1) + \frac{1-t}{2} \log_2(1-t), & t \leq \frac{e^{\Lambda B-1}}{e^{\Lambda B}+1}, \\ 1 - \frac{t+1}{2} \log_2(1+e^{-\Lambda B}) - \frac{1-t}{2} \log_2(1+e^{\Lambda B}), & t > \frac{e^{\Lambda B-1}}{e^{\Lambda B}+1}. \end{cases}$	upper bounded by $\begin{cases} \sqrt{2t}, & t \leq \frac{1}{2} \left(\frac{e^{\Lambda B-1}}{e^{\Lambda B}+1} \right)^2, \\ 2 \left(\frac{e^{\Lambda B+1}}{e^{\Lambda B}-1} \right) t, & t > \frac{1}{2} \left(\frac{e^{\Lambda B-1}}{e^{\Lambda B}+1} \right)^2. \end{cases}$	(43)
Exponential	$\begin{cases} 1 - \sqrt{1-t^2}, & t \leq \frac{e^{2\Lambda B-1}}{e^{2\Lambda B}+1}, \\ 1 - \frac{t+1}{2} e^{-\Lambda B} - \frac{1-t}{2} e^{\Lambda B}, & t > \frac{e^{2\Lambda B-1}}{e^{2\Lambda B}+1}. \end{cases}$	upper bounded by $\begin{cases} \sqrt{2t}, & t \leq \frac{1}{2} \left(\frac{e^{2\Lambda B-1}}{e^{2\Lambda B}+1} \right)^2, \\ 2 \left(\frac{e^{2\Lambda B+1}}{e^{2\Lambda B}-1} \right) t, & t > \frac{1}{2} \left(\frac{e^{2\Lambda B-1}}{e^{2\Lambda B}+1} \right)^2. \end{cases}$	(45)
Quadratic	$\begin{cases} t^2, & t \leq \Lambda B, \\ 2\Lambda Bt - (\Lambda B)^2, & t > \Lambda B. \end{cases}$	$\begin{cases} \sqrt{t}, & t \leq (\Lambda B)^2, \\ \frac{t}{2\Lambda B} + \frac{\Lambda B}{2}, & t > (\Lambda B)^2 \end{cases}$	(47)
Sigmoid	$\tanh(k\Lambda B) t$	$\frac{\tanh(k\Lambda B)}{t}$	(49)
ρ -Margin	$\frac{\min\{\Lambda B, \rho\}}{\rho} t$	$\frac{\rho}{\min\{\Lambda B, \rho\}} t$	(51)

C. Deferred Theorems

Theorem 8 (Non-adversarial distribution-dependent Ψ -bound). *Suppose that \mathcal{H} satisfies the condition of Lemma 1 and that Φ is a margin-based loss function. Assume there exist a convex function $\Psi: \mathbb{R}_+ \rightarrow \mathbb{R}$ with $\Psi(0) = 0$ and $\epsilon \geq 0$ such that the following holds for any $x \in \mathcal{X}$:*

$$\Psi(\langle 2|\Delta\eta(x)| \rangle_\epsilon) \leq \inf_{h \in \mathcal{H}(x)} \Delta\mathcal{C}_{\Phi, \mathcal{H}}(h, x). \quad (13)$$

Then, for any hypothesis $h \in \mathcal{H}$,

$$\Psi(\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}}^* + \mathcal{M}_{\ell_{0-1}, \mathcal{H}}) \leq \mathcal{R}_{\Phi}(h) - \mathcal{R}_{\Phi, \mathcal{H}}^* + \mathcal{M}_{\Phi, \mathcal{H}} + \max\{0, \Psi(\epsilon)\}. \quad (14)$$

Theorem 9 (Non-adversarial distribution-dependent Γ -bound). *Suppose that \mathcal{H} satisfies the condition of Lemma 1 and that Φ is a margin-based loss function. Assume there exist a non-negative and non-decreasing concave function $\Gamma: \mathbb{R}_+ \rightarrow \mathbb{R}$ and $\epsilon \geq 0$ such that the following holds for any $x \in \mathcal{X}$:*

$$\langle 2|\Delta\eta(x)| \rangle_\epsilon \leq \Gamma\left(\inf_{h \in \mathcal{H}(x)} \Delta\mathcal{C}_{\Phi, \mathcal{H}}(h, x)\right). \quad (15)$$

Then, for any hypothesis $h \in \mathcal{H}$,

$$\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}}^* \leq \Gamma(\mathcal{R}_{\Phi}(h) - \mathcal{R}_{\Phi, \mathcal{H}}^* + \mathcal{M}_{\Phi, \mathcal{H}}) - \mathcal{M}_{\ell_{0-1}, \mathcal{H}} + \epsilon. \quad (16)$$

Theorem 10 (Adversarial distribution-dependent Ψ -bound). *Suppose that \mathcal{H} is symmetric and that $\tilde{\Phi}$ is a supremum-based margin loss function. Assume there exist a convex function $\Psi: \mathbb{R}_+ \rightarrow \mathbb{R}$ with $\Psi(0) = 0$ and $\epsilon \geq 0$ such that the following holds for any $x \in \mathcal{X}$:*

$$\begin{aligned} \Psi(\langle |\Delta\eta(x)| + 1/2 \rangle_\epsilon) &\leq \inf_{h \in \mathcal{H}_\gamma(x)} \Delta\mathcal{C}_{\Phi, \mathcal{H}}(h, x), \\ \Psi(\langle 2\Delta\eta(x) \rangle_\epsilon) &\leq \inf_{h \in \mathcal{H}: \bar{h}_\gamma(x) < 0} \Delta\mathcal{C}_{\Phi, \mathcal{H}}(h, x), \\ \Psi(\langle -2\Delta\eta(x) \rangle_\epsilon) &\leq \inf_{h \in \mathcal{H}: \underline{h}_\gamma(x) > 0} \Delta\mathcal{C}_{\Phi, \mathcal{H}}(h, x). \end{aligned} \quad (17)$$

Then, for any hypothesis $h \in \mathcal{H}$,

$$\Psi(\mathcal{R}_{\ell_\gamma}(h) - \mathcal{R}_{\ell_\gamma, \mathcal{H}}^* + \mathcal{M}_{\ell_\gamma, \mathcal{H}}) \leq \mathcal{R}_{\tilde{\Phi}}(h) - \mathcal{R}_{\tilde{\Phi}, \mathcal{H}}^* + \mathcal{M}_{\tilde{\Phi}, \mathcal{H}} + \max\{0, \Psi(\epsilon)\}. \quad (18)$$

Theorem 11 (Adversarial distribution-dependent Γ -bound). *Suppose that \mathcal{H} is symmetric and that $\tilde{\Phi}$ is a supremum-based margin loss function. Assume there exist a non-negative and non-decreasing concave function $\Gamma: \mathbb{R}_+ \rightarrow \mathbb{R}$ and $\epsilon \geq 0$ such that the following holds for any $x \in \mathcal{X}$:*

$$\begin{aligned} \langle |\Delta\eta(x)| + 1/2 \rangle_\epsilon &\leq \Gamma\left(\inf_{h \in \mathcal{H}_\gamma(x)} \Delta\mathcal{C}_{\Phi, \mathcal{H}}(h, x)\right), \\ \langle 2\Delta\eta(x) \rangle_\epsilon &\leq \Gamma\left(\inf_{h \in \mathcal{H}: \bar{h}_\gamma(x) < 0} \Delta\mathcal{C}_{\Phi, \mathcal{H}}(h, x)\right), \\ \langle -2\Delta\eta(x) \rangle_\epsilon &\leq \Gamma\left(\inf_{h \in \mathcal{H}: \underline{h}_\gamma(x) > 0} \Delta\mathcal{C}_{\Phi, \mathcal{H}}(h, x)\right). \end{aligned} \quad (19)$$

Then, for any hypothesis $h \in \mathcal{H}$,

$$\mathcal{R}_{\ell_\gamma}(h) - \mathcal{R}_{\ell_\gamma, \mathcal{H}}^* \leq \Gamma(\mathcal{R}_{\tilde{\Phi}}(h) - \mathcal{R}_{\tilde{\Phi}, \mathcal{H}}^* + \mathcal{M}_{\tilde{\Phi}, \mathcal{H}}) - \mathcal{M}_{\ell_\gamma, \mathcal{H}} + \epsilon. \quad (20)$$

Theorem 12 (Distribution-independent Γ -bound). *Suppose that \mathcal{H} satisfies the condition of Lemma 1 and that Φ is a margin-based loss function. Assume there exist a non-negative and non-decreasing concave function $\Gamma: \mathbb{R}_+ \rightarrow \mathbb{R}$ and $\epsilon \geq 0$ such that the following holds for any for any $t \in [1/2, 1]$:*

$$\langle 2t - 1 \rangle_\epsilon \leq \Gamma\left(\inf_{x \in \mathcal{X}, h \in \mathcal{H}: h(x) < 0} \Delta\mathcal{C}_{\Phi, \mathcal{H}}(h, x, t)\right).$$

Then, for any hypothesis $h \in \mathcal{H}$ and any distribution,

$$\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}}^* \leq \Gamma(\mathcal{R}_{\Phi}(h) - \mathcal{R}_{\Phi, \mathcal{H}}^* + \mathcal{M}_{\Phi, \mathcal{H}}) - \mathcal{M}_{\ell_{0-1}, \mathcal{H}} + \epsilon. \quad (21)$$

Theorem 13 (Adversarial distribution-independent Γ -bound). Suppose that \mathcal{H} is symmetric and that $\tilde{\Phi}$ is a supremum-based margin loss function. Assume there exist a non-negative and non-decreasing concave function $\Gamma: \mathbb{R}_+ \rightarrow \mathbb{R}$ and $\epsilon \geq 0$ such that the following holds for any for any $t \in [1/2, 1]$:

$$\begin{aligned} \langle t \rangle_{\epsilon} &\leq \Gamma \left(\inf_{x \in \mathcal{X}, h \in \overline{\mathcal{H}}_{\gamma}(x) \not\subseteq \mathcal{H}} \Delta \mathcal{C}_{\tilde{\Phi}, \mathcal{H}}(h, x, t) \right), \\ \langle 2t - 1 \rangle_{\epsilon} &\leq \Gamma \left(\inf_{x \in \mathcal{X}, h \in \mathcal{H}: \overline{h}_{\gamma}(x) < 0} \Delta \mathcal{C}_{\tilde{\Phi}, \mathcal{H}}(h, x, t) \right). \end{aligned}$$

Then, for any hypothesis $h \in \mathcal{H}$ and any distribution,

$$\mathcal{R}_{\ell_{\gamma}}(h) - \mathcal{R}_{\ell_{\gamma}, \mathcal{H}}^* \leq \Gamma(\mathcal{R}_{\tilde{\Phi}}(h) - \mathcal{R}_{\tilde{\Phi}, \mathcal{H}}^* + \mathcal{M}_{\tilde{\Phi}, \mathcal{H}}) - \mathcal{M}_{\ell_{\gamma}, \mathcal{H}} + \epsilon. \quad (22)$$

D. Proof of Theorem 1 and Theorem 2

Theorem 1 (Distribution-dependent Ψ -bound). Assume that there exists a convex function $\Psi: \mathbb{R}_+ \rightarrow \mathbb{R}$ with $\Psi(0) \geq 0$ and $\epsilon \geq 0$ such that the following holds for all $h \in \mathcal{H}$ and $x \in \mathcal{X}$:

$$\Psi(\langle \Delta \mathcal{C}_{\ell_2, \mathcal{H}}(h, x) \rangle_{\epsilon}) \leq \Delta \mathcal{C}_{\ell_1, \mathcal{H}}(h, x). \quad (6)$$

Then, the following inequality holds for any $h \in \mathcal{H}$:

$$\Psi(\mathcal{R}_{\ell_2}(h) - \mathcal{R}_{\ell_2, \mathcal{H}}^* + \mathcal{M}_{\ell_2, \mathcal{H}}) \leq \mathcal{R}_{\ell_1}(h) - \mathcal{R}_{\ell_1, \mathcal{H}}^* + \mathcal{M}_{\ell_1, \mathcal{H}} + \max\{\Psi(0), \Psi(\epsilon)\}. \quad (7)$$

Proof. For any $h \in \mathcal{H}$, since $\Psi(\Delta \mathcal{C}_{\ell_2, \mathcal{H}}(h, x) \mathbb{1}_{\Delta \mathcal{C}_{\ell_2, \mathcal{H}}(h, x) > \epsilon}) \leq \Delta \mathcal{C}_{\ell_1, \mathcal{H}}(h, x)$ for all $x \in \mathcal{X}$, we have

$$\begin{aligned} &\Psi(\mathcal{R}_{\ell_2}(h) - \mathcal{R}_{\ell_2, \mathcal{H}}^* + \mathcal{M}_{\ell_2, \mathcal{H}}) \\ &= \Psi(\mathbb{E}_X[\mathcal{C}_{\ell_2}(h, x) - \mathcal{C}_{\ell_2, \mathcal{H}}^*(x)]) \\ &= \Psi(\mathbb{E}_X[\Delta \mathcal{C}_{\ell_2, \mathcal{H}}(h, x)]) \\ &\leq \mathbb{E}_X[\Psi(\Delta \mathcal{C}_{\ell_2, \mathcal{H}}(h, x))] && \text{(Jensen's ineq.)} \\ &= \mathbb{E}_X[\Psi(\Delta \mathcal{C}_{\ell_2, \mathcal{H}}(h, x) \mathbb{1}_{\Delta \mathcal{C}_{\ell_2, \mathcal{H}}(h, x) > \epsilon} + \Delta \mathcal{C}_{\ell_2, \mathcal{H}}(h, x) \mathbb{1}_{\Delta \mathcal{C}_{\ell_2, \mathcal{H}}(h, x) \leq \epsilon})] \\ &\leq \mathbb{E}_X[\Psi(\Delta \mathcal{C}_{\ell_2, \mathcal{H}}(h, x) \mathbb{1}_{\Delta \mathcal{C}_{\ell_2, \mathcal{H}}(h, x) > \epsilon}) + \Psi(\Delta \mathcal{C}_{\ell_2, \mathcal{H}}(h, x) \mathbb{1}_{\Delta \mathcal{C}_{\ell_2, \mathcal{H}}(h, x) \leq \epsilon})] && (\Psi(0) \geq 0) \\ &\leq \mathbb{E}_X[\Delta \mathcal{C}_{\ell_1, \mathcal{H}}(h, x)] + \sup_{t \in [0, \epsilon]} \Psi(t) && \text{(assumption)} \\ &= \mathcal{R}_{\ell_1}(h) - \mathcal{R}_{\ell_1, \mathcal{H}}^* + \mathcal{M}_{\ell_1, \mathcal{H}} + \max\{\Psi(0), \Psi(\epsilon)\}, && \text{(convexity of } \Psi) \end{aligned}$$

which proves the theorem. □

Theorem 2 (Distribution-dependent Γ -bound). Assume that there exists a concave function $\Gamma: \mathbb{R}_+ \rightarrow \mathbb{R}$ and $\epsilon \geq 0$ such that the following holds for all $h \in \mathcal{H}$ and $x \in \mathcal{X}$:

$$\langle \Delta \mathcal{C}_{\ell_2, \mathcal{H}}(h, x) \rangle_{\epsilon} \leq \Gamma(\Delta \mathcal{C}_{\ell_1, \mathcal{H}}(h, x)). \quad (8)$$

Then, the following inequality holds for any $h \in \mathcal{H}$:

$$\mathcal{R}_{\ell_2}(h) - \mathcal{R}_{\ell_2, \mathcal{H}}^* \leq \Gamma(\mathcal{R}_{\ell_1}(h) - \mathcal{R}_{\ell_1, \mathcal{H}}^* + \mathcal{M}_{\ell_1, \mathcal{H}}) - \mathcal{M}_{\ell_2, \mathcal{H}} + \epsilon. \quad (9)$$

Proof. For any $h \in \mathcal{H}$, since $\Delta \mathcal{C}_{\ell_2, \mathcal{H}}(h, x) \mathbb{1}_{\Delta \mathcal{C}_{\ell_2, \mathcal{H}}(h, x) > \epsilon} \leq \Gamma(\Delta \mathcal{C}_{\ell_1, \mathcal{H}}(h, x))$ for all $x \in \mathcal{X}$, we have

$$\begin{aligned}
 & \mathcal{R}_{\ell_2}(h) - \mathcal{R}_{\ell_2, \mathcal{H}}^* + \mathcal{M}_{\ell_2, \mathcal{H}} \\
 &= \mathbb{E}_X[\mathcal{C}_{\ell_2}(h, x) - \mathcal{C}_{\ell_2, \mathcal{H}}^*(x)] \\
 &= \mathbb{E}_X[\Delta \mathcal{C}_{\ell_2, \mathcal{H}}(h, x)] \\
 &= \mathbb{E}_X[\Delta \mathcal{C}_{\ell_2, \mathcal{H}}(h, x) \mathbb{1}_{\Delta \mathcal{C}_{\ell_2, \mathcal{H}}(h, x) > \epsilon} + \Delta \mathcal{C}_{\ell_2, \mathcal{H}}(h, x) \mathbb{1}_{\Delta \mathcal{C}_{\ell_2, \mathcal{H}}(h, x) \leq \epsilon}] \\
 &\leq \mathbb{E}_X[\Gamma(\Delta \mathcal{C}_{\ell_1, \mathcal{H}}(h, x))] + \epsilon && \text{(assumption)} \\
 &\leq \Gamma(\mathbb{E}_X[\Delta \mathcal{C}_{\ell_1, \mathcal{H}}(h, x)]) + \epsilon && \text{(concavity of } \Gamma) \\
 &= \Gamma(\mathcal{R}_{\ell_1}(h) - \mathcal{R}_{\ell_1, \mathcal{H}}^* + \mathcal{M}_{\ell_1, \mathcal{H}}) + \epsilon,
 \end{aligned}$$

which proves the theorem. □

E. Proof of Lemma 1 and Lemma 2

Lemma 1. Assume that \mathcal{H} satisfies the following condition for any $x \in \mathcal{X}$: $\{\text{sign}(h(x)): h \in \mathcal{H}\} = \{-1, +1\}$. Then, the minimal conditional ℓ_{0-1} -risk is

$$\mathcal{C}_{\ell_{0-1}, \mathcal{H}}^*(x) = \mathcal{C}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^*(x) = \min\{\eta(x), 1 - \eta(x)\}.$$

The conditional ϵ -regret for ℓ_{0-1} can be characterized as

$$\langle \Delta \mathcal{C}_{\ell_{0-1}, \mathcal{H}}(h, x) \rangle_{\epsilon} = \langle 2|\Delta\eta(x)| \rangle_{\epsilon} \mathbb{1}_{h \in \overline{\mathcal{H}}(x)}.$$

Proof. By the definition, the conditional ℓ_{0-1} -risk is

$$\begin{aligned}
 \mathcal{C}_{\ell_{0-1}}(h, x) &= \eta(x) \mathbb{1}_{h(x) < 0} + (1 - \eta(x)) \mathbb{1}_{h(x) \geq 0} \\
 &= \begin{cases} \eta(x) & \text{if } h(x) < 0, \\ 1 - \eta(x) & \text{if } h(x) \geq 0. \end{cases}
 \end{aligned}$$

By the assumption, for any $x \in \mathcal{X}$, there exists $h^* \in \mathcal{H}$ such that $\text{sign}(h^*(x)) = \text{sign}(\Delta\eta(x))$, where $\Delta\eta(x)$ is the Bayes classifier such that $\mathcal{C}_{\ell_{0-1}}(\Delta\eta(x), x) = \mathcal{C}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^*(x) = \min\{\eta(x), 1 - \eta(x)\}$. Therefore, the optimal conditional ℓ_{0-1} -risk is

$$\mathcal{C}_{\ell_{0-1}, \mathcal{H}}^*(x) = \mathcal{C}_{\ell_{0-1}}(h^*, x) = \mathcal{C}_{\ell_{0-1}}(\Delta\eta(x), x) = \min\{\eta(x), 1 - \eta(x)\}$$

which proves the first part of lemma. By the definition,

$$\begin{aligned}
 \Delta \mathcal{C}_{\ell_{0-1}, \mathcal{H}}(h, x) &= \mathcal{C}_{\ell_{0-1}}(h, x) - \mathcal{C}_{\ell_{0-1}, \mathcal{H}}^*(x) \\
 &= \eta(x) \mathbb{1}_{h(x) < 0} + (1 - \eta(x)) \mathbb{1}_{h(x) \geq 0} - \min\{\eta(x), 1 - \eta(x)\} \\
 &= \begin{cases} 2|\Delta\eta(x)|, & h \in \overline{\mathcal{H}}(x), \\ 0, & \text{otherwise.} \end{cases}
 \end{aligned}$$

This leads to

$$\langle \Delta \mathcal{C}_{\ell_{0-1}, \mathcal{H}}(h, x) \rangle_{\epsilon} = \langle 2|\Delta\eta(x)| \rangle_{\epsilon} \mathbb{1}_{h \in \overline{\mathcal{H}}(x)}.$$

□

Lemma 2. Assume that \mathcal{H} is symmetric. Then, the minimal conditional ℓ_{γ} -risk is

$$\mathcal{C}_{\ell_{\gamma}, \mathcal{H}}^*(x) = \min\{\eta(x), 1 - \eta(x)\} \mathbb{1}_{\overline{\mathcal{H}}_{\gamma}(x) \neq \mathcal{H}} + \mathbb{1}_{\overline{\mathcal{H}}_{\gamma}(x) = \mathcal{H}}.$$

The conditional ϵ -regret for ℓ_γ can be characterized as

$$\langle \Delta \mathcal{C}_{\ell_\gamma, \mathcal{H}}(h, x) \rangle_\epsilon = \begin{cases} \left\langle |\Delta\eta(x)| + \frac{1}{2} \right\rangle_\epsilon & h \in \overline{\mathcal{H}}_\gamma(x) \not\subseteq \mathcal{H} \\ \langle 2\Delta\eta(x) \rangle_\epsilon & \overline{h}_\gamma(x) < 0 \\ \langle -2\Delta\eta(x) \rangle_\epsilon & \underline{h}_\gamma(x) > 0 \\ 0 & \text{otherwise} \end{cases}$$

Proof. By the definition, the conditional ℓ_γ -risk is

$$\begin{aligned} \mathcal{C}_{\ell_\gamma}(h, x) &= \eta(x) \mathbb{1}_{\{\underline{h}_\gamma(x) \leq 0\}} + (1 - \eta(x)) \mathbb{1}_{\{\overline{h}_\gamma(x) \geq 0\}} \\ &= \begin{cases} 1 & \text{if } h \in \overline{\mathcal{H}}_\gamma(x), \\ \eta(x) & \text{if } \overline{h}_\gamma(x) < 0, \\ 1 - \eta(x) & \text{if } \underline{h}_\gamma(x) > 0. \end{cases} \end{aligned}$$

Since \mathcal{H} is symmetric, for any $x \in \mathcal{X}$, either there exists $h \in \mathcal{H}$ such that $\underline{h}_\gamma(x) > 0$, or $\overline{\mathcal{H}}_\gamma(x) = \mathcal{H}$. When $\overline{\mathcal{H}}_\gamma(x) = \mathcal{H}$, $\{h \in \mathcal{H} : \overline{h}_\gamma(x) < 0\}$ and $\{h \in \mathcal{H} : \underline{h}_\gamma(x) > 0\}$ are both empty sets. Thus $\mathcal{C}_{\ell_\gamma, \mathcal{H}}^*(x) = 1$. When $\overline{\mathcal{H}}_\gamma(x) \neq \mathcal{H}$, there exists $h \in \mathcal{H}$ such that $\mathcal{C}_{\ell_\gamma}(h, x) = \min\{\eta(x), 1 - \eta(x)\} = \mathcal{C}_{\ell_\gamma, \mathcal{H}}^*(x)$. Therefore, the minimal conditional ℓ_γ -risk is

$$\mathcal{C}_{\ell_\gamma, \mathcal{H}}^*(x) = \begin{cases} 1, & \overline{\mathcal{H}}_\gamma(x) = \mathcal{H}, \\ \min\{\eta(x), 1 - \eta(x)\}, & \overline{\mathcal{H}}_\gamma(x) \neq \mathcal{H}. \end{cases}$$

When $\overline{\mathcal{H}}_\gamma(x) = \mathcal{H}$, $\mathcal{C}_{\ell_\gamma}(h, x) \equiv 1$, which implies that $\Delta \mathcal{C}_{\ell_\gamma, \mathcal{H}}(h, x) \equiv 0$. For $h \in \overline{\mathcal{H}}_\gamma(x) \not\subseteq \mathcal{H}$, $\Delta \mathcal{C}_{\ell_\gamma, \mathcal{H}}(h, x) = 1 - \min\{\eta(x), 1 - \eta(x)\} = |\Delta\eta(x)| + 1/2$; for $h \in \mathcal{H}$ such that $\overline{h}_\gamma(x) < 0$, we have $\Delta \mathcal{C}_{\ell_\gamma, \mathcal{H}}(h, x) = \eta(x) - \min\{\eta(x), 1 - \eta(x)\} = \max\{0, 2\Delta\eta(x)\}$; for $h \in \mathcal{H}$ such that $\underline{h}_\gamma(x) > 0$, $\Delta \mathcal{C}_{\ell_\gamma, \mathcal{H}}(h, x) = 1 - \eta(x) - \min\{\eta(x), 1 - \eta(x)\} = \max\{0, -2\Delta\eta(x)\}$. Therefore,

$$\Delta \mathcal{C}_{\ell_\gamma, \mathcal{H}}(h, x) = \begin{cases} |\Delta\eta(x)| + 1/2 & h \in \overline{\mathcal{H}}_\gamma(x) \not\subseteq \mathcal{H}, \\ \max\{0, 2\Delta\eta(x)\} & \overline{h}_\gamma(x) < 0, \\ \max\{0, -2\Delta\eta(x)\} & \underline{h}_\gamma(x) > 0, \\ 0 & \text{otherwise.} \end{cases}$$

This leads to

$$\langle \Delta \mathcal{C}_{\ell_\gamma, \mathcal{H}}(h, x) \rangle_\epsilon = \begin{cases} \left\langle |\Delta\eta(x)| + \frac{1}{2} \right\rangle_\epsilon & h \in \overline{\mathcal{H}}_\gamma(x) \not\subseteq \mathcal{H} \\ \langle 2\Delta\eta(x) \rangle_\epsilon & \overline{h}_\gamma(x) < 0 \\ \langle -2\Delta\eta(x) \rangle_\epsilon & \underline{h}_\gamma(x) > 0 \\ 0 & \text{otherwise} \end{cases}$$

□

F. Comparison with Previous Results when $\mathcal{H} = \mathcal{H}_{\text{all}}$

F.1. Comparison with (Mohri et al., 2018, Theorem 4.7)

Assume Φ is convex and non-increasing. For any $x \in \mathcal{X}$, by the convexity, we have

$$\mathcal{C}_\Phi(h, x) = \eta(x)\Phi(h(x)) + (1 - \eta(x))\Phi(-h(x)) \geq \Phi(2\Delta\eta(x)h(x)). \quad (23)$$

Then,

$$\begin{aligned} \inf_{h \in \overline{\mathcal{H}}_{\text{all}}(x)} \Delta \mathcal{C}_{\Phi, \mathcal{H}_{\text{all}}}(h, x) &\geq \inf_{h \in \mathcal{H}_{\text{all}}: 2\Delta\eta(x)h(x) \leq 0} \Delta \mathcal{C}_{\Phi, \mathcal{H}_{\text{all}}}(h, x) && (h \in \overline{\mathcal{H}}_{\text{all}}(x) \implies h(x)\Delta\eta(x) \leq 0) \\ &\geq \inf_{h \in \mathcal{H}_{\text{all}}: 2\Delta\eta(x)h(x) \leq 0} \Phi(2\Delta\eta(x)h(x)) - \mathcal{C}_{\Phi, \mathcal{H}_{\text{all}}}^*(x) && \text{(eq. (23))} \\ &= \mathcal{C}_\Phi(0, x) - \mathcal{C}_{\Phi, \mathcal{H}_{\text{all}}}^*(x) && (\Phi \text{ is non-increasing}). \end{aligned}$$

Thus the condition of Theorem 4.7 in (Mohri et al., 2018) implies the condition in Corollary 2:

$$|\Delta\eta(x)| \leq c \left[\mathcal{C}_\Phi(0, x) - \mathcal{C}_{\Phi, \mathcal{H}_{\text{all}}}^*(x) \right]^{\frac{1}{s}}, \quad \forall x \in \mathcal{X} \implies |\Delta\eta(x)| \leq c \inf_{h \in \mathcal{H}_{\text{all}}(x)} \left[\Delta\mathcal{C}_{\Phi, \mathcal{H}_{\text{all}}}(h, x) \right]^{\frac{1}{s}}, \quad \forall x \in \mathcal{X}.$$

Therefore, Theorem 4.7 in (Mohri et al., 2018) is a special case of Corollary 2.

F.2. Comparison with (Bartlett et al., 2006, Theorem 1.1)

We show that the ψ -transform in (Bartlett et al., 2006) verifies the condition in Corollary 1 for all distributions. First, by Definition 2 in (Bartlett et al., 2006), we know that ψ is convex, $\psi(0) = 0$ and $\psi \leq \tilde{\psi}$. Then,

$$\begin{aligned} \psi(2|\Delta\eta(x)|) &\leq \tilde{\psi}(2|\Delta\eta(x)|) && (\psi \leq \tilde{\psi}) \\ &= \inf_{\alpha \leq 0} (\max\{\eta(x), 1 - \eta(x)\} \Phi(\alpha) + \min\{\eta(x), 1 - \eta(x)\} \Phi(-\alpha)) \\ &\quad - \inf_{\alpha \in \mathbb{R}} (\max\{\eta(x), 1 - \eta(x)\} \Phi(\alpha) + \min\{\eta(x), 1 - \eta(x)\} \Phi(-\alpha)) && (\text{def. of } \tilde{\psi}) \\ &= \inf_{\alpha \Delta\eta(x) \leq 0} (\eta(x) \Phi(\alpha) + (1 - \eta(x)) \Phi(-\alpha)) - \inf_{\alpha \in \mathbb{R}} (\eta(x) \Phi(\alpha) + (1 - \eta(x)) \Phi(-\alpha)) && (\text{symmetry}) \\ &= \inf_{h \in \mathcal{H}_{\text{all}}: h(x) \Delta\eta(x) \leq 0} \Delta\mathcal{C}_{\Phi, \mathcal{H}_{\text{all}}}(h, x) \\ &\leq \inf_{h \in \overline{\mathcal{H}_{\text{all}}}(x)} \Delta\mathcal{C}_{\Phi, \mathcal{H}_{\text{all}}}(h, x) && (h \in \overline{\mathcal{H}_{\text{all}}}(x) \implies h(x) \Delta\eta(x) \leq 0) \end{aligned}$$

Therefore, Theorem 1.1 in (Bartlett et al., 2006) is a special case of Corollary 1.

G. Proof of Theorem 3 and Theorem 12

Theorem 3 (Distribution-independent Ψ -bound). Assume that \mathcal{H} satisfies the condition of Lemma 1. Assume that there exists a convex function $\Psi: \mathbb{R}_+ \rightarrow \mathbb{R}$ with $\Psi(0) = 0$ and $\epsilon \geq 0$ such that for any $t \in [1/2, 1]$,

$$\Psi(\langle 2t - 1 \rangle_\epsilon) \leq \inf_{x \in \mathcal{X}, h \in \mathcal{H}: h(x) < 0} \Delta\mathcal{C}_{\Phi, \mathcal{H}}(h, x, t).$$

Then, for any hypothesis $h \in \mathcal{H}$ and any distribution,

$$\Psi(\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}}^* + \mathcal{M}_{\ell_{0-1}, \mathcal{H}}) \leq \mathcal{R}_\Phi(h) - \mathcal{R}_{\Phi, \mathcal{H}}^* + \mathcal{M}_{\Phi, \mathcal{H}} + \max\{0, \Psi(\epsilon)\}. \quad (10)$$

Proof. Note the condition (13) in Theorem 8 is symmetric about $\Delta\eta(x) = 0$. Thus, condition (13) uniformly holds for all distributions is equivalent to the following holds for any $t \in [1/2, 1]$:

$$\Psi(\langle 2t - 1 \rangle_\epsilon) \leq \inf_{x \in \mathcal{X}, h \in \mathcal{H}: h(x) < 0} \Delta\mathcal{C}_{\Phi, \mathcal{H}}(h, x, t),$$

which proves the theorem. \square

Theorem 12 (Distribution-independent Γ -bound). Suppose that \mathcal{H} satisfies the condition of Lemma 1 and that Φ is a margin-based loss function. Assume there exist a non-negative and non-decreasing concave function $\Gamma: \mathbb{R}_+ \rightarrow \mathbb{R}$ and $\epsilon \geq 0$ such that the following holds for any for any $t \in [1/2, 1]$:

$$\langle 2t - 1 \rangle_\epsilon \leq \Gamma \left(\inf_{x \in \mathcal{X}, h \in \mathcal{H}: h(x) < 0} \Delta\mathcal{C}_{\Phi, \mathcal{H}}(h, x, t) \right).$$

Then, for any hypothesis $h \in \mathcal{H}$ and any distribution,

$$\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}}^* \leq \Gamma(\mathcal{R}_\Phi(h) - \mathcal{R}_{\Phi, \mathcal{H}}^* + \mathcal{M}_{\Phi, \mathcal{H}}) - \mathcal{M}_{\ell_{0-1}, \mathcal{H}} + \epsilon. \quad (21)$$

Proof. Note the condition (15) in Theorem 9 is symmetric about $\Delta\eta(x) = 0$. Thus, condition (15) uniformly holds for all distributions is equivalent to the following holds for any $t \in [1/2, 1]$:

$$\Psi(\langle 2t - 1 \rangle_\epsilon) \leq \inf_{x \in \mathcal{X}, h \in \mathcal{H}: h(x) < 0} \Delta\mathcal{C}_{\Phi, \mathcal{H}}(h, x, t),$$

which proves the theorem. \square

H. Proof of Theorem 5 and Theorem 13

Theorem 5 (Adversarial distribution-independent Ψ -bound). *Suppose that \mathcal{H} is symmetric. Assume there exist a convex function $\Psi: \mathbb{R}_+ \rightarrow \mathbb{R}$ with $\Psi(0) = 0$ and $\epsilon \geq 0$ such that the following holds for any $t \in [1/2, 1]$:*

$$\begin{aligned}\Psi(\langle t \rangle_\epsilon) &\leq \inf_{x \in \mathcal{X}, h \in \overline{\mathcal{H}}_\gamma(x) \not\subseteq \mathcal{H}} \Delta \mathcal{C}_{\tilde{\Phi}, \mathcal{H}}(h, x, t), \\ \Psi(\langle 2t - 1 \rangle_\epsilon) &\leq \inf_{x \in \mathcal{X}, h \in \mathcal{H}: \bar{h}_\gamma(x) < 0} \Delta \mathcal{C}_{\tilde{\Phi}, \mathcal{H}}(h, x, t).\end{aligned}$$

Then, for any hypothesis $h \in \mathcal{H}$ and any distribution,

$$\Psi(\mathcal{R}_{\ell_\gamma}(h) - \mathcal{R}_{\ell_\gamma, \mathcal{H}}^* + \mathcal{M}_{\ell_\gamma, \mathcal{H}}) \leq \mathcal{R}_{\tilde{\Phi}}(h) - \mathcal{R}_{\tilde{\Phi}, \mathcal{H}}^* + \mathcal{M}_{\tilde{\Phi}, \mathcal{H}} + \max\{0, \Psi(\epsilon)\}. \quad (12)$$

Proof. Note the condition (17) in Theorem 10 is symmetric about $\Delta\eta(x) = 0$. Thus, condition (17) uniformly holds for all distributions is equivalent to the following holds for any $t \in [1/2, 1]$:

$$\begin{aligned}\Psi(\langle t \rangle_\epsilon) &\leq \inf_{x \in \mathcal{X}, h \in \overline{\mathcal{H}}_\gamma(x) \not\subseteq \mathcal{H}} \Delta \mathcal{C}_{\tilde{\Phi}, \mathcal{H}}(h, x, t), \\ \Psi(\langle 2t - 1 \rangle_\epsilon) &\leq \inf_{x \in \mathcal{X}, h \in \mathcal{H}: \bar{h}_\gamma(x) < 0} \Delta \mathcal{C}_{\tilde{\Phi}, \mathcal{H}}(h, x, t),\end{aligned}$$

which proves the theorem. □

Theorem 13 (Adversarial distribution-independent Γ -bound). *Suppose that \mathcal{H} is symmetric and that $\tilde{\Phi}$ is a supremum-based margin loss function. Assume there exist a non-negative and non-decreasing concave function $\Gamma: \mathbb{R}_+ \rightarrow \mathbb{R}$ and $\epsilon \geq 0$ such that the following holds for any for any $t \in [1/2, 1]$:*

$$\begin{aligned}\langle t \rangle_\epsilon &\leq \Gamma \left(\inf_{x \in \mathcal{X}, h \in \overline{\mathcal{H}}_\gamma(x) \not\subseteq \mathcal{H}} \Delta \mathcal{C}_{\tilde{\Phi}, \mathcal{H}}(h, x, t) \right), \\ \langle 2t - 1 \rangle_\epsilon &\leq \Gamma \left(\inf_{x \in \mathcal{X}, h \in \mathcal{H}: \bar{h}_\gamma(x) < 0} \Delta \mathcal{C}_{\tilde{\Phi}, \mathcal{H}}(h, x, t) \right).\end{aligned}$$

Then, for any hypothesis $h \in \mathcal{H}$ and any distribution,

$$\mathcal{R}_{\ell_\gamma}(h) - \mathcal{R}_{\ell_\gamma, \mathcal{H}}^* \leq \Gamma \left(\mathcal{R}_{\tilde{\Phi}}(h) - \mathcal{R}_{\tilde{\Phi}, \mathcal{H}}^* + \mathcal{M}_{\tilde{\Phi}, \mathcal{H}} \right) - \mathcal{M}_{\ell_\gamma, \mathcal{H}} + \epsilon. \quad (22)$$

Proof. Note the condition (19) in Theorem 11 is symmetric about $\Delta\eta(x) = 0$. Thus, condition (19) uniformly holds for all distributions is equivalent to the following holds for any $t \in [1/2, 1]$:

$$\begin{aligned}\langle t \rangle_\epsilon &\leq \Gamma \left(\inf_{x \in \mathcal{X}, h \in \overline{\mathcal{H}}_\gamma(x) \not\subseteq \mathcal{H}} \Delta \mathcal{C}_{\tilde{\Phi}, \mathcal{H}}(h, x, t) \right), \\ \langle 2t - 1 \rangle_\epsilon &\leq \Gamma \left(\inf_{x \in \mathcal{X}, h \in \mathcal{H}: \bar{h}_\gamma(x) < 0} \Delta \mathcal{C}_{\tilde{\Phi}, \mathcal{H}}(h, x, t) \right),\end{aligned}$$

which proves the theorem. □

I. Proof of Theorem 4 and Theorem 6

Theorem 4 (Tightness). *Suppose that \mathcal{H} satisfies the condition of Lemma 1 and that $\epsilon = 0$. If $\mathcal{J}_{\tilde{\Phi}}$ is convex with $\mathcal{J}_{\tilde{\Phi}}(0) = 0$, then, for any $t \in [0, 1]$ and $\delta > 0$, there exist a distribution \mathcal{D} and a hypothesis $h \in \mathcal{H}$ such that $\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}}^* + \mathcal{M}_{\ell_{0-1}, \mathcal{H}} = t$ and $\mathcal{J}_{\tilde{\Phi}}(t) \leq \mathcal{R}_{\tilde{\Phi}}(h) - \mathcal{R}_{\tilde{\Phi}, \mathcal{H}}^* + \mathcal{M}_{\tilde{\Phi}, \mathcal{H}} \leq \mathcal{J}_{\tilde{\Phi}}(t) + \delta$.*

Proof. By Theorem 3, if \mathcal{T}_Φ is convex with $\mathcal{T}_\Phi(0) = 0$, the first inequality holds. For any $t \in [0, 1]$, consider the distribution that supports on a singleton $\{x_0\}$ and satisfies that $\eta(x_0) = \frac{1}{2} + \frac{t}{2}$. Thus

$$\inf_{x \in \mathcal{X}, h \in \mathcal{H}: h(x) < 0} \Delta \mathcal{C}_{\Phi, \mathcal{H}}(h, x, \eta(x_0)) = \inf_{h \in \mathcal{H}: h(x_0) < 0} \Delta \mathcal{C}_{\Phi, \mathcal{H}}(h, x_0, \eta(x_0)) = \inf_{h \in \mathcal{H}: h(x_0) < 0} \Delta \mathcal{C}_{\Phi, \mathcal{H}}(h, x_0).$$

For any $\delta > 0$, take $h_0 \in \mathcal{H}$ such that $h_0(x_0) < 0$ and

$$\Delta \mathcal{C}_{\Phi, \mathcal{H}}(h_0, x_0) \leq \inf_{h \in \mathcal{H}: h(x_0) < 0} \Delta \mathcal{C}_{\Phi, \mathcal{H}}(h, x_0) + \delta = \inf_{x \in \mathcal{X}, h \in \mathcal{H}: h(x) < 0} \Delta \mathcal{C}_{\Phi, \mathcal{H}}(h, x, \eta(x_0)) + \delta.$$

Then, we have

$$\begin{aligned} \mathcal{R}_{\ell_{0-1}}(h_0) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}}^* + \mathcal{M}_{\ell_{0-1}, \mathcal{H}} &= \mathcal{R}_{\ell_{0-1}}(h_0) - \mathbb{E}_X[\mathcal{C}_{\ell_{0-1}, \mathcal{H}}^*(x)] \\ &= \Delta \mathcal{C}_{\ell_{0-1}, \mathcal{H}}(h_0, x_0) \\ &= 2\eta(x_0) - 1 \\ &= t, \\ \mathcal{R}_\Phi(h_0) - \mathcal{R}_{\Phi, \mathcal{H}}^* + \mathcal{M}_{\Phi, \mathcal{H}} &= \mathcal{R}_\Phi(h_0) - \mathbb{E}_X[\mathcal{C}_{\Phi, \mathcal{H}}^*(x)] \\ &= \Delta \mathcal{C}_{\Phi, \mathcal{H}}(h_0, x_0) \\ &\leq \inf_{x \in \mathcal{X}, h \in \mathcal{H}: h(x) < 0} \Delta \mathcal{C}_{\Phi, \mathcal{H}}(h, x, \eta(x_0)) + \delta \\ &= \mathcal{T}_\Phi(2\eta(x_0) - 1) + \delta \\ &= \mathcal{T}_\Phi(t) + \delta, \end{aligned}$$

which completes the proof. □

Theorem 6 (Adversarial tightness). *Suppose that \mathcal{H} is symmetric and that $\epsilon = 0$. If $\mathcal{T}_{\bar{\Phi}} = \min\{\mathcal{T}_1, \mathcal{T}_2\}$ is convex with $\mathcal{T}_{\bar{\Phi}}(0) = 0$ and $\mathcal{T}_2 \leq \mathcal{T}_1$, then, for any $t \in [0, 1]$ and $\delta > 0$, there exist a distribution \mathcal{D} and a hypothesis $h \in \mathcal{H}$ such that $\mathcal{R}_{\ell_\gamma}(h) - \mathcal{R}_{\ell_\gamma, \mathcal{H}}^* + \mathcal{M}_{\ell_\gamma, \mathcal{H}} = t$ and $\mathcal{T}_{\bar{\Phi}}(t) \leq \mathcal{R}_{\bar{\Phi}}(h) - \mathcal{R}_{\bar{\Phi}, \mathcal{H}}^* + \mathcal{M}_{\bar{\Phi}, \mathcal{H}} \leq \mathcal{T}_{\bar{\Phi}}(t) + \delta$.*

Proof. By Theorem 5, if $\mathcal{T}_{\bar{\Phi}}$ is convex with $\mathcal{T}_{\bar{\Phi}}(0) = 0$, the first inequality holds. For any $t \in [0, 1]$, consider the distribution that supports on a singleton $\{x_0\}$, which satisfies that $\eta(x_0) = \frac{1}{2} + \frac{t}{2}$ and $\bar{\mathcal{H}}_\gamma(x_0) \neq \mathcal{H}$. Thus

$$\inf_{x \in \mathcal{X}, h \in \bar{\mathcal{H}}_\gamma: \bar{h}_\gamma(x) < 0} \Delta \mathcal{C}_{\bar{\Phi}, \mathcal{H}}(h, x, \eta(x_0)) = \inf_{h \in \bar{\mathcal{H}}_\gamma: \bar{h}_\gamma(x_0) < 0} \Delta \mathcal{C}_{\bar{\Phi}, \mathcal{H}}(h, x_0, \eta(x_0)) = \inf_{h \in \bar{\mathcal{H}}_\gamma: \bar{h}_\gamma(x_0) < 0} \Delta \mathcal{C}_{\bar{\Phi}, \mathcal{H}}(h, x_0).$$

For any $\delta > 0$, take $h \in \bar{\mathcal{H}}_\gamma$ such that $\bar{h}_\gamma(x_0) < 0$ and

$$\Delta \mathcal{C}_{\bar{\Phi}, \mathcal{H}}(h, x_0) \leq \inf_{h \in \bar{\mathcal{H}}_\gamma: \bar{h}_\gamma(x_0) < 0} \Delta \mathcal{C}_{\bar{\Phi}, \mathcal{H}}(h, x_0) + \delta = \inf_{x \in \mathcal{X}, h \in \bar{\mathcal{H}}_\gamma: \bar{h}_\gamma(x) < 0} \Delta \mathcal{C}_{\bar{\Phi}, \mathcal{H}}(h, x, \eta(x_0)) + \delta.$$

Then, we have

$$\begin{aligned} \mathcal{R}_{\ell_\gamma}(h) - \mathcal{R}_{\ell_\gamma, \mathcal{H}}^* + \mathcal{M}_{\ell_\gamma, \mathcal{H}} &= \mathcal{R}_{\ell_\gamma}(h) - \mathbb{E}_X[\mathcal{C}_{\ell_\gamma, \mathcal{H}}^*(x)] \\ &= \Delta \mathcal{C}_{\ell_\gamma, \mathcal{H}}(h, x_0) \\ &= 2\eta(x_0) - 1 \\ &= t, \\ \mathcal{R}_{\bar{\Phi}}(h) - \mathcal{R}_{\bar{\Phi}, \mathcal{H}}^* + \mathcal{M}_{\bar{\Phi}, \mathcal{H}} &= \mathcal{R}_{\bar{\Phi}}(h) - \mathbb{E}_X[\mathcal{C}_{\bar{\Phi}, \mathcal{H}}^*(x)] \\ &= \Delta \mathcal{C}_{\bar{\Phi}, \mathcal{H}}(h, x_0) \\ &\leq \inf_{x \in \mathcal{X}, h \in \bar{\mathcal{H}}_\gamma: \bar{h}_\gamma(x) < 0} \Delta \mathcal{C}_{\bar{\Phi}, \mathcal{H}}(h, x, \eta(x_0)) + \delta \\ &= \mathcal{T}_2(2\eta(x_0) - 1) + \delta \\ &= \mathcal{T}_{\bar{\Phi}}(2\eta(x_0) - 1) + \delta && (\mathcal{T}_2 \leq \mathcal{T}_1) \\ &= \mathcal{T}_{\bar{\Phi}}(t) + \delta \end{aligned}$$

which completes the proof. □

J. Proof of Theorem 7

Theorem 7 (Negative results for robustness). *Suppose that \mathcal{H} contains 0 and is regular for adversarial calibration. Let ℓ_1 be supremum-based convex loss or supremum-based symmetric loss and $\ell_2 = \ell_\gamma$. Then, $f(t) \geq 1/2$ for any $t \geq 0$ are the only non-decreasing functions f such that (3) holds.*

Proof. Assume $x_0 \in \mathcal{X}$ is distinguishing. Consider the distribution that supports on $\{x_0\}$. Let $\eta(x_0) = 1/2$ and $h_0 = 0 \in \mathcal{H}$. Then, for any $h \in \mathcal{H}$,

$$\mathcal{R}_{\ell_\gamma}(h) = \mathcal{C}_{\ell_\gamma}(h, x_0) = 1/2 \mathbb{1}_{\underline{h}_\gamma(x_0) \leq 0} + 1/2 \mathbb{1}_{\bar{h}_\gamma(x_0) \geq 0} \geq 1/2,$$

where the equality can be achieved for some $h \in \mathcal{H}$ since x_0 is distinguishing. Therefore,

$$\mathcal{R}_{\ell_\gamma, \mathcal{H}}^* = \mathcal{C}_{\ell_\gamma, \mathcal{H}}^*(x_0) = \inf_{h \in \mathcal{H}} \mathcal{C}_{\ell_\gamma}(h, x_0) = 1/2.$$

Note $\mathcal{R}_{\ell_\gamma}(h_0) = 1/2 + 1/2 = 1$. For the supremum-based convex loss $\tilde{\Phi}$, for any $h \in \mathcal{H}$,

$$\begin{aligned} \mathcal{R}_{\tilde{\Phi}}(h) &= \mathcal{C}_{\tilde{\Phi}}(h, x_0) = 1/2 \Phi(\underline{h}_\gamma(x_0)) + 1/2 \Phi(-\bar{h}_\gamma(x_0)) \\ &\geq \Phi(1/2 \underline{h}_\gamma(x_0) - 1/2 \bar{h}_\gamma(x_0)) && \text{(convexity of } \Phi) \\ &\geq \Phi(0), && \text{(} \Phi \text{ is non-increasing)} \end{aligned}$$

where both equality can be achieved by $h_0 = 0$. Therefore,

$$\mathcal{R}_{\tilde{\Phi}, \mathcal{H}}^* = \mathcal{C}_{\tilde{\Phi}, \mathcal{H}}^*(x_0) = \mathcal{R}_{\tilde{\Phi}}(h_0) = \Phi(0).$$

If (3) holds for some non-decreasing function f , then, we obtain for any $h \in \mathcal{H}$,

$$\mathcal{R}_{\ell_\gamma}(h) - 1/2 \leq f(\mathcal{R}_{\tilde{\Phi}}(h) - \Phi(0)).$$

Let $h = h_0$, then $f(0) \geq 1/2$. Since f is non-decreasing, for any $t \in [0, 1]$, $f(t) \geq 1/2$.

For the supremum-based symmetric loss $\tilde{\Phi}_{\text{sym}}$, there exists a constant $C \geq 0$ such that, for any $h \in \mathcal{H}$,

$$\begin{aligned} \mathcal{R}_{\tilde{\Phi}_{\text{sym}}}(h) &= \mathcal{C}_{\tilde{\Phi}_{\text{sym}}}(h, x_0) = 1/2 \Phi_{\text{sym}}(\underline{h}_\gamma(x_0)) + 1/2 \Phi_{\text{sym}}(-\bar{h}_\gamma(x_0)) \\ &\geq 1/2 \Phi_{\text{sym}}(\bar{h}_\gamma(x_0)) + 1/2 \Phi_{\text{sym}}(-\bar{h}_\gamma(x_0)) \\ &\geq \frac{C}{2} \end{aligned}$$

where the equality can be achieved by $h_0 = 0$. Therefore,

$$\mathcal{R}_{\tilde{\Phi}_{\text{sym}}, \mathcal{H}}^* = \mathcal{C}_{\tilde{\Phi}_{\text{sym}}, \mathcal{H}}^*(x_0) = \mathcal{R}_{\tilde{\Phi}_{\text{sym}}}(h_0) = \frac{C}{2}.$$

If (3) holds for some non-decreasing function f , then, we obtain for any $h \in \mathcal{H}$,

$$\mathcal{R}_{\ell_\gamma}(h) - 1/2 \leq f\left(\mathcal{R}_{\tilde{\Phi}_{\text{sym}}}(h) - \frac{C}{2}\right).$$

Let $h = h_0$, then $f(0) \geq 1/2$. Since f is non-decreasing, for any $t \in [0, 1]$, $f(t) \geq 1/2$. □

K. Derivation of Non-Adversarial \mathcal{H} -Consistency Bounds

K.1. Linear Hypotheses

Since \mathcal{H}_{lin} satisfies the condition of Lemma 1, by Lemma 1 the $(\ell_{0-1}, \mathcal{H}_{\text{lin}})$ -minimizability gap can be expressed as follows:

$$\begin{aligned} \mathcal{M}_{\ell_{0-1}, \mathcal{H}_{\text{lin}}} &= \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}] \\ &= \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{lin}}}^* - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^*. \end{aligned} \tag{24}$$

Therefore, the $(\ell_{0-1}, \mathcal{H}_{\text{lin}})$ -minimizability gap coincides with the $(\ell_{0-1}, \mathcal{H}_{\text{lin}})$ -approximation error. By the definition of \mathcal{H}_{lin} , for any $x \in \mathcal{X}$, $\{h(x) \mid h \in \mathcal{H}_{\text{lin}}\} = [-W\|x\|_p - B, W\|x\|_p + B]$.

K.1.1. HINGE LOSS

For the hinge loss $\Phi_{\text{hinge}}(\alpha) := \max\{0, 1 - \alpha\}$, for all $h \in \mathcal{H}_{\text{lin}}$ and $x \in \mathcal{X}$:

$$\begin{aligned} \mathcal{C}_{\Phi_{\text{hinge}}}(h, x, t) &= t\Phi_{\text{hinge}}(h(x)) + (1-t)\Phi_{\text{hinge}}(-h(x)) \\ &= t\max\{0, 1 - h(x)\} + (1-t)\max\{0, 1 + h(x)\}. \\ \inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\Phi_{\text{hinge}}}(h, x, t) &= 1 - |2t - 1| \min\{W\|x\|_p + B, 1\}. \end{aligned}$$

Therefore, the $(\Phi_{\text{hinge}}, \mathcal{H}_{\text{lin}})$ -minimizability gap can be expressed as follows:

$$\begin{aligned} \mathcal{M}_{\Phi_{\text{hinge}}, \mathcal{H}_{\text{lin}}} &= \mathcal{R}_{\Phi_{\text{hinge}}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X \left[1 - \inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\Phi_{\text{hinge}}}(h, x, \eta(x)) \right]. \\ &= \mathcal{R}_{\Phi_{\text{hinge}}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X \left[1 - |2\eta(x) - 1| \min\{W\|x\|_p + B, 1\} \right]. \end{aligned} \quad (25)$$

Note the $(\Phi_{\text{hinge}}, \mathcal{H}_{\text{lin}})$ -minimizability gap coincides with the $(\Phi_{\text{hinge}}, \mathcal{H}_{\text{lin}})$ -approximation error $\mathcal{R}_{\Phi_{\text{hinge}}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X[1 - |2\eta(x) - 1|]$ for $B \geq 1$.

For $\frac{1}{2} < t \leq 1$, we have

$$\begin{aligned} \inf_{h \in \mathcal{H}_{\text{lin}}: h(x) < 0} \mathcal{C}_{\Phi_{\text{hinge}}}(h, x, t) &= t\max\{0, 1 - 0\} + (1-t)\max\{0, 1 + 0\} \\ &= 1. \\ \inf_{x \in \mathcal{X}} \inf_{h \in \mathcal{H}_{\text{lin}}: h(x) < 0} \Delta \mathcal{C}_{\Phi_{\text{hinge}}, \mathcal{H}_{\text{lin}}}(h, x, t) &= \inf_{x \in \mathcal{X}} \left\{ \inf_{h \in \mathcal{H}_{\text{lin}}: h(x) < 0} \mathcal{C}_{\Phi_{\text{hinge}}}(h, x, t) - \inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\Phi_{\text{hinge}}}(h, x, t) \right\} \\ &= \inf_{x \in \mathcal{X}} (2t - 1) \min\{W\|x\|_p + B, 1\} \\ &= (2t - 1) \min\{B, 1\} \\ &= \mathcal{T}(2t - 1), \end{aligned}$$

where \mathcal{T} is the increasing and convex function on $[0, 1]$ defined by

$$\forall t \in [0, 1], \quad \mathcal{T}(t) = \min\{B, 1\}t.$$

By Definition 3, for any $\epsilon \geq 0$, the \mathcal{H}_{lin} -estimation error transformation of the hinge loss is as follows:

$$\mathcal{T}_{\Phi_{\text{hinge}}} = \min\{B, 1\}t, \quad t \in [0, 1],$$

Therefore, $\mathcal{T}_{\Phi_{\text{hinge}}}$ is convex, non-decreasing, invertible and satisfies that $\mathcal{T}_{\Phi_{\text{hinge}}}(0) = 0$. By Theorem 4, we can choose $\Psi(t) = \min\{B, 1\}t$ in Theorem 3, or, equivalently, $\Gamma(t) = \frac{t}{\min\{B, 1\}}$ in Theorem 12, which are optimal when $\epsilon = 0$. Thus, by Theorem 3 or Theorem 12, setting $\epsilon = 0$ yields the \mathcal{H}_{lin} -consistency bound for the hinge loss, valid for all $h \in \mathcal{H}_{\text{lin}}$:

$$\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{lin}}}^* \leq \frac{\mathcal{R}_{\Phi_{\text{hinge}}}(h) - \mathcal{R}_{\Phi_{\text{hinge}}, \mathcal{H}_{\text{lin}}}^* + \mathcal{M}_{\Phi_{\text{hinge}}, \mathcal{H}_{\text{lin}}}}{\min\{B, 1\}} - \mathcal{M}_{\ell_{0-1}, \mathcal{H}_{\text{lin}}}. \quad (26)$$

Since the $(\ell_{0-1}, \mathcal{H}_{\text{lin}})$ -minimizability gap coincides with the $(\ell_{0-1}, \mathcal{H}_{\text{lin}})$ -approximation error and $(\Phi_{\text{hinge}}, \mathcal{H}_{\text{lin}})$ -minimizability gap coincides with the $(\Phi_{\text{hinge}}, \mathcal{H}_{\text{lin}})$ -approximation error for $B \geq 1$, the inequality can be rewritten as follows:

$$\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^* \leq \begin{cases} \mathcal{R}_{\Phi_{\text{hinge}}}(h) - \mathcal{R}_{\Phi_{\text{hinge}}, \mathcal{H}_{\text{all}}}^* & \text{if } B \geq 1 \\ \frac{1}{B} \left[\mathcal{R}_{\Phi_{\text{hinge}}}(h) - \mathbb{E}_X \left[1 - |2\eta(x) - 1| \min\{W\|x\|_p + B, 1\} \right] \right] & \text{otherwise.} \end{cases}$$

The inequality for $B \geq 1$ coincides with the consistency excess error bound known for the hinge loss (Zhang, 2004a; Bartlett et al., 2006; Mohri et al., 2018) but the one for $B < 1$ is distinct and novel. For $B < 1$, we have

$$\mathbb{E}_X \left[1 - |2\eta(x) - 1| \min\{W\|x\|_p + B, 1\} \right] > \mathbb{E}_X \left[1 - |2\eta(x) - 1| \right] = 2\mathbb{E}_X \left[\min\{\eta(x), 1 - \eta(x)\} \right] = \mathcal{R}_{\Phi_{\text{hinge}}, \mathcal{H}_{\text{all}}}^*.$$

Therefore for $B < 1$,

$$\mathcal{R}_{\Phi_{\text{hinge}}}(h) - \mathbb{E}_X[1 - |2\eta(x) - 1| \min\{W\|x\|_p + B, 1\}] < \mathcal{R}_{\Phi_{\text{hinge}}}(h) - \mathcal{R}_{\Phi_{\text{hinge}}, \mathcal{H}_{\text{all}}}^*.$$

Note that: $\mathcal{R}_{\Phi_{\text{hinge}}, \mathcal{H}_{\text{all}}}^* = 2\mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^* = 2\mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}]$. Thus, the first inequality (case $B \geq 1$) can be equivalently written as follows:

$$\forall h \in \mathcal{H}_{\text{lin}}, \mathcal{R}_{\ell_{0-1}}(h) \leq \mathcal{R}_{\Phi_{\text{hinge}}}(h) - \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}],$$

which is a more informative upper bound than the standard inequality $\mathcal{R}_{\ell_{0-1}}(h) \leq \mathcal{R}_{\Phi_{\text{hinge}}}(h)$.

K.1.2. LOGISTIC LOSS

For the logistic loss $\Phi_{\log}(\alpha) := \log_2(1 + e^{-\alpha})$, for all $h \in \mathcal{H}_{\text{lin}}$ and $x \in \mathcal{X}$:

$$\begin{aligned} \mathcal{C}_{\Phi_{\log}}(h, x, t) &= t\Phi_{\log}(h(x)) + (1-t)\Phi_{\log}(-h(x)) \\ &= t \log_2(1 + e^{-h(x)}) + (1-t) \log_2(1 + e^{h(x)}). \\ \inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\Phi_{\log}}(h, x, t) &= \begin{cases} -t \log_2(t) - (1-t) \log_2(1-t) & \text{if } \log\left|\frac{t}{1-t}\right| \leq W\|x\|_p + B, \\ \max\{t, 1-t\} \log_2(1 + e^{-(W\|x\|_p + B)}) + \min\{t, 1-t\} \log_2(1 + e^{W\|x\|_p + B}) & \text{if } \log\left|\frac{t}{1-t}\right| > W\|x\|_p + B. \end{cases} \end{aligned}$$

Therefore, the $(\Phi_{\log}, \mathcal{H}_{\text{lin}})$ -minimizability gap can be expressed as follows:

$$\begin{aligned} \mathcal{M}_{\Phi_{\log}, \mathcal{H}_{\text{lin}}} &= \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X \left[\inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\Phi_{\log}}(h, x, \eta(x)) \right] \\ &= \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X \left[-\eta(x) \log_2(\eta(x)) - (1-\eta(x)) \log_2(1-\eta(x)) \mathbb{1}_{\log\left|\frac{\eta(x)}{1-\eta(x)}\right| \leq W\|x\|_p + B} \right] \\ &\quad - \mathbb{E}_X \left[\max\{\eta(x), 1-\eta(x)\} \log_2(1 + e^{-(W\|x\|_p + B)}) \mathbb{1}_{\log\left|\frac{\eta(x)}{1-\eta(x)}\right| > W\|x\|_p + B} \right] \\ &\quad - \mathbb{E}_X \left[\min\{\eta(x), 1-\eta(x)\} \log_2(1 + e^{W\|x\|_p + B}) \mathbb{1}_{\log\left|\frac{\eta(x)}{1-\eta(x)}\right| > W\|x\|_p + B} \right] \end{aligned} \quad (27)$$

Note $(\Phi_{\log}, \mathcal{H}_{\text{lin}})$ -minimizability gap coincides with the $(\Phi_{\log}, \mathcal{H}_{\text{lin}})$ -approximation error $\mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X[-\eta(x) \log_2(\eta(x)) - (1-\eta(x)) \log_2(1-\eta(x))]$ for $B = +\infty$.

For $\frac{1}{2} < t \leq 1$, we have

$$\begin{aligned} \inf_{h \in \mathcal{H}_{\text{lin}}: h(x) < 0} \mathcal{C}_{\Phi_{\log}}(h, x, t) &= t \log_2(1 + e^{-0}) + (1-t) \log_2(1 + e^0) \\ &= 1, \\ \inf_{x \in \mathcal{X}} \inf_{h \in \mathcal{H}_{\text{lin}}: h(x) < 0} \Delta \mathcal{C}_{\Phi_{\log}, \mathcal{H}_{\text{lin}}}(h, x, t) &= \inf_{x \in \mathcal{X}} \left(\inf_{h \in \mathcal{H}_{\text{lin}}: h(x) < 0} \mathcal{C}_{\Phi_{\log}}(h, x, t) - \inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\Phi_{\log}}(h, x, t) \right) \\ &= \inf_{x \in \mathcal{X}} \begin{cases} 1 + t \log_2(t) + (1-t) \log_2(1-t) \\ \text{if } \log\left|\frac{t}{1-t}\right| \leq W\|x\|_p + B, \\ 1 - t \log_2(1 + e^{-(W\|x\|_p + B)}) - (1-t) \log_2(1 + e^{W\|x\|_p + B}) \\ \text{if } \log\left|\frac{t}{1-t}\right| > W\|x\|_p + B. \end{cases} \\ &= \begin{cases} 1 + t \log_2(t) + (1-t) \log_2(1-t) & \text{if } \log\left|\frac{t}{1-t}\right| \leq B, \\ 1 - t \log_2(1 + e^{-B}) - (1-t) \log_2(1 + e^B) & \text{if } \log\left|\frac{t}{1-t}\right| > B. \end{cases} \\ &= \mathcal{J}(2t - 1), \end{aligned}$$

where \mathcal{J} is the increasing and convex function on $[0, 1]$ defined by

$$\forall t \in [0, 1], \quad \mathcal{J}(t) = \begin{cases} \frac{t+1}{2} \log_2(t+1) + \frac{1-t}{2} \log_2(1-t), & t \leq \frac{e^B - 1}{e^B + 1}, \\ 1 - \frac{t+1}{2} \log_2(1 + e^{-B}) - \frac{1-t}{2} \log_2(1 + e^B), & t > \frac{e^B - 1}{e^B + 1}. \end{cases}$$

By Definition 3, for any $\epsilon \geq 0$, the \mathcal{H}_{lin} -estimation error transformation of the logistic loss is as follows:

$$\mathcal{T}_{\Phi_{\log}} = \begin{cases} \mathcal{T}(t), & t \in [\epsilon, 1], \\ \frac{\mathcal{T}(\epsilon)}{\epsilon} t, & t \in [0, \epsilon]. \end{cases}$$

Therefore, when $\epsilon = 0$, $\mathcal{T}_{\Phi_{\log}}$ is convex, non-decreasing, invertible and satisfies that $\mathcal{T}_{\Phi_{\log}}(0) = 0$. By Theorem 4, we can choose $\Psi(t) = \mathcal{T}_{\Phi_{\log}}(t)$ in Theorem 3, or equivalently $\Gamma(t) = \mathcal{T}_{\Phi_{\log}}^{-1}(t)$ in Theorem 12, which are optimal. To simplify the expression, using the fact that

$$\begin{aligned} \frac{t+1}{2} \log_2(t+1) + \frac{1-t}{2} \log_2(1-t) &= 1 - \left(-\frac{t+1}{2} \log_2\left(\frac{t+1}{2}\right) - \frac{1-t}{2} \log_2\left(\frac{1-t}{2}\right) \right) \\ &\geq 1 - \sqrt{4 \frac{1-t}{2} \frac{t+1}{2}} \\ &= 1 - \sqrt{1-t^2} \\ &\geq \frac{t^2}{2}, \\ 1 - \frac{t+1}{2} \log_2(1+e^{-B}) - \frac{1-t}{2} \log_2(1+e^B) &= \frac{1}{2} \log_2\left(\frac{4}{2+e^{-B}+e^B}\right) + 1/2 \log_2\left(\frac{1+e^B}{1+e^{-B}}\right)t, \end{aligned}$$

$\mathcal{T}_{\Phi_{\log}}$ can be lower bounded by

$$\tilde{\mathcal{T}}_{\Phi_{\log}}(t) = \begin{cases} \frac{t^2}{2}, & t \leq \frac{e^B-1}{e^B+1}, \\ \frac{1}{2} \left(\frac{e^B-1}{e^B+1} \right) t, & t > \frac{e^B-1}{e^B+1}. \end{cases}$$

Thus, we adopt an upper bound of $\mathcal{T}_{\Phi_{\log}}^{-1}$ as follows:

$$\tilde{\mathcal{T}}_{\Phi_{\log}}^{-1}(t) = \begin{cases} \sqrt{2t}, & t \leq \frac{1}{2} \left(\frac{e^B-1}{e^B+1} \right)^2, \\ 2 \left(\frac{e^B+1}{e^B-1} \right) t, & t > \frac{1}{2} \left(\frac{e^B-1}{e^B+1} \right)^2. \end{cases}$$

Therefore, by Theorem 3 or Theorem 12, setting $\epsilon = 0$ yields the \mathcal{H}_{lin} -consistency bound for the logistic loss, valid for all $h \in \mathcal{H}_{\text{lin}}$:

$$\begin{aligned} \mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{lin}}}^* + \mathcal{M}_{\ell_{0-1}, \mathcal{H}_{\text{lin}}} &\leq \begin{cases} \sqrt{2} \left(\mathcal{R}_{\Phi_{\log}}(h) - \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{lin}}}^* + \mathcal{M}_{\Phi_{\log}, \mathcal{H}_{\text{lin}}} \right)^{\frac{1}{2}}, & \text{if } \mathcal{R}_{\Phi_{\log}}(h) - \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{lin}}}^* \leq \frac{1}{2} \left(\frac{e^B-1}{e^B+1} \right)^2 - \mathcal{M}_{\Phi_{\log}, \mathcal{H}_{\text{lin}}} \\ 2 \left(\frac{e^B+1}{e^B-1} \right) \left(\mathcal{R}_{\Phi_{\log}}(h) - \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{lin}}}^* + \mathcal{M}_{\Phi_{\log}, \mathcal{H}_{\text{lin}}} \right), & \text{otherwise} \end{cases} \quad (28) \end{aligned}$$

Since the $(\ell_{0-1}, \mathcal{H}_{\text{lin}})$ -minimizability gap coincides with the $(\ell_{0-1}, \mathcal{H}_{\text{lin}})$ -approximation error and $(\Phi_{\log}, \mathcal{H}_{\text{lin}})$ -minimizability gap coincides with the $(\Phi_{\log}, \mathcal{H}_{\text{lin}})$ -approximation error for $B = +\infty$, the inequality can be rewritten as follows:

$$\begin{aligned} \mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^* &\leq \begin{cases} \sqrt{2} \left[\mathcal{R}_{\Phi_{\log}}(h) - \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{all}}}^* \right]^{\frac{1}{2}} & \text{if } B = +\infty \\ \sqrt{2} \left[\mathcal{R}_{\Phi_{\log}}(h) - \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{lin}}}^* + \mathcal{M}_{\Phi_{\log}, \mathcal{H}_{\text{lin}}} \right]^{\frac{1}{2}} & \text{if } \mathcal{R}_{\Phi_{\log}}(h) - \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{lin}}}^* \leq \frac{1}{2} \left(\frac{e^B-1}{e^B+1} \right)^2 - \mathcal{M}_{\Phi_{\log}, \mathcal{H}_{\text{lin}}} \\ 2 \left(\frac{e^B+1}{e^B-1} \right) \left(\mathcal{R}_{\Phi_{\log}}(h) - \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{lin}}}^* + \mathcal{M}_{\Phi_{\log}, \mathcal{H}_{\text{lin}}} \right) & \text{otherwise} \end{cases} \end{aligned}$$

where the $(\Phi_{\log}, \mathcal{H}_{\text{lin}})$ -minimizability gap $\mathcal{M}_{\Phi_{\log}, \mathcal{H}_{\text{lin}}}$ is characterized as below, which is less than the $(\Phi_{\log}, \mathcal{H}_{\text{lin}})$ -

approximation error when $B < +\infty$:

$$\begin{aligned}
 \mathcal{M}_{\Phi_{\log}, \mathcal{H}_{\text{lin}}} &= \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X \left[-\eta(x) \log_2(\eta(x)) - (1 - \eta(x)) \log_2(1 - \eta(x)) \mathbb{1}_{\log \left| \frac{\eta(x)}{1-\eta(x)} \right| \leq W \|x\|_p + B} \right] \\
 &\quad - \mathbb{E}_X \left[\max\{\eta(x), 1 - \eta(x)\} \log_2(1 + e^{-(W \|x\|_p + B)}) \mathbb{1}_{\log \left| \frac{\eta(x)}{1-\eta(x)} \right| > W \|x\|_p + B} \right] \\
 &\quad - \mathbb{E}_X \left[\min\{\eta(x), 1 - \eta(x)\} \log_2(1 + e^{W \|x\|_p + B}) \mathbb{1}_{\log \left| \frac{\eta(x)}{1-\eta(x)} \right| > W \|x\|_p + B} \right] \\
 &< \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X [-\eta(x) \log_2(\eta(x)) - (1 - \eta(x)) \log_2(1 - \eta(x))] \\
 &= \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{lin}}}^* - \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{all}}}^*.
 \end{aligned}$$

Therefore, the inequality for $B = +\infty$ coincides with the consistency excess error bound known for the logistic loss (Zhang, 2004a; Mohri et al., 2018) but the one for $B < +\infty$ is distinct and novel.

K.1.3. EXPONENTIAL LOSS

For the exponential loss $\Phi_{\text{exp}}(\alpha) := e^{-\alpha}$, for all $h \in \mathcal{H}_{\text{lin}}$ and $x \in \mathcal{X}$:

$$\begin{aligned}
 \mathcal{C}_{\Phi_{\text{exp}}}(h, x, t) &= t\Phi_{\text{exp}}(h(x)) + (1 - t)\Phi_{\text{exp}}(-h(x)) \\
 &= te^{-h(x)} + (1 - t)e^{h(x)}. \\
 \inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\Phi_{\text{exp}}}(h, x, t) &= \begin{cases} 2\sqrt{t(1-t)} & \text{if } 1/2 \log \left| \frac{t}{1-t} \right| \leq W \|x\|_p + B \\ \max\{t, 1-t\}e^{-(W \|x\|_p + B)} + \min\{t, 1-t\}e^{W \|x\|_p + B} & \text{if } 1/2 \log \left| \frac{t}{1-t} \right| > W \|x\|_p + B. \end{cases}
 \end{aligned}$$

Therefore, the $(\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}})$ -minimizability gap can be expressed as follows:

$$\begin{aligned}
 \mathcal{M}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}}} &= \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X \left[\inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\Phi_{\text{exp}}}(h, x, \eta(x)) \right] \\
 &= \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X \left[2\sqrt{\eta(x)(1-\eta(x))} \mathbb{1}_{1/2 \log \left| \frac{\eta(x)}{1-\eta(x)} \right| \leq W \|x\|_p + B} \right] \\
 &\quad - \mathbb{E}_X \left[\max\{\eta(x), 1 - \eta(x)\} e^{-(W \|x\|_p + B)} \mathbb{1}_{1/2 \log \left| \frac{\eta(x)}{1-\eta(x)} \right| > W \|x\|_p + B} \right] \\
 &\quad - \mathbb{E}_X \left[\min\{\eta(x), 1 - \eta(x)\} e^{W \|x\|_p + B} \mathbb{1}_{1/2 \log \left| \frac{\eta(x)}{1-\eta(x)} \right| > W \|x\|_p + B} \right].
 \end{aligned} \tag{29}$$

Note $(\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}})$ -minimizability gap coincides with the $(\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}})$ -approximation error $\mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X [2\sqrt{\eta(x)(1-\eta(x))}]$ for $B = +\infty$.

For $\frac{1}{2} < t \leq 1$, we have

$$\begin{aligned}
 \inf_{h \in \mathcal{H}_{\text{lin}}: h(x) < 0} \mathcal{C}_{\Phi_{\text{exp}}}(h, x, t) &= te^{-0} + (1 - t)e^0 \\
 &= 1. \\
 \inf_{x \in \mathcal{X}} \inf_{h \in \mathcal{H}_{\text{lin}}: h(x) < 0} \Delta \mathcal{C}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}}}(h, x, t) &= \inf_{x \in \mathcal{X}} \left(\inf_{h \in \mathcal{H}_{\text{lin}}: h(x) < 0} \mathcal{C}_{\Phi_{\text{exp}}}(h, x, t) - \inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\Phi_{\text{exp}}}(h, x, t) \right) \\
 &= \inf_{x \in \mathcal{X}} \begin{cases} 1 - 2\sqrt{t(1-t)} & \text{if } 1/2 \log \frac{t}{1-t} \leq W \|x\|_p + B, \\ 1 - te^{-(W \|x\|_p + B)} - (1 - t)e^{W \|x\|_p + B} & \text{if } 1/2 \log \frac{t}{1-t} > W \|x\|_p + B. \end{cases} \\
 &= \begin{cases} 1 - 2\sqrt{t(1-t)}, & 1/2 \log \frac{t}{1-t} \leq B \\ 1 - te^{-B} - (1 - t)e^B, & 1/2 \log \frac{t}{1-t} > B \end{cases} \\
 &= \mathcal{T}(2t - 1),
 \end{aligned}$$

where \mathcal{T} is the increasing and convex function on $[0, 1]$ defined by

$$\forall t \in [0, 1], \quad \mathcal{T}(t) = \begin{cases} 1 - \sqrt{1 - t^2}, & t \leq \frac{e^{2B} - 1}{e^{2B} + 1}, \\ 1 - \frac{t+1}{2}e^{-B} - \frac{1-t}{2}e^B, & t > \frac{e^{2B} - 1}{e^{2B} + 1}. \end{cases}$$

By Definition 3, for any $\epsilon \geq 0$, the \mathcal{H}_{lin} -estimation error transformation of the exponential loss is as follows:

$$\mathcal{T}_{\Phi_{\text{exp}}} = \begin{cases} \mathcal{T}(t), & t \in [\epsilon, 1], \\ \frac{\mathcal{T}(\epsilon)}{\epsilon} t, & t \in [0, \epsilon]. \end{cases}$$

Therefore, when $\epsilon = 0$, $\mathcal{T}_{\Phi_{\text{exp}}}$ is convex, non-decreasing, invertible and satisfies that $\mathcal{T}_{\Phi_{\text{exp}}}(0) = 0$. By Theorem 4, we can choose $\Psi(t) = \mathcal{T}_{\Phi_{\text{exp}}}(t)$ in Theorem 3, or equivalently $\Gamma(t) = \mathcal{T}_{\Phi_{\text{exp}}}^{-1}(t)$ in Theorem 12, which are optimal. To simplify the expression, using the fact that

$$1 - \sqrt{1 - t^2} \geq \frac{t^2}{2},$$

$$1 - \frac{t+1}{2}e^{-B} - \frac{1-t}{2}e^B = 1 - 1/2e^B - 1/2e^{-B} + \frac{e^B - e^{-B}}{2}t,$$

$\mathcal{T}_{\Phi_{\text{exp}}}$ can be lower bounded by

$$\tilde{\mathcal{T}}_{\Phi_{\text{exp}}}(t) = \begin{cases} \frac{t^2}{2}, & t \leq \frac{e^{2B}-1}{e^{2B+1}}, \\ \frac{1}{2} \left(\frac{e^{2B}-1}{e^{2B+1}} \right) t, & t > \frac{e^{2B}-1}{e^{2B+1}}. \end{cases}$$

Thus, we adopt an upper bound of $\mathcal{T}_{\Phi_{\text{exp}}}^{-1}$ as follows:

$$\tilde{\mathcal{T}}_{\Phi_{\text{exp}}}^{-1}(t) = \begin{cases} \sqrt{2t}, & t \leq \frac{1}{2} \left(\frac{e^{2B}-1}{e^{2B+1}} \right)^2, \\ 2 \left(\frac{e^{2B}+1}{e^{2B-1}} \right) t, & t > \frac{1}{2} \left(\frac{e^{2B}-1}{e^{2B+1}} \right)^2. \end{cases}$$

Therefore, by Theorem 3 or Theorem 12, setting $\epsilon = 0$ yields the \mathcal{H}_{lin} -consistency bound for the exponential loss, valid for all $h \in \mathcal{H}_{\text{lin}}$:

$$\begin{aligned} & \mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{lin}}}^* + \mathcal{M}_{\ell_{0-1}, \mathcal{H}_{\text{lin}}} \\ & \leq \begin{cases} \sqrt{2} \left(\mathcal{R}_{\Phi_{\text{exp}}}(h) - \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}}}^* + \mathcal{M}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}}} \right)^{\frac{1}{2}}, & \text{if } \mathcal{R}_{\Phi_{\text{exp}}}(h) - \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}}}^* \leq \frac{1}{2} \left(\frac{e^{2B}-1}{e^{2B+1}} \right)^2 - \mathcal{M}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}}}, \\ 2 \left(\frac{e^{2B}+1}{e^{2B-1}} \right) \left(\mathcal{R}_{\Phi_{\text{exp}}}(h) - \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}}}^* + \mathcal{M}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}}} \right), & \text{otherwise.} \end{cases} \quad (30) \end{aligned}$$

Since the $(\ell_{0-1}, \mathcal{H}_{\text{lin}})$ -minimizability gap coincides with the $(\ell_{0-1}, \mathcal{H}_{\text{lin}})$ -approximation error and $(\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}})$ -minimizability gap coincides with the $(\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}})$ -approximation error for $B = +\infty$, the inequality can be rewritten as follows:

$$\begin{aligned} & \mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^* \\ & \leq \begin{cases} \sqrt{2} \left[\mathcal{R}_{\Phi_{\text{exp}}}(h) - \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{all}}}^* \right]^{\frac{1}{2}} & \text{if } B = +\infty, \\ \begin{cases} \sqrt{2} \left[\mathcal{R}_{\Phi_{\text{exp}}}(h) - \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}}}^* + \mathcal{M}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}}} \right]^{\frac{1}{2}} & \text{if } \mathcal{R}_{\Phi_{\text{exp}}}(h) - \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}}}^* \leq \frac{1}{2} \left(\frac{e^{2B}-1}{e^{2B+1}} \right)^2 - \mathcal{M}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}}}, \\ 2 \left(\frac{e^{2B}+1}{e^{2B-1}} \right) \left(\mathcal{R}_{\Phi_{\text{exp}}}(h) - \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}}}^* + \mathcal{M}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}}} \right) & \text{otherwise.} \end{cases} & \text{otherwise.} \end{cases} \end{aligned}$$

where the $(\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}})$ -minimizability gap $\mathcal{M}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}}}$ is characterized as below, which is less than the $(\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}})$ -approximation error when $B < +\infty$:

$$\begin{aligned} \mathcal{M}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}}} &= \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X \left[2\sqrt{\eta(x)(1-\eta(x))} \mathbb{1}_{1/2 \log \left| \frac{\eta(x)}{1-\eta(x)} \right| \leq W\|x\|_p + B} \right] \\ & \quad - \mathbb{E}_X \left[\max\{\eta(x), 1-\eta(x)\} e^{-(W\|x\|_p + B)} \mathbb{1}_{1/2 \log \left| \frac{\eta(x)}{1-\eta(x)} \right| > W\|x\|_p + B} \right] \\ & \quad - \mathbb{E}_X \left[\min\{\eta(x), 1-\eta(x)\} e^{W\|x\|_p + B} \mathbb{1}_{1/2 \log \left| \frac{\eta(x)}{1-\eta(x)} \right| > W\|x\|_p + B} \right] \\ & < \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X \left[2\sqrt{\eta(x)(1-\eta(x))} \right] \\ & = \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{lin}}}^* - \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{all}}}^*. \end{aligned}$$

Therefore, the inequality for $B = +\infty$ coincides with the consistency excess error bound known for the exponential loss (Zhang, 2004a; Mohri et al., 2018) but the one for $B < +\infty$ is distinct and novel.

K.1.4. QUADRATIC LOSS

For the quadratic loss $\Phi_{\text{quad}}(\alpha) := (1 - \alpha)^2 \mathbb{1}_{\alpha \leq 1}$, for all $h \in \mathcal{H}_{\text{lin}}$ and $x \in \mathcal{X}$:

$$\begin{aligned} \mathcal{C}_{\Phi_{\text{quad}}}(h, x, t) &= t\Phi_{\text{quad}}(h(x)) + (1-t)\Phi_{\text{quad}}(-h(x)) \\ &= t(1-h(x))^2 \mathbb{1}_{h(x) \leq 1} + (1-t)(1+h(x))^2 \mathbb{1}_{h(x) \geq -1}. \end{aligned}$$

$$\inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\Phi_{\text{quad}}}(h, x, t) = \begin{cases} 4t(1-t), & |2t-1| \leq W\|x\|_p + B, \\ \max\{t, 1-t\}(1 - (W\|x\|_p + B))^2 + \min\{t, 1-t\}(1 + W\|x\|_p + B)^2, & |2t-1| > W\|x\|_p + B. \end{cases}$$

Therefore, the $(\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}})$ -minimizability gap can be expressed as follows:

$$\begin{aligned} \mathcal{M}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}}} &= \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X \left[4\eta(x)(1-\eta(x)) \mathbb{1}_{|2\eta(x)-1| \leq W\|x\|_p + B} \right] \\ &\quad - \mathbb{E}_X \left[\max\{\eta(x), 1-\eta(x)\} (1 - (W\|x\|_p + B))^2 \mathbb{1}_{|2\eta(x)-1| > W\|x\|_p + B} \right] \\ &\quad - \mathbb{E}_X \left[\min\{\eta(x), 1-\eta(x)\} (1 + (W\|x\|_p + B))^2 \mathbb{1}_{|2\eta(x)-1| > W\|x\|_p + B} \right] \end{aligned} \quad (31)$$

Note $(\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}})$ -minimizability gap coincides with the $(\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}})$ -approximation error $\mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X [4\eta(x)(1-\eta(x))]$ for $B \geq 1$.

For $\frac{1}{2} < t \leq 1$, we have

$$\begin{aligned} \inf_{h \in \mathcal{H}_{\text{lin}}: h(x) < 0} \mathcal{C}_{\Phi_{\text{quad}}}(h, x, t) &= t + (1-t) \\ &= 1 \end{aligned}$$

$$\begin{aligned} \inf_{x \in \mathcal{X}} \inf_{h \in \mathcal{H}_{\text{lin}}: h(x) < 0} \Delta \mathcal{C}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}}}(h, x, t) &= \inf_{x \in \mathcal{X}} \left(\inf_{h \in \mathcal{H}_{\text{lin}}: h(x) < 0} \mathcal{C}_{\Phi_{\text{quad}}}(h, x, t) - \inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\Phi_{\text{quad}}}(h, x, t) \right) \\ &= \inf_{x \in \mathcal{X}} \begin{cases} 1 - 4t(1-t), & 2t-1 \leq W\|x\|_p + B, \\ 1 - t(1 - (W\|x\|_p + B))^2 - (1-t)(1 + W\|x\|_p + B)^2, & 2t-1 > W\|x\|_p + B. \end{cases} \\ &= \begin{cases} 1 - 4t(1-t), & 2t-1 \leq B, \\ 1 - t(1-B)^2 - (1-t)(1+B)^2, & 2t-1 > B. \end{cases} \\ &= \mathcal{T}(2t-1) \end{aligned}$$

where \mathcal{T} is the increasing and convex function on $[0, 1]$ defined by

$$\forall t \in [0, 1], \quad \mathcal{T}(t) = \begin{cases} t^2, & t \leq B, \\ 2Bt - B^2, & t > B. \end{cases}$$

By Definition 3, for any $\epsilon \geq 0$, the \mathcal{H}_{lin} -estimation error transformation of the quadratic loss is as follows:

$$\mathcal{T}_{\Phi_{\text{quad}}} = \begin{cases} \mathcal{T}(t), & t \in [\epsilon, 1], \\ \frac{\mathcal{T}(\epsilon)}{\epsilon} t, & t \in [0, \epsilon]. \end{cases}$$

Therefore, when $\epsilon = 0$, $\mathcal{T}_{\Phi_{\text{quad}}}$ is convex, non-decreasing, invertible and satisfies that $\mathcal{T}_{\Phi_{\text{quad}}}(0) = 0$. By Theorem 4, we can choose $\Psi(t) = \mathcal{T}_{\Phi_{\text{quad}}}(t)$ in Theorem 3, or equivalently $\Gamma(t) = \mathcal{T}_{\Phi_{\text{quad}}}^{-1}(t) = \begin{cases} \sqrt{t}, & t \leq B^2 \\ \frac{t}{2B} + \frac{B}{2}, & t > B^2 \end{cases}$, in Theorem 12, which are optimal. Thus, by Theorem 3 or Theorem 12, setting $\epsilon = 0$ yields the \mathcal{H}_{lin} -consistency bound for the quadratic loss, valid for all $h \in \mathcal{H}_{\text{lin}}$:

$$\begin{aligned} \mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{lin}}}^* + \mathcal{M}_{\ell_{0-1}, \mathcal{H}_{\text{lin}}} &\leq \begin{cases} \left[\mathcal{R}_{\Phi_{\text{quad}}}(h) - \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}}}^* + \mathcal{M}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}}} \right]^{\frac{1}{2}} & \text{if } \mathcal{R}_{\Phi_{\text{quad}}}(h) - \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}}}^* \leq B^2 - \mathcal{M}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}}} \\ \frac{\mathcal{R}_{\Phi_{\text{quad}}}(h) - \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}}}^* + \mathcal{M}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}}}}{2B} + \frac{B}{2} & \text{otherwise} \end{cases} \end{aligned} \quad (32)$$

Since the $(\ell_{0-1}, \mathcal{H}_{\text{lin}})$ -minimizability gap coincides with the $(\ell_{0-1}, \mathcal{H}_{\text{lin}})$ -approximation error and $(\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}})$ -minimizability gap coincides with the $(\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}})$ -approximation error for $B \geq 1$, the inequality can be rewritten as follows:

$$\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^* \leq \begin{cases} \left[\mathcal{R}_{\Phi_{\text{quad}}}(h) - \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{all}}}^* \right]^{\frac{1}{2}} & \text{if } B \geq 1 \\ \begin{cases} \left[\mathcal{R}_{\Phi_{\text{quad}}}(h) - \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}}}^* + \mathcal{M}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}}} \right]^{\frac{1}{2}} & \text{if } \mathcal{R}_{\Phi_{\text{quad}}}(h) - \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}}}^* \leq B^2 - \mathcal{M}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}}} \\ \frac{\mathcal{R}_{\Phi_{\text{quad}}}(h) - \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}}}^* + \mathcal{M}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}}}}{2B} + \frac{B}{2} & \text{otherwise} \end{cases} & \text{otherwise} \end{cases}$$

where the $(\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}})$ -minimizability gap $\mathcal{M}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}}}$ is characterized as below, which is less than the $(\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}})$ -approximation error when $B < 1$:

$$\begin{aligned} \mathcal{M}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}}} &= \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X \left[4\eta(x)(1 - \eta(x)) \mathbb{1}_{|2\eta(x)-1| \leq W\|x\|_p + B} \right] \\ &\quad - \mathbb{E}_X \left[\max\{\eta(x), 1 - \eta(x)\} (1 - (W\|x\|_p + B))^2 \mathbb{1}_{|2\eta(x)-1| > W\|x\|_p + B} \right] \\ &\quad - \mathbb{E}_X \left[\min\{\eta(x), 1 - \eta(x)\} (1 + (W\|x\|_p + B))^2 \mathbb{1}_{|2\eta(x)-1| > W\|x\|_p + B} \right] \\ &\leq \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X [4\eta(x)(1 - \eta(x))] \\ &= \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{lin}}}^* - \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{all}}}^*. \end{aligned}$$

Therefore, the inequality for $B \geq 1$ coincides with the consistency excess error bound known for the quadratic loss (Zhang, 2004a; Bartlett et al., 2006) but the one for $B < 1$ is distinct and novel.

K.1.5. SIGMOID LOSS

For the sigmoid loss $\Phi_{\text{sig}}(\alpha) := 1 - \tanh(k\alpha)$, $k > 0$, for all $h \in \mathcal{H}_{\text{lin}}$ and $x \in \mathcal{X}$:

$$\begin{aligned} \mathcal{C}_{\Phi_{\text{sig}}}(h, x, t) &= t\Phi_{\text{sig}}(h(x)) + (1 - t)\Phi_{\text{sig}}(-h(x)), \\ &= t(1 - \tanh(kh(x))) + (1 - t)(1 + \tanh(kh(x))). \\ \inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\Phi_{\text{sig}}}(h, x, t) &= 1 - |1 - 2t| \tanh(k(W\|x\|_p + B)) \end{aligned}$$

Therefore, the $(\Phi_{\text{sig}}, \mathcal{H}_{\text{lin}})$ -minimizability gap can be expressed as follows:

$$\begin{aligned} \mathcal{M}_{\Phi_{\text{sig}}, \mathcal{H}_{\text{lin}}} &= \mathcal{R}_{\Phi_{\text{sig}}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X \left[\inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\Phi_{\text{sig}}}(h, x, \eta(x)) \right] \\ &= \mathcal{R}_{\Phi_{\text{sig}}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X [1 - |1 - 2\eta(x)| \tanh(k(W\|x\|_p + B))]. \end{aligned} \tag{33}$$

Note $(\Phi_{\text{sig}}, \mathcal{H}_{\text{lin}})$ -minimizability gap coincides with the $(\Phi_{\text{sig}}, \mathcal{H}_{\text{lin}})$ -approximation error $\mathcal{R}_{\Phi_{\text{sig}}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X [1 - |1 - 2\eta(x)|]$ for $B = +\infty$.

For $\frac{1}{2} < t \leq 1$, we have

$$\begin{aligned} \inf_{h \in \mathcal{H}_{\text{lin}}: h(x) < 0} \mathcal{C}_{\Phi_{\text{sig}}}(h, x, t) &= 1 - |1 - 2t| \tanh(0) \\ &= 1. \\ \inf_{x \in \mathcal{X}} \inf_{h \in \mathcal{H}_{\text{lin}}: h(x) < 0} \Delta \mathcal{C}_{\Phi_{\text{sig}}, \mathcal{H}_{\text{lin}}}(h, x, t) &= \inf_{x \in \mathcal{X}} \left(\inf_{h \in \mathcal{H}_{\text{lin}}: h(x) < 0} \mathcal{C}_{\Phi_{\text{sig}}}(h, x, t) - \inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\Phi_{\text{sig}}}(h, x, t) \right) \\ &= \inf_{x \in \mathcal{X}} (2t - 1) \tanh(k(W\|x\|_p + B)) \\ &= (2t - 1) \tanh(kB) \\ &= \mathcal{J}(2t - 1) \end{aligned}$$

where \mathcal{T} is the increasing and convex function on $[0, 1]$ defined by

$$\forall t \in [0, 1], \mathcal{T}(t) = \tanh(kB)t.$$

By Definition 3, for any $\epsilon \geq 0$, the \mathcal{H}_{lin} -estimation error transformation of the sigmoid loss is as follows:

$$\mathcal{T}_{\Phi_{\text{sig}}} = \tanh(kB)t, \quad t \in [0, 1],$$

Therefore, $\mathcal{T}_{\Phi_{\text{sig}}}$ is convex, non-decreasing, invertible and satisfies that $\mathcal{T}_{\Phi_{\text{sig}}}(0) = 0$. By Theorem 4, we can choose $\Psi(t) = \tanh(kB)t$ in Theorem 3, or equivalently $\Gamma(t) = \frac{t}{\tanh(kB)}$ in Theorem 12, which are optimal when $\epsilon = 0$. Thus, by Theorem 3 or Theorem 12, setting $\epsilon = 0$ yields the \mathcal{H}_{lin} -consistency bound for the sigmoid loss, valid for all $h \in \mathcal{H}_{\text{lin}}$:

$$\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{lin}}}^* \leq \frac{\mathcal{R}_{\Phi_{\text{sig}}}(h) - \mathcal{R}_{\Phi_{\text{sig}}, \mathcal{H}_{\text{lin}}}^* + \mathcal{M}_{\Phi_{\text{sig}}, \mathcal{H}_{\text{lin}}}}{\tanh(kB)} - \mathcal{M}_{\ell_{0-1}, \mathcal{H}_{\text{lin}}}. \quad (34)$$

Since the $(\ell_{0-1}, \mathcal{H}_{\text{lin}})$ -minimizability gap coincides with the $(\ell_{0-1}, \mathcal{H}_{\text{lin}})$ -approximation error and $(\Phi_{\text{sig}}, \mathcal{H}_{\text{lin}})$ -minimizability gap coincides with the $(\Phi_{\text{sig}}, \mathcal{H}_{\text{lin}})$ -approximation error for $B = +\infty$, the inequality can be rewritten as follows:

$$\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^* \leq \begin{cases} \mathcal{R}_{\Phi_{\text{sig}}}(h) - \mathcal{R}_{\Phi_{\text{sig}}, \mathcal{H}_{\text{all}}}^* & \text{if } B = +\infty \\ \frac{1}{\tanh(kB)} \left[\mathcal{R}_{\Phi_{\text{sig}}}(h) - \mathbb{E}_X[1 - |1 - 2\eta(x)| \tanh(k(W\|x\|_p + B))] \right] & \text{otherwise.} \end{cases} \quad (35)$$

The inequality for $B = +\infty$ coincides with the consistency excess error bound known for the sigmoid loss (Zhang, 2004a; Bartlett et al., 2006; Mohri et al., 2018) but the one for $B < +\infty$ is distinct and novel. For $B < +\infty$, we have

$$\mathbb{E}_X[1 - |1 - 2\eta(x)| \tanh(k(W\|x\|_p + B))] > \mathbb{E}_X[1 - |2\eta(x) - 1|] = 2\mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}] = \mathcal{R}_{\Phi_{\text{hinge}}, \mathcal{H}_{\text{all}}}^*.$$

Therefore for $B < +\infty$,

$$\mathcal{R}_{\Phi_{\text{sig}}}(h) - \mathbb{E}_X[1 - |1 - 2\eta(x)| \tanh(k(W\|x\|_p + B))] < \mathcal{R}_{\Phi_{\text{sig}}}(h) - \mathcal{R}_{\Phi_{\text{sig}}, \mathcal{H}_{\text{all}}}^*.$$

Note that: $\mathcal{R}_{\Phi_{\text{sig}}, \mathcal{H}_{\text{all}}}^* = 2\mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^* = 2\mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}]$. Thus, the first inequality (case $B = +\infty$) can be equivalently written as follows:

$$\forall h \in \mathcal{H}_{\text{lin}}, \quad \mathcal{R}_{\ell_{0-1}}(h) \leq \mathcal{R}_{\Phi_{\text{sig}}}(h) - \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}],$$

which is a more informative upper bound than the standard inequality $\mathcal{R}_{\ell_{0-1}}(h) \leq \mathcal{R}_{\Phi_{\text{sig}}}(h)$.

K.1.6. ρ -MARGIN LOSS

For the ρ -margin loss $\Phi_\rho(\alpha) := \min\left\{1, \max\left\{0, 1 - \frac{\alpha}{\rho}\right\}\right\}$, $\rho > 0$, for all $h \in \mathcal{H}_{\text{lin}}$ and $x \in \mathcal{X}$:

$$\begin{aligned} \mathcal{C}_{\Phi_\rho}(h, x, t) &= t\Phi_\rho(h(x)) + (1-t)\Phi_\rho(-h(x)), \\ &= t \min\left\{1, \max\left\{0, 1 - \frac{h(x)}{\rho}\right\}\right\} + (1-t) \min\left\{1, \max\left\{0, 1 + \frac{h(x)}{\rho}\right\}\right\}. \\ \inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\Phi_\rho}(h, x, t) &= \min\{t, 1-t\} + \max\{t, 1-t\} \left(1 - \frac{\min\{W\|x\|_p + B, \rho\}}{\rho}\right). \end{aligned}$$

Therefore, the $(\Phi_\rho, \mathcal{H}_{\text{lin}})$ -minimizability gap can be expressed as follows:

$$\begin{aligned} \mathcal{M}_{\Phi_\rho, \mathcal{H}_{\text{lin}}} &= \mathcal{R}_{\Phi_\rho, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X \left[\inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\Phi_\rho}(h, x, \eta(x)) \right] \\ &= \mathcal{R}_{\Phi_\rho, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X \left[\min\{\eta(x), 1 - \eta(x)\} + \max\{\eta(x), 1 - \eta(x)\} \left(1 - \frac{\min\{W\|x\|_p + B, \rho\}}{\rho}\right) \right]. \end{aligned} \quad (36)$$

Note the $(\Phi_\rho, \mathcal{H}_{\text{lin}})$ -minimizability gap coincides with the $(\Phi_\rho, \mathcal{H}_{\text{lin}})$ -approximation error $\mathcal{R}_{\Phi_\rho, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}]$ for $B \geq \rho$.

For $\frac{1}{2} < t \leq 1$, we have

$$\begin{aligned} \inf_{h \in \mathcal{H}_{\text{lin}}: h(x) < 0} \mathcal{C}_{\Phi_\rho}(h, x, t) &= t + (1-t) \left(1 - \frac{\min\{W\|x\|_p + B, \rho\}}{\rho} \right). \\ \inf_{x \in \mathcal{X}} \inf_{h \in \mathcal{H}_{\text{lin}}: h(x) < 0} \Delta \mathcal{C}_{\Phi_\rho, \mathcal{H}_{\text{lin}}}(h, x) &= \inf_{x \in \mathcal{X}} \left(\inf_{h \in \mathcal{H}_{\text{lin}}: h(x) < 0} \mathcal{C}_{\Phi_\rho}(h, x, t) - \inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\Phi_\rho}(h, x, t) \right) \\ &= \inf_{x \in \mathcal{X}} (2t-1) \frac{\min\{W\|x\|_p + B, \rho\}}{\rho} \\ &= (2t-1) \frac{\min\{B, \rho\}}{\rho} \\ &= \mathcal{J}(2t-1) \end{aligned}$$

where \mathcal{J} is the increasing and convex function on $[0, 1]$ defined by

$$\forall t \in [0, 1], \mathcal{J}(t) = \frac{\min\{B, \rho\}}{\rho} t.$$

By Definition 3, for any $\epsilon \geq 0$, the \mathcal{H}_{lin} -estimation error transformation of the ρ -margin loss is as follows:

$$\mathcal{T}_{\Phi_\rho} = \frac{\min\{B, \rho\}}{\rho} t, \quad t \in [0, 1],$$

Therefore, \mathcal{T}_{Φ_ρ} is convex, non-decreasing, invertible and satisfies that $\mathcal{T}_{\Phi_\rho}(0) = 0$. By Theorem 4, we can choose $\Psi(t) = \frac{\min\{B, \rho\}}{\rho} t$ in Theorem 3, or equivalently $\Gamma(t) = \frac{\rho}{\min\{B, \rho\}} t$ in Theorem 12, which are optimal when $\epsilon = 0$. Thus, by Theorem 3 or Theorem 12, setting $\epsilon = 0$ yields the \mathcal{H}_{lin} -consistency bound for the ρ -margin loss, valid for all $h \in \mathcal{H}_{\text{lin}}$:

$$\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{lin}}}^* \leq \frac{\rho \left(\mathcal{R}_{\Phi_\rho}(h) - \mathcal{R}_{\Phi_\rho, \mathcal{H}_{\text{lin}}}^* + \mathcal{M}_{\Phi_\rho, \mathcal{H}_{\text{lin}}} \right)}{\min\{B, \rho\}} - \mathcal{M}_{\ell_{0-1}, \mathcal{H}_{\text{lin}}}. \quad (37)$$

Since the $(\ell_{0-1}, \mathcal{H}_{\text{lin}})$ -minimizability gap coincides with the $(\ell_{0-1}, \mathcal{H}_{\text{lin}})$ -approximation error and $(\Phi_\rho, \mathcal{H}_{\text{lin}})$ -minimizability gap coincides with the $(\Phi_\rho, \mathcal{H}_{\text{lin}})$ -approximation error for $B \geq \rho$, the inequality can be rewritten as follows:

$$\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^* \leq \begin{cases} \mathcal{R}_{\Phi_\rho}(h) - \mathcal{R}_{\Phi_\rho, \mathcal{H}_{\text{all}}}^* & \text{if } B \geq \rho \\ \rho \left(\mathcal{R}_{\Phi_\rho}(h) - \mathbb{E}_X \left[\min\{\eta(x), 1 - \eta(x)\} + \max\{\eta(x), 1 - \eta(x)\} \left(1 - \frac{\min\{W\|x\|_p + B, \rho\}}{\rho} \right) \right] \right) & \text{otherwise.} \end{cases}$$

Note that: $\mathcal{R}_{\Phi_\rho, \mathcal{H}_{\text{all}}}^* = \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^* = \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}]$. Thus, the first inequality (case $B \geq \rho$) can be equivalently written as follows:

$$\forall h \in \mathcal{H}_{\text{lin}}, \quad \mathcal{R}_{\ell_{0-1}}(h) \leq \mathcal{R}_{\Phi_\rho}(h). \quad (38)$$

The case $B \geq \rho$ is one of the ‘‘trivial cases’’ mentioned in Section 4, where the trivial inequality $\mathcal{R}_{\ell_{0-1}}(h) \leq \mathcal{R}_{\Phi_\rho}(h)$ can be obtained directly using the fact that ℓ_{0-1} is upper bounded by Φ_ρ . This, however, does not imply that non-adversarial \mathcal{H}_{lin} -consistency bound for the ρ -margin loss is trivial when $B > \rho$ since it is optimal.

K.2. One-Hidden-Layer ReLU Neural Network

As with the linear case, \mathcal{H}_{NN} also satisfies the condition of Lemma 1 and thus the $(\ell_{0-1}, \mathcal{H}_{\text{NN}})$ -minimizability gap coincides with the $(\ell_{0-1}, \mathcal{H}_{\text{NN}})$ -approximation error:

$$\begin{aligned} \mathcal{M}_{\ell_{0-1}, \mathcal{H}_{\text{NN}}} &= \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}] \\ &= \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{NN}}}^* - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^*. \end{aligned} \quad (39)$$

By the definition of \mathcal{H}_{NN} , for any $x \in \mathcal{X}$,

$$\{h(x) \mid h \in \mathcal{H}_{\text{NN}}\} = [-\Lambda(W\|x\|_p + B), \Lambda(W\|x\|_p + B)].$$

K.2.1. HINGE LOSS

For the hinge loss $\Phi_{\text{hinge}}(\alpha) := \max\{0, 1 - \alpha\}$, for all $h \in \mathcal{H}_{\text{NN}}$ and $x \in \mathcal{X}$:

$$\begin{aligned} \mathcal{C}_{\Phi_{\text{hinge}}}(h, x, t) &= t\Phi_{\text{hinge}}(h(x)) + (1-t)\Phi_{\text{hinge}}(-h(x)) \\ &= t \max\{0, 1 - h(x)\} + (1-t) \max\{0, 1 + h(x)\}. \\ \inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\Phi_{\text{hinge}}}(h, x, t) &= 1 - |2t - 1| \min\{\Lambda W \|x\|_p + \Lambda B, 1\}. \end{aligned}$$

Therefore, the $(\Phi_{\text{hinge}}, \mathcal{H}_{\text{NN}})$ -minimizability gap can be expressed as follows:

$$\begin{aligned} \mathcal{M}_{\Phi_{\text{hinge}}, \mathcal{H}_{\text{NN}}} &= \mathcal{R}_{\Phi_{\text{hinge}}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X \left[1 - \inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\Phi_{\text{hinge}}}(h, x, \eta(x)) \right]. \\ &= \mathcal{R}_{\Phi_{\text{hinge}}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X [1 - |2\eta(x) - 1| \min\{\Lambda W \|x\|_p + \Lambda B, 1\}]. \end{aligned} \quad (40)$$

Note the $(\Phi_{\text{hinge}}, \mathcal{H}_{\text{NN}})$ -minimizability gap coincides with the $(\Phi_{\text{hinge}}, \mathcal{H}_{\text{NN}})$ -approximation error $\mathcal{R}_{\Phi_{\text{hinge}}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X [1 - |2\eta(x) - 1|]$ for $\Lambda B \geq 1$.

For $\frac{1}{2} < t \leq 1$, we have

$$\begin{aligned} \inf_{h \in \mathcal{H}_{\text{NN}}: h(x) < 0} \mathcal{C}_{\Phi_{\text{hinge}}}(h, x, t) &= t \max\{0, 1 - 0\} + (1-t) \max\{0, 1 + 0\} \\ &= 1. \\ \inf_{x \in \mathcal{X}} \inf_{h \in \mathcal{H}_{\text{NN}}: h(x) < 0} \Delta \mathcal{C}_{\Phi_{\text{hinge}}, \mathcal{H}_{\text{NN}}}(h, x, t) &= \inf_{x \in \mathcal{X}} \left\{ \inf_{h \in \mathcal{H}_{\text{NN}}: h(x) < 0} \mathcal{C}_{\Phi_{\text{hinge}}}(h, x, t) - \inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\Phi_{\text{hinge}}}(h, x, t) \right\} \\ &= \inf_{x \in \mathcal{X}} (2t - 1) \min\{\Lambda W \|x\|_p + \Lambda B, 1\} \\ &= (2t - 1) \min\{\Lambda B, 1\} \\ &= \mathcal{T}(2t - 1), \end{aligned}$$

where \mathcal{T} is the increasing and convex function on $[0, 1]$ defined by

$$\forall t \in [0, 1], \mathcal{T}(t) = \min\{\Lambda B, 1\} t.$$

By Definition 3, for any $\epsilon \geq 0$, the \mathcal{H}_{NN} -estimation error transformation of the hinge loss is as follows:

$$\mathcal{T}_{\Phi_{\text{hinge}}} = \min\{\Lambda B, 1\} t, \quad t \in [0, 1],$$

Therefore, $\mathcal{T}_{\Phi_{\text{hinge}}}$ is convex, non-decreasing, invertible and satisfies that $\mathcal{T}_{\Phi_{\text{hinge}}}(0) = 0$. By Theorem 4, we can choose $\Psi(t) = \min\{\Lambda B, 1\} t$ in Theorem 3, or, equivalently, $\Gamma(t) = \frac{t}{\min\{\Lambda B, 1\}}$ in Theorem 12, which are optimal when $\epsilon = 0$. Thus, by Theorem 3 or Theorem 12, setting $\epsilon = 0$ yields the \mathcal{H}_{NN} -consistency bound for the hinge loss, valid for all $h \in \mathcal{H}_{\text{NN}}$:

$$\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{NN}}}^* \leq \frac{\mathcal{R}_{\Phi_{\text{hinge}}}(h) - \mathcal{R}_{\Phi_{\text{hinge}}, \mathcal{H}_{\text{NN}}}^* + \mathcal{M}_{\Phi_{\text{hinge}}, \mathcal{H}_{\text{NN}}}}{\min\{\Lambda B, 1\}} - \mathcal{M}_{\ell_{0-1}, \mathcal{H}_{\text{NN}}}. \quad (41)$$

Since the $(\ell_{0-1}, \mathcal{H}_{\text{NN}})$ -minimizability gap coincides with the $(\ell_{0-1}, \mathcal{H}_{\text{NN}})$ -approximation error and $(\Phi_{\text{hinge}}, \mathcal{H}_{\text{NN}})$ -minimizability gap coincides with the $(\Phi_{\text{hinge}}, \mathcal{H}_{\text{NN}})$ -approximation error for $\Lambda B \geq 1$, the inequality can be rewritten as follows:

$$\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^* \leq \begin{cases} \mathcal{R}_{\Phi_{\text{hinge}}}(h) - \mathcal{R}_{\Phi_{\text{hinge}}, \mathcal{H}_{\text{all}}}^* & \text{if } \Lambda B \geq 1 \\ \frac{1}{\Lambda B} \left[\mathcal{R}_{\Phi_{\text{hinge}}}(h) - \mathbb{E}_X [1 - |2\eta(x) - 1| \min\{\Lambda W \|x\|_p + \Lambda B, 1\}] \right] & \text{otherwise.} \end{cases}$$

The inequality for $\Lambda B \geq 1$ coincides with the consistency excess error bound known for the hinge loss (Zhang, 2004a; Bartlett et al., 2006; Mohri et al., 2018) but the one for $\Lambda B < 1$ is distinct and novel. For $\Lambda B < 1$, we have

$$\mathbb{E}_X [1 - |2\eta(x) - 1| \min\{\Lambda W \|x\|_p + \Lambda B, 1\}] > \mathbb{E}_X [1 - |2\eta(x) - 1|] = 2\mathbb{E}_X [\min\{\eta(x), 1 - \eta(x)\}] = \mathcal{R}_{\Phi_{\text{hinge}}, \mathcal{H}_{\text{all}}}^*.$$

Therefore for $\Lambda B < 1$,

$$\mathcal{R}_{\Phi_{\text{hinge}}}(h) - \mathbb{E}_X[1 - |2\eta(x) - 1| \min\{\Lambda W \|x\|_p + \Lambda B, 1\}] < \mathcal{R}_{\Phi_{\text{hinge}}}(h) - \mathcal{R}_{\Phi_{\text{hinge}}, \mathcal{H}_{\text{all}}}^*.$$

Note that: $\mathcal{R}_{\Phi_{\text{hinge}}, \mathcal{H}_{\text{all}}}^* = 2\mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^* = 2\mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}]$. Thus, the first inequality (case $\Lambda B \geq 1$) can be equivalently written as follows:

$$\forall h \in \mathcal{H}_{\text{NN}}, \quad \mathcal{R}_{\ell_{0-1}}(h) \leq \mathcal{R}_{\Phi_{\text{hinge}}}(h) - \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}],$$

which is a more informative upper bound than the standard inequality $\mathcal{R}_{\ell_{0-1}}(h) \leq \mathcal{R}_{\Phi_{\text{hinge}}}(h)$.

K.2.2. LOGISTIC LOSS

For the logistic loss $\Phi_{\log}(\alpha) := \log_2(1 + e^{-\alpha})$, for all $h \in \mathcal{H}_{\text{NN}}$ and $x \in \mathcal{X}$:

$$\begin{aligned} \mathcal{C}_{\Phi_{\log}}(h, x, t) &= t\Phi_{\log}(h(x)) + (1-t)\Phi_{\log}(-h(x)), \\ &= t \log_2(1 + e^{-h(x)}) + (1-t) \log_2(1 + e^{h(x)}). \end{aligned}$$

$$\inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\Phi_{\log}}(h, x, t) = \begin{cases} -t \log_2(t) - (1-t) \log_2(1-t) & \text{if } \log\left|\frac{t}{1-t}\right| \leq \Lambda W \|x\|_p + \Lambda B, \\ \max\{t, 1-t\} \log_2(1 + e^{-(\Lambda W \|x\|_p + \Lambda B)}) + \min\{t, 1-t\} \log_2(1 + e^{\Lambda W \|x\|_p + \Lambda B}) & \text{if } \log\left|\frac{t}{1-t}\right| > \Lambda W \|x\|_p + \Lambda B. \end{cases}$$

Therefore, the $(\Phi_{\log}, \mathcal{H}_{\text{NN}})$ -minimizability gap can be expressed as follows:

$$\begin{aligned} \mathcal{M}_{\Phi_{\log}, \mathcal{H}_{\text{NN}}} &= \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X \left[\inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\Phi_{\log}}(h, x, \eta(x)) \right] \\ &= \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X \left[-\eta(x) \log_2(\eta(x)) - (1-\eta(x)) \log_2(1-\eta(x)) \mathbb{1}_{\log\left|\frac{\eta(x)}{1-\eta(x)}\right| \leq \Lambda W \|x\|_p + \Lambda B} \right] \\ &\quad - \mathbb{E}_X \left[\max\{\eta(x), 1-\eta(x)\} \log_2(1 + e^{-(\Lambda W \|x\|_p + \Lambda B)}) \mathbb{1}_{\log\left|\frac{\eta(x)}{1-\eta(x)}\right| > \Lambda W \|x\|_p + \Lambda B} \right] \\ &\quad - \mathbb{E}_X \left[\min\{\eta(x), 1-\eta(x)\} \log_2(1 + e^{\Lambda W \|x\|_p + \Lambda B}) \mathbb{1}_{\log\left|\frac{\eta(x)}{1-\eta(x)}\right| > \Lambda W \|x\|_p + \Lambda B} \right] \end{aligned} \quad (42)$$

Note $(\Phi_{\log}, \mathcal{H}_{\text{NN}})$ -minimizability gap coincides with the $(\Phi_{\log}, \mathcal{H}_{\text{NN}})$ -approximation error $\mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X[-\eta(x) \log_2(\eta(x)) - (1-\eta(x)) \log_2(1-\eta(x))]$ for $\Lambda B = +\infty$.

For $\frac{1}{2} < t \leq 1$, we have

$$\begin{aligned} \inf_{h \in \mathcal{H}_{\text{NN}}: h(x) < 0} \mathcal{C}_{\Phi_{\log}}(h, x, t) &= t \log_2(1 + e^{-0}) + (1-t) \log_2(1 + e^0) \\ &= 1, \\ \inf_{x \in \mathcal{X}} \inf_{h \in \mathcal{H}_{\text{NN}}: h(x) < 0} \Delta \mathcal{C}_{\Phi_{\log}, \mathcal{H}_{\text{NN}}}(h, x, t) &= \inf_{x \in \mathcal{X}} \left(\inf_{h \in \mathcal{H}_{\text{NN}}: h(x) < 0} \mathcal{C}_{\Phi_{\log}}(h, x, t) - \inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\Phi_{\log}}(h, x, t) \right) \\ &= \inf_{x \in \mathcal{X}} \begin{cases} 1 + t \log_2(t) + (1-t) \log_2(1-t) \\ \text{if } \log\left|\frac{t}{1-t}\right| \leq \Lambda W \|x\|_p + \Lambda B, \\ 1 - t \log_2(1 + e^{-(\Lambda W \|x\|_p + \Lambda B)}) - (1-t) \log_2(1 + e^{\Lambda W \|x\|_p + \Lambda B}) \\ \text{if } \log\left|\frac{t}{1-t}\right| > \Lambda W \|x\|_p + \Lambda B. \end{cases} \\ &= \begin{cases} 1 + t \log_2(t) + (1-t) \log_2(1-t) & \text{if } \log\left|\frac{t}{1-t}\right| \leq \Lambda B, \\ 1 - t \log_2(1 + e^{-\Lambda B}) - (1-t) \log_2(1 + e^{\Lambda B}) & \text{if } \log\left|\frac{t}{1-t}\right| > \Lambda B. \end{cases} \\ &= \mathcal{J}(2t - 1), \end{aligned}$$

where \mathcal{J} is the increasing and convex function on $[0, 1]$ defined by

$$\forall t \in [0, 1], \quad \mathcal{J}(t) = \begin{cases} \frac{t+1}{2} \log_2(t+1) + \frac{1-t}{2} \log_2(1-t), & t \leq \frac{e^{\Lambda B} - 1}{e^{\Lambda B} + 1}, \\ 1 - \frac{t+1}{2} \log_2(1 + e^{-\Lambda B}) - \frac{1-t}{2} \log_2(1 + e^{\Lambda B}), & t > \frac{e^{\Lambda B} - 1}{e^{\Lambda B} + 1}. \end{cases}$$

By Definition 3, for any $\epsilon \geq 0$, the \mathcal{H}_{NN} -estimation error transformation of the logistic loss is as follows:

$$\mathcal{J}_{\Phi_{\log}} = \begin{cases} \mathcal{J}(t), & t \in [\epsilon, 1], \\ \frac{\mathcal{J}(\epsilon)}{\epsilon} t, & t \in [0, \epsilon]. \end{cases}$$

Therefore, when $\epsilon = 0$, $\mathcal{J}_{\Phi_{\log}}$ is convex, non-decreasing, invertible and satisfies that $\mathcal{J}_{\Phi_{\log}}(0) = 0$. By Theorem 4, we can choose $\Psi(t) = \mathcal{J}_{\Phi_{\log}}(t)$ in Theorem 3, or equivalently $\Gamma(t) = \mathcal{J}_{\Phi_{\log}}^{-1}(t)$ in Theorem 12, which are optimal. To simplify the expression, using the fact that

$$\begin{aligned} \frac{t+1}{2} \log_2(t+1) + \frac{1-t}{2} \log_2(1-t) &= 1 - \left(-\frac{t+1}{2} \log_2\left(\frac{t+1}{2}\right) - \frac{1-t}{2} \log_2\left(\frac{1-t}{2}\right) \right) \\ &\geq 1 - \sqrt{4 \frac{1-t}{2} \frac{t+1}{2}} \\ &= 1 - \sqrt{1-t^2} \\ &\geq \frac{t^2}{2}, \\ 1 - \frac{t+1}{2} \log_2(1+e^{-\Lambda B}) - \frac{1-t}{2} \log_2(1+e^{\Lambda B}) &= \frac{1}{2} \log_2\left(\frac{4}{2+e^{-\Lambda B}+e^{\Lambda B}}\right) + 1/2 \log_2\left(\frac{1+e^{\Lambda B}}{1+e^{-\Lambda B}}\right)t, \end{aligned}$$

$\mathcal{J}_{\Phi_{\log}}$ can be lower bounded by

$$\tilde{\mathcal{J}}_{\Phi_{\log}}(t) = \begin{cases} \frac{t^2}{2}, & t \leq \frac{e^{\Lambda B}-1}{e^{\Lambda B}+1}, \\ \frac{1}{2} \left(\frac{e^{\Lambda B}-1}{e^{\Lambda B}+1} \right) t, & t > \frac{e^{\Lambda B}-1}{e^{\Lambda B}+1}. \end{cases}$$

Thus, we adopt an upper bound of $\mathcal{J}_{\Phi_{\log}}^{-1}$ as follows:

$$\tilde{\mathcal{J}}_{\Phi_{\log}}^{-1}(t) = \begin{cases} \sqrt{2t}, & t \leq \frac{1}{2} \left(\frac{e^{\Lambda B}-1}{e^{\Lambda B}+1} \right)^2, \\ 2 \left(\frac{e^{\Lambda B}+1}{e^{\Lambda B}-1} \right) t, & t > \frac{1}{2} \left(\frac{e^{\Lambda B}-1}{e^{\Lambda B}+1} \right)^2. \end{cases}$$

Therefore, by Theorem 3 or Theorem 12, setting $\epsilon = 0$ yields the \mathcal{H}_{NN} -consistency bound for the logistic loss, valid for all $h \in \mathcal{H}_{\text{NN}}$:

$$\begin{aligned} \mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{NN}}}^* + \mathcal{M}_{\ell_{0-1}, \mathcal{H}_{\text{NN}}} \\ \leq \begin{cases} \sqrt{2} \left(\mathcal{R}_{\Phi_{\log}}(h) - \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{NN}}}^* + \mathcal{M}_{\Phi_{\log}, \mathcal{H}_{\text{NN}}} \right)^{\frac{1}{2}}, & \text{if } \mathcal{R}_{\Phi_{\log}}(h) - \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{NN}}}^* \leq \frac{1}{2} \left(\frac{e^{\Lambda B}-1}{e^{\Lambda B}+1} \right)^2 - \mathcal{M}_{\Phi_{\log}, \mathcal{H}_{\text{NN}}} \\ 2 \left(\frac{e^{\Lambda B}+1}{e^{\Lambda B}-1} \right) \left(\mathcal{R}_{\Phi_{\log}}(h) - \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{NN}}}^* + \mathcal{M}_{\Phi_{\log}, \mathcal{H}_{\text{NN}}} \right), & \text{otherwise} \end{cases} \quad (43) \end{aligned}$$

Since the $(\ell_{0-1}, \mathcal{H}_{\text{NN}})$ -minimizability gap coincides with the $(\ell_{0-1}, \mathcal{H}_{\text{NN}})$ -approximation error and $(\Phi_{\log}, \mathcal{H}_{\text{NN}})$ -minimizability gap coincides with the $(\Phi_{\log}, \mathcal{H}_{\text{NN}})$ -approximation error for $\Lambda B = +\infty$, the inequality can be rewritten as follows:

$$\begin{aligned} \mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^* \\ \leq \begin{cases} \sqrt{2} \left[\mathcal{R}_{\Phi_{\log}}(h) - \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{all}}}^* \right]^{\frac{1}{2}} & \text{if } \Lambda B = +\infty \\ \begin{cases} \sqrt{2} \left[\mathcal{R}_{\Phi_{\log}}(h) - \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{NN}}}^* + \mathcal{M}_{\Phi_{\log}, \mathcal{H}_{\text{NN}}} \right]^{\frac{1}{2}} & \text{if } \mathcal{R}_{\Phi_{\log}}(h) - \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{NN}}}^* \leq \frac{1}{2} \left(\frac{e^{\Lambda B}-1}{e^{\Lambda B}+1} \right)^2 - \mathcal{M}_{\Phi_{\log}, \mathcal{H}_{\text{NN}}} \\ 2 \left(\frac{e^{\Lambda B}+1}{e^{\Lambda B}-1} \right) \left(\mathcal{R}_{\Phi_{\log}}(h) - \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{NN}}}^* + \mathcal{M}_{\Phi_{\log}, \mathcal{H}_{\text{NN}}} \right) & \text{otherwise} \end{cases} & \text{otherwise} \end{cases} \end{aligned}$$

where the $(\Phi_{\log}, \mathcal{H}_{\text{NN}})$ -minimizability gap $\mathcal{M}_{\Phi_{\log}, \mathcal{H}_{\text{NN}}}$ is characterized as below, which is less than the $(\Phi_{\log}, \mathcal{H}_{\text{NN}})$ -

approximation error when $\Lambda B < +\infty$:

$$\begin{aligned}
 \mathcal{M}_{\Phi_{\log}, \mathcal{H}_{\text{NN}}} &= \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X \left[-\eta(x) \log_2(\eta(x)) - (1 - \eta(x)) \log_2(1 - \eta(x)) \mathbb{1}_{\log_2 \left| \frac{\eta(x)}{1 - \eta(x)} \right| \leq \Lambda W \|x\|_p + \Lambda B} \right] \\
 &\quad - \mathbb{E}_X \left[\max\{\eta(x), 1 - \eta(x)\} \log_2(1 + e^{-(\Lambda W \|x\|_p + \Lambda B)}) \mathbb{1}_{\log_2 \left| \frac{\eta(x)}{1 - \eta(x)} \right| > \Lambda W \|x\|_p + \Lambda B} \right] \\
 &\quad - \left[\min\{\eta(x), 1 - \eta(x)\} \log_2(1 + e^{\Lambda W \|x\|_p + \Lambda B}) \mathbb{1}_{\log_2 \left| \frac{\eta(x)}{1 - \eta(x)} \right| > \Lambda W \|x\|_p + \Lambda B} \right] \\
 &< \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X \left[-\eta(x) \log_2(\eta(x)) - (1 - \eta(x)) \log_2(1 - \eta(x)) \right] \\
 &= \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{NN}}}^* - \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{all}}}^*.
 \end{aligned}$$

Therefore, the inequality for $\Lambda B = +\infty$ coincides with the consistency excess error bound known for the logistic loss (Zhang, 2004a; Mohri et al., 2018) but the one for $\Lambda B < +\infty$ is distinct and novel.

K.2.3. EXPONENTIAL LOSS

For the exponential loss $\Phi_{\text{exp}}(\alpha) := e^{-\alpha}$, for all $h \in \mathcal{H}_{\text{NN}}$ and $x \in \mathcal{X}$:

$$\begin{aligned}
 \mathcal{C}_{\Phi_{\text{exp}}}(h, x, t) &= t\Phi_{\text{exp}}(h(x)) + (1 - t)\Phi_{\text{exp}}(-h(x)) \\
 &= te^{-h(x)} + (1 - t)e^{h(x)}.
 \end{aligned}$$

$$\inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\Phi_{\text{exp}}}(h, x, t) = \begin{cases} 2\sqrt{t(1-t)} & \text{if } 1/2 \log \left| \frac{t}{1-t} \right| \leq \Lambda W \|x\|_p + \Lambda B \\ \max\{t, 1-t\} e^{-(\Lambda W \|x\|_p + \Lambda B)} + \min\{t, 1-t\} e^{\Lambda W \|x\|_p + \Lambda B} & \text{if } 1/2 \log \left| \frac{t}{1-t} \right| > \Lambda W \|x\|_p + \Lambda B. \end{cases}$$

Therefore, the $(\Phi_{\text{exp}}, \mathcal{H}_{\text{NN}})$ -minimizability gap can be expressed as follows:

$$\begin{aligned}
 \mathcal{M}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{NN}}} &= \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X \left[\inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\Phi_{\text{exp}}}(h, x, \eta(x)) \right] \\
 &= \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X \left[2\sqrt{\eta(x)(1 - \eta(x))} \mathbb{1}_{1/2 \log \left| \frac{\eta(x)}{1 - \eta(x)} \right| \leq \Lambda W \|x\|_p + \Lambda B} \right] \\
 &\quad - \mathbb{E}_X \left[\max\{\eta(x), 1 - \eta(x)\} e^{-(\Lambda W \|x\|_p + \Lambda B)} \mathbb{1}_{1/2 \log \left| \frac{\eta(x)}{1 - \eta(x)} \right| > \Lambda W \|x\|_p + \Lambda B} \right] \\
 &\quad - \mathbb{E}_X \left[\min\{\eta(x), 1 - \eta(x)\} e^{\Lambda W \|x\|_p + \Lambda B} \mathbb{1}_{1/2 \log \left| \frac{\eta(x)}{1 - \eta(x)} \right| > \Lambda W \|x\|_p + \Lambda B} \right].
 \end{aligned} \tag{44}$$

Note $(\Phi_{\text{exp}}, \mathcal{H}_{\text{NN}})$ -minimizability gap coincides with the $(\Phi_{\text{exp}}, \mathcal{H}_{\text{NN}})$ -approximation error $\mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X \left[2\sqrt{\eta(x)(1 - \eta(x))} \right]$ for $\Lambda B = +\infty$.

For $\frac{1}{2} < t \leq 1$, we have

$$\begin{aligned}
 \inf_{h \in \mathcal{H}_{\text{NN}}: h(x) < 0} \mathcal{C}_{\Phi_{\text{exp}}}(h, x, t) &= te^{-0} + (1 - t)e^0 \\
 &= 1.
 \end{aligned}$$

$$\begin{aligned}
 \inf_{x \in \mathcal{X}} \inf_{h \in \mathcal{H}_{\text{NN}}: h(x) < 0} \Delta \mathcal{C}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{NN}}}(h, x, t) &= \inf_{x \in \mathcal{X}} \left(\inf_{h \in \mathcal{H}_{\text{NN}}: h(x) < 0} \mathcal{C}_{\Phi_{\text{exp}}}(h, x, t) - \inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\Phi_{\text{exp}}}(h, x, t) \right) \\
 &= \inf_{x \in \mathcal{X}} \begin{cases} 1 - 2\sqrt{t(1-t)} & \text{if } 1/2 \log \frac{t}{1-t} \leq \Lambda W \|x\|_p + \Lambda B, \\ 1 - te^{-(\Lambda W \|x\|_p + \Lambda B)} - (1 - t)e^{\Lambda W \|x\|_p + \Lambda B} & \text{if } 1/2 \log \frac{t}{1-t} > \Lambda W \|x\|_p + \Lambda B. \end{cases} \\
 &= \begin{cases} 1 - 2\sqrt{t(1-t)}, & 1/2 \log \frac{t}{1-t} \leq \Lambda B \\ 1 - te^{-\Lambda B} - (1 - t)e^{\Lambda B}, & 1/2 \log \frac{t}{1-t} > \Lambda B \end{cases} \\
 &= \mathcal{T}(2t - 1),
 \end{aligned}$$

where \mathcal{T} is the increasing and convex function on $[0, 1]$ defined by

$$\forall t \in [0, 1], \quad \mathcal{T}(t) = \begin{cases} 1 - \sqrt{1 - t^2}, & t \leq \frac{e^{2\Lambda B} - 1}{e^{2\Lambda B} + 1}, \\ 1 - \frac{t+1}{2} e^{-\Lambda B} - \frac{1-t}{2} e^{\Lambda B}, & t > \frac{e^{2\Lambda B} - 1}{e^{2\Lambda B} + 1}. \end{cases}$$

By Definition 3, for any $\epsilon \geq 0$, the \mathcal{H}_{lin} -estimation error transformation of the exponential loss is as follows:

$$\mathcal{T}_{\Phi_{\text{exp}}} = \begin{cases} \mathcal{T}(t), & t \in [\epsilon, 1], \\ \frac{\mathcal{T}(\epsilon)}{\epsilon} t, & t \in [0, \epsilon]. \end{cases}$$

Therefore, when $\epsilon = 0$, $\mathcal{T}_{\Phi_{\text{exp}}}$ is convex, non-decreasing, invertible and satisfies that $\mathcal{T}_{\Phi_{\text{exp}}}(0) = 0$. By Theorem 4, we can choose $\Psi(t) = \mathcal{T}_{\Phi_{\text{exp}}}(t)$ in Theorem 3, or equivalently $\Gamma(t) = \mathcal{T}_{\Phi_{\text{exp}}}^{-1}(t)$ in Theorem 12, which are optimal. To simplify the expression, using the fact that

$$1 - \sqrt{1 - t^2} \geq \frac{t^2}{2},$$

$$1 - \frac{t+1}{2}e^{-\Lambda B} - \frac{1-t}{2}e^{\Lambda B} = 1 - 1/2e^{\Lambda B} - 1/2e^{-\Lambda B} + \frac{e^{\Lambda B} - e^{-\Lambda B}}{2}t,$$

$\mathcal{T}_{\Phi_{\text{exp}}}$ can be lower bounded by

$$\tilde{\mathcal{T}}_{\Phi_{\text{exp}}}(t) = \begin{cases} \frac{t^2}{2}, & t \leq \frac{e^{2\Lambda B} - 1}{e^{2\Lambda B} + 1}, \\ \frac{1}{2} \left(\frac{e^{2\Lambda B} - 1}{e^{2\Lambda B} + 1} \right) t, & t > \frac{e^{2\Lambda B} - 1}{e^{2\Lambda B} + 1}. \end{cases}$$

Thus, we adopt an upper bound of $\mathcal{T}_{\Phi_{\text{exp}}}^{-1}$ as follows:

$$\tilde{\mathcal{T}}_{\Phi_{\text{exp}}}^{-1}(t) = \begin{cases} \sqrt{2t}, & t \leq \frac{1}{2} \left(\frac{e^{2\Lambda B} - 1}{e^{2\Lambda B} + 1} \right)^2, \\ 2 \left(\frac{e^{2\Lambda B} + 1}{e^{2\Lambda B} - 1} \right) t, & t > \frac{1}{2} \left(\frac{e^{2\Lambda B} - 1}{e^{2\Lambda B} + 1} \right)^2. \end{cases}$$

Therefore, by Theorem 3 or Theorem 12, setting $\epsilon = 0$ yields the \mathcal{H}_{NN} -consistency bound for the exponential loss, valid for all $h \in \mathcal{H}_{\text{NN}}$:

$$\begin{aligned} & \mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{NN}}}^* + \mathcal{M}_{\ell_{0-1}, \mathcal{H}_{\text{NN}}} \\ & \leq \begin{cases} \sqrt{2} \left(\mathcal{R}_{\Phi_{\text{exp}}}(h) - \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{NN}}}^* + \mathcal{M}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{NN}}} \right)^{\frac{1}{2}}, & \text{if } \mathcal{R}_{\Phi_{\text{exp}}}(h) - \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{NN}}}^* \leq \frac{1}{2} \left(\frac{e^{2\Lambda B} - 1}{e^{2\Lambda B} + 1} \right)^2 - \mathcal{M}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{NN}}}, \\ 2 \left(\frac{e^{2\Lambda B} + 1}{e^{2\Lambda B} - 1} \right) \left(\mathcal{R}_{\Phi_{\text{exp}}}(h) - \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{NN}}}^* + \mathcal{M}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{NN}}} \right), & \text{otherwise.} \end{cases} \quad (45) \end{aligned}$$

Since the $(\ell_{0-1}, \mathcal{H}_{\text{NN}})$ -minimizability gap coincides with the $(\ell_{0-1}, \mathcal{H}_{\text{NN}})$ -approximation error and $(\Phi_{\log}, \mathcal{H}_{\text{NN}})$ -minimizability gap coincides with the $(\Phi_{\log}, \mathcal{H}_{\text{NN}})$ -approximation error for $\Lambda B = +\infty$, the inequality can be rewritten as follows:

$$\begin{aligned} & \mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^* \leq \\ & \begin{cases} \sqrt{2} \left[\mathcal{R}_{\Phi_{\text{exp}}}(h) - \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{all}}}^* \right]^{\frac{1}{2}} & B = \infty \\ \begin{cases} \sqrt{2} \left[\mathcal{R}_{\Phi_{\text{exp}}}(h) - \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{NN}}}^* + \mathcal{M}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{NN}}} \right]^{\frac{1}{2}} & \text{if } \mathcal{R}_{\Phi_{\text{exp}}}(h) - \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{NN}}}^* \leq \frac{1}{2} \left(\frac{e^{2\Lambda B} - 1}{e^{2\Lambda B} + 1} \right)^2 - \mathcal{M}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{NN}}} \\ 2 \left(\frac{e^{2\Lambda B} + 1}{e^{2\Lambda B} - 1} \right) \left(\mathcal{R}_{\Phi_{\text{exp}}}(h) - \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{NN}}}^* + \mathcal{M}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{NN}}} \right) & \text{o/w} \end{cases} & \text{o/w} \end{cases} \end{aligned}$$

where the $(\Phi_{\text{exp}}, \mathcal{H}_{\text{NN}})$ -minimizability gap $\mathcal{M}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{NN}}}$ is characterized as below, which is less than the $(\Phi_{\text{exp}}, \mathcal{H}_{\text{NN}})$ -approximation error when $\Lambda B < +\infty$:

$$\begin{aligned} \mathcal{M}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{NN}}} &= \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X \left[2\sqrt{\eta(x)(1-\eta(x))} \mathbb{1}_{1/2 \log_2 \left| \frac{\eta(x)}{1-\eta(x)} \right| \leq \Lambda W \|x\|_p + \Lambda B} \right] \\ &\quad - \mathbb{E}_X \left[\max\{\eta(x), 1-\eta(x)\} e^{-(\Lambda W \|x\|_p + \Lambda B)} \mathbb{1}_{1/2 \log_2 \left| \frac{\eta(x)}{1-\eta(x)} \right| > \Lambda W \|x\|_p + \Lambda B} \right] \\ &\quad - \left[\min\{\eta(x), 1-\eta(x)\} e^{\Lambda W \|x\|_p + \Lambda B} \mathbb{1}_{1/2 \log_2 \left| \frac{\eta(x)}{1-\eta(x)} \right| > \Lambda W \|x\|_p + \Lambda B} \right] \\ &< \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X \left[2\sqrt{\eta(x)(1-\eta(x))} \right] \\ &= \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{NN}}}^* - \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{all}}}^*. \end{aligned}$$

Therefore, the inequality for $\Lambda B = +\infty$ coincides with the consistency excess error bound known for the exponential loss (Zhang, 2004a; Mohri et al., 2018) but the one for $\Lambda B < +\infty$ is distinct and novel.

K.2.4. QUADRATIC LOSS

For the quadratic loss $\Phi_{\text{quad}}(\alpha) := (1 - \alpha)^2 \mathbb{1}_{\alpha \leq 1}$, for all $h \in \mathcal{H}_{\text{NN}}$ and $x \in \mathcal{X}$:

$$\begin{aligned} & \mathcal{C}_{\Phi_{\text{quad}}}(h, x, t) \\ &= t\Phi_{\text{quad}}(h(x)) + (1-t)\Phi_{\text{quad}}(-h(x)) \\ &= t(1-h(x))^2 \mathbb{1}_{h(x) \leq 1} + (1-t)(1+h(x))^2 \mathbb{1}_{h(x) \geq -1}. \\ & \inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\Phi_{\text{quad}}}(h, x, t) \\ &= \begin{cases} 4t(1-t), & |2t-1| \leq \Lambda W \|x\|_p + \Lambda B, \\ \max\{t, 1-t\} (1 - (\Lambda W \|x\|_p + \Lambda B))^2 + \min\{t, 1-t\} (1 + \Lambda W \|x\|_p + \Lambda B)^2, & |2t-1| > \Lambda W \|x\|_p + \Lambda B. \end{cases} \end{aligned}$$

Therefore, the $(\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}})$ -minimizability gap can be expressed as follows:

$$\begin{aligned} \mathcal{M}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}}} &= \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X \left[4\eta(x)(1-\eta(x)) \mathbb{1}_{|2\eta(x)-1| \leq \Lambda W \|x\|_p + \Lambda B} \right] \\ &\quad - \mathbb{E}_X \left[\max\{\eta(x), 1-\eta(x)\} (1 - (\Lambda W \|x\|_p + \Lambda B))^2 \mathbb{1}_{|2\eta(x)-1| > \Lambda W \|x\|_p + \Lambda B} \right] \\ &\quad - \mathbb{E}_X \left[\min\{\eta(x), 1-\eta(x)\} (1 + (\Lambda W \|x\|_p + \Lambda B))^2 \mathbb{1}_{|2\eta(x)-1| > \Lambda W \|x\|_p + \Lambda B} \right] \end{aligned} \quad (46)$$

Note $(\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}})$ -minimizability gap coincides with the $(\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}})$ -approximation error $\mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X [4\eta(x)(1-\eta(x))]$ for $\Lambda B \geq 1$.

For $\frac{1}{2} < t \leq 1$, we have

$$\begin{aligned} & \inf_{h \in \mathcal{H}_{\text{NN}}: h(x) < 0} \mathcal{C}_{\Phi_{\text{quad}}}(h, x, t) \\ &= t + (1-t) \\ &= 1 \\ & \inf_{x \in \mathcal{X}} \inf_{h \in \mathcal{H}_{\text{NN}}: h(x) < 0} \Delta \mathcal{C}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}}}(h, x, t) \\ &= \inf_{x \in \mathcal{X}} \left(\inf_{h \in \mathcal{H}_{\text{NN}}: h(x) < 0} \mathcal{C}_{\Phi_{\text{quad}}}(h, x, t) - \inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\Phi_{\text{quad}}}(h, x, t) \right) \\ &= \inf_{x \in \mathcal{X}} \begin{cases} 1 - 4t(1-t), & 2t-1 \leq \Lambda W \|x\|_p + \Lambda B, \\ 1 - t(1 - (\Lambda W \|x\|_p + \Lambda B))^2 - (1-t)(1 + \Lambda W \|x\|_p + \Lambda B)^2, & 2t-1 > \Lambda W \|x\|_p + \Lambda B. \end{cases} \\ &= \begin{cases} 1 - 4t(1-t), & 2t-1 \leq \Lambda B, \\ 1 - t(1 - \Lambda B)^2 - (1-t)(1 + \Lambda B)^2, & 2t-1 > \Lambda B. \end{cases} \\ &= \mathcal{J}(2t-1) \end{aligned}$$

where \mathcal{J} is the increasing and convex function on $[0, 1]$ defined by

$$\forall t \in [0, 1], \quad \mathcal{J}(t) = \begin{cases} t^2, & t \leq \Lambda B, \\ 2\Lambda B t - (\Lambda B)^2, & t > \Lambda B. \end{cases}$$

By Definition 3, for any $\epsilon \geq 0$, the \mathcal{H}_{NN} -estimation error transformation of the quadratic loss is as follows:

$$\mathcal{J}_{\Phi_{\text{quad}}} = \begin{cases} \mathcal{J}(t), & t \in [\epsilon, 1], \\ \frac{\mathcal{J}(\epsilon)}{\epsilon} t, & t \in [0, \epsilon]. \end{cases}$$

Therefore, when $\epsilon = 0$, $\mathcal{J}_{\Phi_{\text{quad}}}$ is convex, non-decreasing, invertible and satisfies that $\mathcal{J}_{\Phi_{\text{quad}}}(0) = 0$. By Theorem 4, we can choose $\Psi(t) = \mathcal{J}_{\Phi_{\text{quad}}}(t)$ in Theorem 3, or equivalently $\Gamma(t) = \mathcal{J}_{\Phi_{\text{quad}}}^{-1}(t) = \begin{cases} \sqrt{t}, & t \leq (\Lambda B)^2 \\ \frac{t}{2\Lambda B} + \frac{\Lambda B}{2}, & t > (\Lambda B)^2 \end{cases}$, in Theorem 12, which are optimal. Thus, by Theorem 3 or Theorem 12, setting $\epsilon = 0$ yields the \mathcal{H}_{NN} -consistency bound for the quadratic

loss, valid for all $h \in \mathcal{H}_{\text{NN}}$:

$$\begin{aligned} & \mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{NN}}}^* + \mathcal{M}_{\ell_{0-1}, \mathcal{H}_{\text{NN}}} \\ & \leq \begin{cases} \left[\mathcal{R}_{\Phi_{\text{quad}}}(h) - \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}}}^* + \mathcal{M}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}}} \right]^{\frac{1}{2}} & \text{if } \mathcal{R}_{\Phi_{\text{quad}}}(h) - \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}}}^* \leq (\Lambda B)^2 - \mathcal{M}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}}} \\ \frac{\mathcal{R}_{\Phi_{\text{quad}}}(h) - \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}}}^* + \mathcal{M}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}}}}{2\Lambda B} + \frac{\Lambda B}{2} & \text{otherwise} \end{cases} \end{aligned} \quad (47)$$

Since the $(\ell_{0-1}, \mathcal{H}_{\text{NN}})$ -minimizability gap coincides with the $(\ell_{0-1}, \mathcal{H}_{\text{NN}})$ -approximation error and $(\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}})$ -minimizability gap coincides with the $(\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}})$ -approximation error for $\Lambda B \geq 1$, the inequality can be rewritten as follows:

$$\begin{aligned} & \mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^* \\ & \leq \begin{cases} \left[\mathcal{R}_{\Phi_{\text{quad}}}(h) - \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{all}}}^* \right]^{\frac{1}{2}} & \text{if } \Lambda B \geq 1 \\ \begin{cases} \left[\mathcal{R}_{\Phi_{\text{quad}}}(h) - \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}}}^* + \mathcal{M}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}}} \right]^{\frac{1}{2}} & \text{if } \mathcal{R}_{\Phi_{\text{quad}}}(h) - \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}}}^* \leq (\Lambda B)^2 - \mathcal{M}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}}} \\ \frac{\mathcal{R}_{\Phi_{\text{quad}}}(h) - \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}}}^* + \mathcal{M}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}}}}{2\Lambda B} + \frac{\Lambda B}{2} & \text{otherwise} \end{cases} & \text{otherwise} \end{cases} \end{aligned}$$

where the $(\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}})$ -minimizability gap $\mathcal{M}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}}}$ is characterized as below, which is less than the $(\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}})$ -approximation error when $\Lambda B < 1$:

$$\begin{aligned} \mathcal{M}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}}} &= \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X \left[4\eta(x)(1 - \eta(x)) \mathbb{1}_{|2\eta(x)-1| \leq \Lambda W \|x\|_p + \Lambda B} \right] \\ &\quad - \mathbb{E}_X \left[\max\{\eta(x), 1 - \eta(x)\} (1 - (\Lambda W \|x\|_p + \Lambda B))^2 \mathbb{1}_{|2\eta(x)-1| > \Lambda W \|x\|_p + \Lambda B} \right] \\ &\quad - \mathbb{E}_X \left[\min\{\eta(x), 1 - \eta(x)\} (1 + (\Lambda W \|x\|_p + \Lambda B))^2 \mathbb{1}_{|2\eta(x)-1| > \Lambda W \|x\|_p + \Lambda B} \right] \\ &< \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X [4\eta(x)(1 - \eta(x))] \\ &= \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{NN}}}^* - \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{all}}}^*. \end{aligned}$$

Therefore, the inequality for $\Lambda B \geq 1$ coincides with the consistency excess error bound known for the quadratic loss (Zhang, 2004a; Bartlett et al., 2006) but the one for $\Lambda B < 1$ is distinct and novel.

K.2.5. SIGMOID LOSS

For the sigmoid loss $\Phi_{\text{sig}}(\alpha) := 1 - \tanh(k\alpha)$, $k > 0$, for all $h \in \mathcal{H}_{\text{NN}}$ and $x \in \mathcal{X}$:

$$\begin{aligned} \mathcal{C}_{\Phi_{\text{sig}}}(h, x, t) &= t\Phi_{\text{sig}}(h(x)) + (1-t)\Phi_{\text{sig}}(-h(x)), \\ &= t(1 - \tanh(kh(x))) + (1-t)(1 + \tanh(kh(x))). \\ \inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\Phi_{\text{sig}}}(h, x, t) &= 1 - |1 - 2t| \tanh(k(\Lambda W \|x\|_p + \Lambda B)) \end{aligned}$$

Therefore, the $(\Phi_{\text{sig}}, \mathcal{H}_{\text{NN}})$ -minimizability gap can be expressed as follows:

$$\begin{aligned} \mathcal{M}_{\Phi_{\text{sig}}, \mathcal{H}_{\text{NN}}} &= \mathcal{R}_{\Phi_{\text{sig}}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X \left[\inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\Phi_{\text{sig}}}(h, x, \eta(x)) \right] \\ &= \mathcal{R}_{\Phi_{\text{sig}}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X [1 - |1 - 2\eta(x)| \tanh(k(\Lambda W \|x\|_p + \Lambda B))]. \end{aligned} \quad (48)$$

Note $(\Phi_{\text{sig}}, \mathcal{H}_{\text{NN}})$ -minimizability gap coincides with the $(\Phi_{\text{sig}}, \mathcal{H}_{\text{NN}})$ -approximation error $\mathcal{R}_{\Phi_{\text{sig}}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X [1 - |1 - 2\eta(x)|]$ for $\Lambda B = +\infty$.

For $\frac{1}{2} < t \leq 1$, we have

$$\begin{aligned} \inf_{h \in \mathcal{H}_{\text{NN}}: h(x) < 0} \mathcal{C}_{\Phi_{\text{sig}}}(h, x, t) &= 1 - |1 - 2t| \tanh(0) \\ &= 1. \\ \inf_{x \in \mathcal{X}} \inf_{h \in \mathcal{H}_{\text{NN}}: h(x) < 0} \Delta \mathcal{C}_{\Phi_{\text{sig}}, \mathcal{H}_{\text{NN}}}(h, x, t) &= \inf_{x \in \mathcal{X}} \left(\inf_{h \in \mathcal{H}_{\text{NN}}: h(x) < 0} \mathcal{C}_{\Phi_{\text{sig}}}(h, x, t) - \inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\Phi_{\text{sig}}}(h, x, t) \right) \\ &= \inf_{x \in \mathcal{X}} (2t - 1) \tanh(k(\Lambda W \|x\|_p + \Lambda B)) \\ &= (2t - 1) \tanh(k\Lambda B) \\ &= \mathcal{J}(2t - 1) \end{aligned}$$

where \mathcal{J} is the increasing and convex function on $[0, 1]$ defined by

$$\forall t \in [0, 1], \mathcal{J}(t) = \tanh(k\Lambda B) t.$$

By Definition 3, for any $\epsilon \geq 0$, the \mathcal{H}_{NN} -estimation error transformation of the sigmoid loss is as follows:

$$\mathcal{T}_{\Phi_{\text{sig}}} = \tanh(k\Lambda B) t, \quad t \in [0, 1],$$

Therefore, $\mathcal{T}_{\Phi_{\text{sig}}}$ is convex, non-decreasing, invertible and satisfies that $\mathcal{T}_{\Phi_{\text{sig}}}(0) = 0$. By Theorem 4, we can choose $\Psi(t) = \tanh(k\Lambda B) t$ in Theorem 3, or equivalently $\Gamma(t) = \frac{t}{\tanh(k\Lambda B)}$ in Theorem 12, which are optimal when $\epsilon = 0$. Thus, by Theorem 3 or Theorem 12, setting $\epsilon = 0$ yields the \mathcal{H}_{NN} -consistency bound for the sigmoid loss, valid for all $h \in \mathcal{H}_{\text{NN}}$:

$$\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{NN}}}^* \leq \frac{\mathcal{R}_{\Phi_{\text{sig}}}(h) - \mathcal{R}_{\Phi_{\text{sig}}, \mathcal{H}_{\text{NN}}}^* + \mathcal{M}_{\Phi_{\text{sig}}, \mathcal{H}_{\text{NN}}} - \mathcal{M}_{\ell_{0-1}, \mathcal{H}_{\text{NN}}}}{\tanh(k\Lambda B)}. \quad (49)$$

Since the $(\ell_{0-1}, \mathcal{H}_{\text{NN}})$ -minimizability gap coincides with the $(\ell_{0-1}, \mathcal{H}_{\text{NN}})$ -approximation error, and since $(\Phi_{\text{sig}}, \mathcal{H}_{\text{NN}})$ -minimizability gap coincides with the $(\Phi_{\text{sig}}, \mathcal{H}_{\text{NN}})$ -approximation error for $\Lambda B = +\infty$, the inequality can be rewritten as follows:

$$\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^* \leq \begin{cases} \mathcal{R}_{\Phi_{\text{sig}}}(h) - \mathcal{R}_{\Phi_{\text{sig}}, \mathcal{H}_{\text{all}}}^* & \text{if } \Lambda B = +\infty \\ \frac{1}{\tanh(k\Lambda B)} \left[\mathcal{R}_{\Phi_{\text{sig}}}(h) - \mathbb{E}_X [1 - |1 - 2\eta(x)| \tanh(k(\Lambda W \|x\|_p + \Lambda B))] \right] & \text{otherwise.} \end{cases}$$

The inequality for $\Lambda B = +\infty$ coincides with the consistency excess error bound known for the sigmoid loss (Zhang, 2004a; Bartlett et al., 2006; Mohri et al., 2018) but the one for $\Lambda B < +\infty$ is distinct and novel. For $\Lambda B < +\infty$, we have

$$\mathbb{E}_X [1 - |1 - 2\eta(x)| \tanh(k(\Lambda W \|x\|_p + \Lambda B))] > \mathbb{E}_X [1 - |2\eta(x) - 1|] = 2\mathbb{E}_X [\min\{\eta(x), 1 - \eta(x)\}] = \mathcal{R}_{\Phi_{\text{hinge}}, \mathcal{H}_{\text{all}}}^*.$$

Therefore for $\Lambda B < +\infty$,

$$\mathcal{R}_{\Phi_{\text{sig}}}(h) - \mathbb{E}_X [1 - |1 - 2\eta(x)| \tanh(k(\Lambda W \|x\|_p + \Lambda B))] < \mathcal{R}_{\Phi_{\text{sig}}}(h) - \mathcal{R}_{\Phi_{\text{sig}}, \mathcal{H}_{\text{all}}}^*.$$

Note that: $\mathcal{R}_{\Phi_{\text{sig}}, \mathcal{H}_{\text{all}}}^* = 2\mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^* = 2\mathbb{E}_X [\min\{\eta(x), 1 - \eta(x)\}]$. Thus, the first inequality (case $\Lambda B = +\infty$) can be equivalently written as follows:

$$\forall h \in \mathcal{H}_{\text{NN}}, \mathcal{R}_{\ell_{0-1}}(h) \leq \mathcal{R}_{\Phi_{\text{sig}}}(h) - \mathbb{E}_X [\min\{\eta(x), 1 - \eta(x)\}],$$

which is a more informative upper bound than the standard inequality $\mathcal{R}_{\ell_{0-1}}(h) \leq \mathcal{R}_{\Phi_{\text{sig}}}(h)$.

K.2.6. ρ -MARGIN LOSS

For the ρ -margin loss $\Phi_\rho(\alpha) := \min\left\{1, \max\left\{0, 1 - \frac{\alpha}{\rho}\right\}\right\}$, $\rho > 0$, for all $h \in \mathcal{H}_{\text{NN}}$ and $x \in \mathcal{X}$:

$$\begin{aligned} \mathcal{C}_{\Phi_\rho}(h, x, t) &= t\Phi_\rho(h(x)) + (1-t)\Phi_\rho(-h(x)), \\ &= t \min\left\{1, \max\left\{0, 1 - \frac{h(x)}{\rho}\right\}\right\} + (1-t) \min\left\{1, \max\left\{0, 1 + \frac{h(x)}{\rho}\right\}\right\}. \\ \inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\Phi_\rho}(h, x, t) &= \min\{t, 1-t\} + \max\{t, 1-t\} \left(1 - \frac{\min\{\Lambda W \|x\|_p + \Lambda B, \rho\}}{\rho}\right). \end{aligned}$$

Therefore, the $(\Phi_\rho, \mathcal{H}_{\text{NN}})$ -minimizability gap can be expressed as follows:

$$\begin{aligned} \mathcal{M}_{\Phi_\rho, \mathcal{H}_{\text{NN}}} &= \mathcal{R}_{\Phi_\rho, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X \left[\inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\Phi_\rho}(h, x, \eta(x)) \right] \\ &= \mathcal{R}_{\Phi_\rho, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X \left[\min\{\eta(x), 1 - \eta(x)\} + \max\{\eta(x), 1 - \eta(x)\} \left(1 - \frac{\min\{\Lambda W \|x\|_p + \Lambda B, \rho\}}{\rho} \right) \right]. \end{aligned} \quad (50)$$

Note the $(\Phi_\rho, \mathcal{H}_{\text{NN}})$ -minimizability gap coincides with the $(\Phi_\rho, \mathcal{H}_{\text{NN}})$ -approximation error $\mathcal{R}_{\Phi_\rho, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}]$ for $\Lambda B \geq \rho$.

For $\frac{1}{2} < t \leq 1$, we have

$$\begin{aligned} \inf_{h \in \mathcal{H}_{\text{NN}}: h(x) < 0} \mathcal{C}_{\Phi_\rho}(h, x, t) &= t + (1 - t) \left(1 - \frac{\min\{\Lambda W \|x\|_p + \Lambda B, \rho\}}{\rho} \right). \\ \inf_{x \in \mathcal{X}} \inf_{h \in \mathcal{H}_{\text{NN}}: h(x) < 0} \Delta \mathcal{C}_{\Phi_\rho, \mathcal{H}_{\text{NN}}}(h, x) &= \inf_{x \in \mathcal{X}} \left(\inf_{h \in \mathcal{H}_{\text{NN}}: h(x) < 0} \mathcal{C}_{\Phi_\rho}(h, x, t) - \inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\Phi_\rho}(h, x, t) \right) \\ &= \inf_{x \in \mathcal{X}} (2t - 1) \frac{\min\{\Lambda W \|x\|_p + \Lambda B, \rho\}}{\rho} \\ &= (2t - 1) \frac{\min\{\Lambda B, \rho\}}{\rho} \\ &= \mathcal{T}(2t - 1) \end{aligned}$$

where \mathcal{T} is the increasing and convex function on $[0, 1]$ defined by

$$\forall t \in [0, 1], \quad \mathcal{T}(t) = \frac{\min\{\Lambda B, \rho\}}{\rho} t.$$

By Definition 3, for any $\epsilon \geq 0$, the \mathcal{H}_{NN} -estimation error transformation of the ρ -margin loss is as follows:

$$\mathcal{T}_{\Phi_\rho} = \frac{\min\{\Lambda B, \rho\}}{\rho} t, \quad t \in [0, 1],$$

Therefore, \mathcal{T}_{Φ_ρ} is convex, non-decreasing, invertible and satisfies that $\mathcal{T}_{\Phi_\rho}(0) = 0$. By Theorem 4, we can choose $\Psi(t) = \frac{\min\{\Lambda B, \rho\}}{\rho} t$ in Theorem 3, or equivalently $\Gamma(t) = \frac{\rho}{\min\{\Lambda B, \rho\}} t$ in Theorem 12, which are optimal when $\epsilon = 0$. Thus, by Theorem 3 or Theorem 12, setting $\epsilon = 0$ yields the \mathcal{H}_{NN} -consistency bound for the ρ -margin loss, valid for all $h \in \mathcal{H}_{\text{NN}}$:

$$\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{NN}}}^* \leq \frac{\rho \left(\mathcal{R}_{\Phi_\rho}(h) - \mathcal{R}_{\Phi_\rho, \mathcal{H}_{\text{NN}}}^* + \mathcal{M}_{\Phi_\rho, \mathcal{H}_{\text{NN}}} \right)}{\min\{\Lambda B, \rho\}} - \mathcal{M}_{\ell_{0-1}, \mathcal{H}_{\text{NN}}}. \quad (51)$$

Since the $(\ell_{0-1}, \mathcal{H}_{\text{NN}})$ -minimizability gap coincides with the $(\ell_{0-1}, \mathcal{H}_{\text{NN}})$ -approximation error and $(\Phi_\rho, \mathcal{H}_{\text{NN}})$ -minimizability gap coincides with the $(\Phi_\rho, \mathcal{H}_{\text{NN}})$ -approximation error for $\Lambda B \geq \rho$, the inequality can be rewritten as follows:

$$\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^* \leq \begin{cases} \mathcal{R}_{\Phi_\rho}(h) - \mathcal{R}_{\Phi_\rho, \mathcal{H}_{\text{all}}}^* & \text{if } \Lambda B \geq \rho \\ \frac{\rho \left(\mathcal{R}_{\Phi_\rho}(h) - \mathbb{E}_X \left[\min\{\eta(x), 1 - \eta(x)\} + \max\{\eta(x), 1 - \eta(x)\} \left(1 - \frac{\min\{\Lambda W \|x\|_p + \Lambda B, \rho\}}{\rho} \right) \right] \right)}{\Lambda B} & \text{otherwise.} \end{cases}$$

Note that: $\mathcal{R}_{\Phi_\rho, \mathcal{H}_{\text{all}}}^* = \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^* = \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}]$. Thus, the first inequality (case $\Lambda B \geq \rho$) can be equivalently written as follows:

$$\forall h \in \mathcal{H}_{\text{NN}}, \quad \mathcal{R}_{\ell_{0-1}}(h) \leq \mathcal{R}_{\Phi_\rho}(h).$$

The case $\Lambda B \geq \rho$ is one of the ‘‘trivial cases’’ mentioned in Section 4, where the trivial inequality $\mathcal{R}_{\ell_{0-1}}(h) \leq \mathcal{R}_{\Phi_\rho}(h)$ can be obtained directly using the fact that ℓ_{0-1} is upper bounded by Φ_ρ . This, however, does not imply that non-adversarial \mathcal{H}_{NN} -consistency bound for the ρ -margin loss is trivial when $\Lambda B > \rho$ since it is optimal.

L. Derivation of Adversarial \mathcal{H} -Consistency Bounds

L.1. Linear Hypotheses

By the definition of \mathcal{H}_{lin} , for any $x \in \mathcal{X}$,

$$\begin{aligned} \underline{h}_\gamma(x) &= w \cdot x - \gamma \|w\|_q + b \\ &\in \begin{cases} [-W\|x\|_p - \gamma W - B, W\|x\|_p - \gamma W + B] & \|x\|_p \geq \gamma \\ [-W\|x\|_p - \gamma W - B, B] & \|x\|_p < \gamma \end{cases}, \\ \bar{h}_\gamma(x) &= w \cdot x + \gamma \|w\|_q + b \\ &\in \begin{cases} [-W\|x\|_p + \gamma W - B, W\|x\|_p + \gamma W + B] & \|x\|_p \geq \gamma \\ [-B, W\|x\|_p + \gamma W + B] & \|x\|_p < \gamma \end{cases}. \end{aligned}$$

Note \mathcal{H}_{lin} is symmetric. For any $x \in \mathcal{X}$, there exist $w = 0$ and any $0 < b \leq B$ such that $w \cdot x - \gamma \|w\|_q + b > 0$. Thus by Lemma 2, for any $x \in \mathcal{X}$, $\mathcal{C}_{\ell_\gamma, \mathcal{H}_{\text{lin}}}^*(x) = \min\{\eta(x), 1 - \eta(x)\}$. The $(\ell_\gamma, \mathcal{H}_{\text{lin}})$ -minimizability gap can be expressed as follows:

$$\mathcal{M}_{\ell_\gamma, \mathcal{H}_{\text{lin}}} = \mathcal{R}_{\ell_\gamma, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}]. \quad (52)$$

L.1.1. SUPREMUM-BASED ρ -MARGIN LOSS

For the supremum-based ρ -margin loss

$$\tilde{\Phi}_\rho := \sup_{x': \|x-x'\|_p \leq \gamma} \Phi_\rho(yh(x')), \quad \text{where } \Phi_\rho(\alpha) = \min\left\{1, \max\left\{0, 1 - \frac{\alpha}{\rho}\right\}\right\}, \quad \rho > 0,$$

for all $h \in \mathcal{H}_{\text{lin}}$ and $x \in \mathcal{X}$:

$$\begin{aligned} \mathcal{C}_{\tilde{\Phi}_\rho}(h, x, t) &= t\tilde{\Phi}_\rho(h(x)) + (1-t)\tilde{\Phi}_\rho(-h(x)) \\ &= t\Phi_\rho(\underline{h}_\gamma(x)) + (1-t)\Phi_\rho(-\bar{h}_\gamma(x)) \\ &= t \min\left\{1, \max\left\{0, 1 - \frac{\underline{h}_\gamma(x)}{\rho}\right\}\right\} + (1-t) \min\left\{1, \max\left\{0, 1 + \frac{\bar{h}_\gamma(x)}{\rho}\right\}\right\}. \\ \inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\tilde{\Phi}_\rho}(h, x, t) &= \max\{t, 1-t\} \left(1 - \frac{\min\{W \max\{\|x\|_p, \gamma\} - \gamma W + B, \rho\}}{\rho}\right) + \min\{t, 1-t\}. \end{aligned}$$

Therefore, the $(\tilde{\Phi}_\rho, \mathcal{H}_{\text{lin}})$ -minimizability gap can be expressed as follows:

$$\begin{aligned} \mathcal{M}_{\tilde{\Phi}_\rho, \mathcal{H}_{\text{lin}}} &= \mathcal{R}_{\tilde{\Phi}_\rho, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X \left[\inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\tilde{\Phi}_\rho}(h, x, \eta(x)) \right] \\ &= \mathcal{R}_{\tilde{\Phi}_\rho, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X \left[\max\{\eta(x), 1 - \eta(x)\} \left(1 - \frac{\min\{W \max\{\|x\|_p, \gamma\} - \gamma W + B, \rho\}}{\rho}\right) \right] \\ &\quad - \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}]. \end{aligned} \quad (53)$$

For $\frac{1}{2} < t \leq 1$, we have

$$\begin{aligned}
 \inf_{h \in \mathcal{H}_{\text{lin}}; \underline{h}_\gamma(x) \leq 0 \leq \bar{h}_\gamma(x)} \mathcal{C}_{\tilde{\Phi}_\rho}(h, x, t) &= t + (1 - t) \\
 &= 1 \\
 \inf_{x \in \mathcal{X}} \inf_{h \in \mathcal{H}_{\text{lin}}; \underline{h}_\gamma(x) \leq 0 \leq \bar{h}_\gamma(x)} \Delta \mathcal{C}_{\tilde{\Phi}_\rho, \mathcal{H}_{\text{lin}}}(h, x, t) &= \inf_{x \in \mathcal{X}} \left\{ \inf_{h \in \mathcal{H}_{\text{lin}}; \underline{h}_\gamma(x) \leq 0 \leq \bar{h}_\gamma(x)} \mathcal{C}_{\tilde{\Phi}_\rho}(h, x, t) - \inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\tilde{\Phi}_\rho}(h, x, t) \right\} \\
 &= \inf_{x \in \mathcal{X}} \frac{\min\{W \max\{\|x\|_p, \gamma\} - \gamma W + B, \rho\}}{\rho} t \\
 &= \frac{\min\{B, \rho\}}{\rho} t \\
 &= \mathcal{T}_1(t),
 \end{aligned}$$

where \mathcal{T}_1 is the increasing and convex function on $[0, 1]$ defined by

$$\forall t \in [0, 1], \quad \mathcal{T}_1(t) = \frac{\min\{B, \rho\}}{\rho} t;$$

$$\begin{aligned}
 \inf_{h \in \mathcal{H}_{\text{lin}}; \bar{h}_\gamma(x) < 0} \mathcal{C}_{\tilde{\Phi}_\rho}(h, x, t) &= t + (1 - t) \left(1 - \frac{\min\{W \max\{\|x\|_p, \gamma\} - \gamma W + B, \rho\}}{\rho} \right) \\
 \inf_{x \in \mathcal{X}} \inf_{h \in \mathcal{H}_{\text{lin}}; \bar{h}_\gamma(x) < 0} \Delta \mathcal{C}_{\tilde{\Phi}_\rho, \mathcal{H}_{\text{lin}}}(h, x, t) &= \inf_{x \in \mathcal{X}} \left\{ \inf_{h \in \mathcal{H}_{\text{lin}}; \bar{h}_\gamma(x) < 0} \mathcal{C}_{\tilde{\Phi}_\rho}(h, x, t) - \inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\tilde{\Phi}_\rho}(h, x, t) \right\} \\
 &= \inf_{x \in \mathcal{X}} (2t - 1) \frac{\min\{W \max\{\|x\|_p, \gamma\} - \gamma W + B, \rho\}}{\rho} \\
 &= (2t - 1) \frac{\min\{B, \rho\}}{\rho} \\
 &= \mathcal{T}_2(2t - 1),
 \end{aligned}$$

where \mathcal{T}_2 is the increasing and convex function on $[0, 1]$ defined by

$$\forall t \in [0, 1], \quad \mathcal{T}_2(t) = \frac{\min\{B, \rho\}}{\rho} t;$$

By Definition 5, for any $\epsilon \geq 0$, the adversarial \mathcal{H}_{lin} -estimation error transformation of the supremum-based ρ -margin loss is as follows:

$$\mathcal{T}_{\tilde{\Phi}_\rho} = \frac{\min\{B, \rho\}}{\rho} t, \quad t \in [0, 1],$$

Therefore, $\mathcal{T}_1 = \mathcal{T}_2$ and $\mathcal{T}_{\tilde{\Phi}_\rho}$ is convex, non-decreasing, invertible and satisfies that $\mathcal{T}_{\tilde{\Phi}_\rho}(0) = 0$. By Theorem 6, we can choose $\Psi(t) = \frac{\min\{B, \rho\}}{\rho} t$ in Theorem 5, or equivalently $\Gamma(t) = \frac{\rho}{\min\{B, \rho\}} t$ in Theorem 13, which are optimal when $\epsilon = 0$. Thus, by Theorem 5 or Theorem 13, setting $\epsilon = 0$ yields the adversarial \mathcal{H}_{lin} -consistency bound for the supremum-based ρ -margin loss, valid for all $h \in \mathcal{H}_{\text{lin}}$:

$$\mathcal{R}_{\ell_\gamma}(h) - \mathcal{R}_{\ell_\gamma, \mathcal{H}_{\text{lin}}}^* \leq \frac{\rho \left(\mathcal{R}_{\tilde{\Phi}_\rho}(h) - \mathcal{R}_{\tilde{\Phi}_\rho, \mathcal{H}_{\text{lin}}}^* + \mathcal{M}_{\tilde{\Phi}_\rho, \mathcal{H}_{\text{lin}}} \right)}{\min\{B, \rho\}} - \mathcal{M}_{\ell_\gamma, \mathcal{H}_{\text{lin}}}. \quad (54)$$

Since

$$\begin{aligned}
 \mathcal{M}_{\ell_\gamma, \mathcal{H}_{\text{lin}}} &= \mathcal{R}_{\ell_\gamma, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}], \\
 \mathcal{M}_{\tilde{\Phi}_\rho, \mathcal{H}_{\text{lin}}} &= \mathcal{R}_{\tilde{\Phi}_\rho, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X \left[\max\{\eta(x), 1 - \eta(x)\} \left(1 - \frac{\min\{W \max\{\|x\|_p, \gamma\} - \gamma W + B, \rho\}}{\rho} \right) \right] \\
 &\quad - \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}],
 \end{aligned}$$

inequality (54) can be rewritten as follows:

$$\mathcal{R}_{\ell_\gamma}(h) \leq \begin{cases} \mathcal{R}_{\tilde{\Phi}_\rho}(h) - \mathbb{E}_X \left[\max\{\eta(x), 1 - \eta(x)\} \left(1 - \frac{\min\{W \max\{\|x\|_p, \gamma\} - \gamma W + B, \rho\}}{\rho} \right) \right] & \text{if } B \geq \rho \\ \rho \left(\mathcal{R}_{\tilde{\Phi}_\rho}(h) - \mathbb{E}_X \left[\max\{\eta(x), 1 - \eta(x)\} \left(1 - \frac{\min\{W \max\{\|x\|_p, \gamma\} - \gamma W + B, \rho\}}{\rho} \right) \right] \right) & \\ \frac{\phantom{\rho \left(\mathcal{R}_{\tilde{\Phi}_\rho}(h) - \mathbb{E}_X \left[\max\{\eta(x), 1 - \eta(x)\} \left(1 - \frac{\min\{W \max\{\|x\|_p, \gamma\} - \gamma W + B, \rho\}}{\rho} \right) \right] \right)}}{\min\{B, \rho\}} & \\ + \left(1 - \frac{\rho}{\min\{B, \rho\}} \right) \mathbb{E}_X [\min\{\eta(x), 1 - \eta(x)\}] & \text{otherwise.} \end{cases} \quad (55)$$

Note that: $\min\{W \max\{\|x\|_p, \gamma\} - \gamma W + B, \rho\} = \rho$ if $B \geq \rho$. Thus, the first inequality (case $B \geq \rho$) can be equivalently written as follows:

$$\forall h \in \mathcal{H}_{\text{lin}}, \quad \mathcal{R}_{\ell_\gamma}(h) \leq \mathcal{R}_{\tilde{\Phi}_\rho}(h). \quad (56)$$

The case $B \geq \rho$ is one of the ‘‘trivial cases’’ mentioned in Section 4, where the trivial inequality $\mathcal{R}_{\ell_\gamma}(h) \leq \mathcal{R}_{\tilde{\Phi}_\rho}(h)$ can be obtained directly using the fact that ℓ_γ is upper bounded by $\tilde{\Phi}_\rho$. This, however, does not imply that adversarial \mathcal{H}_{lin} -consistency bound for the supremum-based ρ -margin loss is trivial when $B > \rho$ since it is optimal.

L.2. One-Hidden-Layer ReLU Neural Networks

By the definition of \mathcal{H}_{NN} , for any $x \in \mathcal{X}$,

$$\begin{aligned} \underline{h}_\gamma(x) &= \inf_{x': \|x-x'\|_p \leq \gamma} \sum_{j=1}^n u_j(w_j \cdot x' + b)_+ \\ \bar{h}_\gamma(x) &= \sup_{x': \|x-x'\|_p \leq \gamma} \sum_{j=1}^n u_j(w_j \cdot x' + b)_+ \end{aligned}$$

Note \mathcal{H}_{NN} is symmetric. For any $x \in \mathcal{X}$, there exist $u = (\frac{1}{\Lambda}, \dots, \frac{1}{\Lambda})$, $w = 0$ and any $0 < b \leq B$ satisfy that $\underline{h}_\gamma(x) > 0$. Thus by Lemma 2, for any $x \in \mathcal{X}$, $\mathcal{C}_{\ell_\gamma, \mathcal{H}_{\text{NN}}}^*(x) = \min\{\eta(x), 1 - \eta(x)\}$. The $(\ell_\gamma, \mathcal{H}_{\text{NN}})$ -minimizability gap can be expressed as follows:

$$\mathcal{M}_{\ell_\gamma, \mathcal{H}_{\text{NN}}} = \mathcal{R}_{\ell_\gamma, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X [\min\{\eta(x), 1 - \eta(x)\}]. \quad (57)$$

L.2.1. SUPREMUM-BASED ρ -MARGIN LOSS

For the supremum-based ρ -margin loss

$$\tilde{\Phi}_\rho = \sup_{x': \|x-x'\|_p \leq \gamma} \Phi_\rho(yh(x')), \quad \text{where } \Phi_\rho(\alpha) = \min\left\{1, \max\left\{0, 1 - \frac{\alpha}{\rho}\right\}\right\}, \quad \rho > 0,$$

for all $h \in \mathcal{H}_{\text{NN}}$ and $x \in \mathcal{X}$:

$$\begin{aligned} \mathcal{C}_{\tilde{\Phi}_\rho}(h, x, t) &= t\tilde{\Phi}_\rho(h(x)) + (1-t)\tilde{\Phi}_\rho(-h(x)) \\ &= t\Phi_\rho(\underline{h}_\gamma(x)) + (1-t)\Phi_\rho(-\bar{h}_\gamma(x)) \\ &= t \min\left\{1, \max\left\{0, 1 - \frac{\underline{h}_\gamma(x)}{\rho}\right\}\right\} + (1-t) \min\left\{1, \max\left\{0, 1 + \frac{\bar{h}_\gamma(x)}{\rho}\right\}\right\}. \\ \inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\tilde{\Phi}_\rho}(h, x, t) &= \max\{t, 1-t\} \left(1 - \frac{\min\{\sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x), \rho\}}{\rho} \right) + \min\{t, 1-t\}. \end{aligned}$$

Therefore, the $(\tilde{\Phi}_\rho, \mathcal{H}_{\text{NN}})$ -minimizability gap can be expressed as follows:

$$\begin{aligned} \mathcal{M}_{\tilde{\Phi}_\rho, \mathcal{H}_{\text{NN}}} &= \mathcal{R}_{\tilde{\Phi}_\rho, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X \left[\inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\tilde{\Phi}_\rho}(h, x, \eta(x)) \right] \\ &= \mathcal{R}_{\tilde{\Phi}_\rho, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X \left[\max\{\eta(x), 1 - \eta(x)\} \left(1 - \frac{\min\{\sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x), \rho\}}{\rho} \right) \right] \\ &\quad - \mathbb{E}_X [\min\{\eta(x), 1 - \eta(x)\}]. \end{aligned} \quad (58)$$

For $\frac{1}{2} < t \leq 1$, we have

$$\begin{aligned}
 & \inf_{h \in \mathcal{H}_{\text{NN}}: \underline{h}_\gamma(x) \leq 0 \leq \bar{h}_\gamma(x)} \mathcal{C}_{\bar{\Phi}_\rho}(h, x, t) = t + (1 - t) \\
 & = 1 \\
 \inf_{x \in \mathcal{X}} \inf_{h \in \mathcal{H}_{\text{NN}}: \underline{h}_\gamma(x) \leq 0 \leq \bar{h}_\gamma(x)} \Delta \mathcal{C}_{\bar{\Phi}_\rho, \mathcal{H}_{\text{NN}}}(h, x, t) &= \inf_{x \in \mathcal{X}} \left\{ \inf_{h \in \mathcal{H}_{\text{NN}}: \underline{h}_\gamma(x) \leq 0 \leq \bar{h}_\gamma(x)} \mathcal{C}_{\bar{\Phi}_\rho}(h, x, t) - \inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\bar{\Phi}_\rho}(h, x, t) \right\} \\
 &= \inf_{x \in \mathcal{X}} \frac{\min\{\sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x), \rho\}}{\rho} t \\
 &= \frac{\min\{\inf_{x \in \mathcal{X}} \sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x), \rho\}}{\rho} t \\
 &= \mathcal{T}_1(\eta(x)),
 \end{aligned}$$

where \mathcal{T}_1 is the increasing and convex function on $[0, 1]$ defined by

$$\forall t \in [0, 1], \quad \mathcal{T}_1(t) = \frac{\min\{\inf_{x \in \mathcal{X}} \sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x), \rho\}}{\rho} t;$$

$$\begin{aligned}
 \inf_{h \in \mathcal{H}_{\text{NN}}: \bar{h}_\gamma(x) < 0} \mathcal{C}_{\bar{\Phi}_\rho}(h, x, t) &= t + (1 - t) \left(1 - \frac{\min\{\sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x), \rho\}}{\rho} \right) \\
 \inf_{x \in \mathcal{X}} \inf_{h \in \mathcal{H}_{\text{NN}}: \bar{h}_\gamma(x) < 0} \Delta \mathcal{C}_{\bar{\Phi}_\rho, \mathcal{H}_{\text{NN}}}(h, x, t) &= \inf_{x \in \mathcal{X}} \left\{ \inf_{h \in \mathcal{H}_{\text{NN}}: \bar{h}_\gamma(x) < 0} \mathcal{C}_{\bar{\Phi}_\rho}(h, x, t) - \inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\bar{\Phi}_\rho}(h, x, t) \right\} \\
 &= \inf_{x \in \mathcal{X}} (2t - 1) \frac{\min\{\sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x), \rho\}}{\rho} \\
 &= (2t - 1) \frac{\min\{\inf_{x \in \mathcal{X}} \sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x), \rho\}}{\rho} \\
 &= \mathcal{T}_2(2t - 1),
 \end{aligned}$$

where \mathcal{T}_2 is the increasing and convex function on $[0, 1]$ defined by

$$\forall t \in [0, 1], \quad \mathcal{T}_2(t) = \frac{\min\{\inf_{x \in \mathcal{X}} \sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x), \rho\}}{\rho} t;$$

By Definition 5, for any $\epsilon \geq 0$, the adversarial \mathcal{H}_{NN} -estimation error transformation of the supremum-based ρ -margin loss is as follows:

$$\mathcal{T}_{\bar{\Phi}_\rho} = \frac{\min\{\inf_{x \in \mathcal{X}} \sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x), \rho\}}{\rho} t, \quad t \in [0, 1],$$

Therefore, $\mathcal{T}_1 = \mathcal{T}_2$ and $\mathcal{T}_{\bar{\Phi}_\rho}$ is convex, non-decreasing, invertible and satisfies that $\mathcal{T}_{\bar{\Phi}_\rho}(0) = 0$. By Theorem 6, we can choose $\Psi(t) = \frac{\min\{\inf_{x \in \mathcal{X}} \sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x), \rho\}}{\rho} t$ in Theorem 5, or equivalently $\Gamma(t) = \frac{\rho}{\min\{\inf_{x \in \mathcal{X}} \sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x), \rho\}} t$ in Theorem 13, which are optimal when $\epsilon = 0$. Thus, by Theorem 5 or Theorem 13, setting $\epsilon = 0$ yields the adversarial \mathcal{H}_{NN} -consistency bound for the supremum-based ρ -margin loss, valid for all $h \in \mathcal{H}_{\text{NN}}$:

$$\mathcal{R}_{\ell_\gamma}(h) - \mathcal{R}_{\ell_\gamma, \mathcal{H}_{\text{NN}}}^* \leq \frac{\rho \left(\mathcal{R}_{\bar{\Phi}_\rho}(h) - \mathcal{R}_{\bar{\Phi}_\rho, \mathcal{H}_{\text{NN}}}^* + \mathcal{M}_{\bar{\Phi}_\rho, \mathcal{H}_{\text{NN}}} \right)}{\min\{\inf_{x \in \mathcal{X}} \sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x), \rho\}} - \mathcal{M}_{\ell_\gamma, \mathcal{H}_{\text{NN}}}. \quad (59)$$

Observe that

$$\begin{aligned}
 \inf_{x \in \mathcal{X}} \sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x) &\geq \sup_{h \in \mathcal{H}_{\text{NN}}} \inf_{x \in \mathcal{X}} \underline{h}_\gamma(x) \\
 &= \sup_{\|u\|_1 \leq \Lambda, \|w_j\|_q \leq W, |b| \leq B} \inf_{x \in \mathcal{X}} \inf_{\|s\|_p \leq \gamma} \sum_{j=1}^n u_j (w_j \cdot x + w_j \cdot s + b)_+ \\
 &\geq \sup_{\|u\|_1 \leq \Lambda, |b| \leq B} \inf_{x \in \mathcal{X}} \inf_{\|s\|_p \leq \gamma} \sum_{j=1}^n u_j (0 \cdot x + 0 \cdot s + b)_+ \\
 &= \sup_{\|u\|_1 \leq \Lambda, |b| \leq B} \sum_{j=1}^n u_j (b)_+ \\
 &= \Lambda B.
 \end{aligned}$$

Thus, the inequality can be relaxed as follows:

$$\mathcal{R}_{\ell_\gamma}(h) - \mathcal{R}_{\ell_\gamma, \mathcal{H}_{\text{NN}}}^* \leq \frac{\rho \left(\mathcal{R}_{\Phi_\rho}(h) - \mathcal{R}_{\Phi_\rho, \mathcal{H}_{\text{NN}}}^* + \mathcal{M}_{\Phi_\rho, \mathcal{H}_{\text{NN}}} \right)}{\min\{\Lambda B, \rho\}} - \mathcal{M}_{\ell_\gamma, \mathcal{H}_{\text{NN}}}. \quad (60)$$

Since

$$\begin{aligned}
 \mathcal{M}_{\ell_\gamma, \mathcal{H}_{\text{NN}}} &= \mathcal{R}_{\ell_\gamma, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}], \\
 \mathcal{M}_{\Phi_\rho, \mathcal{H}_{\text{NN}}} &= \mathcal{R}_{\Phi_\rho, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X \left[\max\{\eta(x), 1 - \eta(x)\} \left(1 - \frac{\min\{\sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x), \rho\}}{\rho} \right) \right] \\
 &\quad - \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}],
 \end{aligned}$$

inequality (59) can be rewritten as follows:

$$\mathcal{R}_{\ell_\gamma}(h) \leq \begin{cases} \mathcal{R}_{\Phi_\rho}(h) - \mathbb{E}_X \left[\max\{\eta(x), 1 - \eta(x)\} \left(1 - \frac{\min\{\sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x), \rho\}}{\rho} \right) \right] & \text{if } \Lambda B \geq \rho \\ \frac{\rho \left(\mathcal{R}_{\Phi_\rho}(h) - \mathbb{E}_X \left[\max\{\eta(x), 1 - \eta(x)\} \left(1 - \frac{\min\{\sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x), \rho\}}{\rho} \right) \right] \right)}{\min\{\Lambda B, \rho\}} & \\ + \left(1 - \frac{\rho}{\min\{\Lambda B, \rho\}} \right) \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}] & \text{otherwise.} \end{cases}$$

Observe that

$$\begin{aligned}
 \sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x) &= \sup_{\|u\|_1 \leq \Lambda, \|w_j\|_q \leq W, |b| \leq B} \inf_{x' : \|x - x'\|_p \leq \gamma} \sum_{j=1}^n u_j (w_j \cdot x' + b)_+ \\
 &\leq \inf_{x' : \|x - x'\|_p \leq \gamma} \sup_{\|u\|_1 \leq \Lambda, \|w_j\|_q \leq W, |b| \leq B} \sum_{j=1}^n u_j (w_j \cdot x' + b)_+ \\
 &= \inf_{x' : \|x - x'\|_p \leq \gamma} \Lambda (W \|x'\|_p + B) \\
 &= \begin{cases} \Lambda (W \|x\|_p - \gamma W + B) & \text{if } \|x\|_p \geq \gamma \\ \Lambda B & \text{if } \|x\|_p < \gamma \end{cases} \\
 &= \Lambda (W \max\{\|x\|_p, \gamma\} - \gamma W + B).
 \end{aligned}$$

Thus, the inequality can be further relaxed as follows:

$$\mathcal{R}_{\ell_\gamma}(h) \leq \begin{cases} \mathcal{R}_{\Phi_\rho}(h) - \mathbb{E}_X \left[\max\{\eta(x), 1 - \eta(x)\} \left(1 - \frac{\min\{\Lambda (W \max\{\|x\|_p, \gamma\} - \gamma W + B), \rho\}}{\rho} \right) \right] & \text{if } \Lambda B \geq \rho \\ \frac{\rho \left(\mathcal{R}_{\Phi_\rho}(h) - \mathbb{E}_X \left[\max\{\eta(x), 1 - \eta(x)\} \left(1 - \frac{\min\{\Lambda (W \max\{\|x\|_p, \gamma\} - \gamma W + B), \rho\}}{\rho} \right) \right] \right)}{\min\{\Lambda B, \rho\}} & \\ + \left(1 - \frac{\rho}{\min\{\Lambda B, \rho\}} \right) \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}] & \text{otherwise.} \end{cases} \quad (61)$$

Note the relaxed adversarial \mathcal{H}_{NN} -consistency bounds (59) and (61) for the supremum-based ρ -margin loss are identical to the bounds (54) and (55) in the linear case respectively modulo the replacement of B by ΛB .

M. Derivation of Non-Adversarial \mathcal{H}_{all} -Consistency Bounds under Massart's Noise Condition

With Massart's noise condition, we introduce a modified \mathcal{H} -estimation error transformation. We assume that $\epsilon = 0$ throughout this section.

Proposition 1. *Under Massart's noise condition with β , the modified \mathcal{H} -estimation error transformation of Φ for $\epsilon = 0$ is defined on $t \in [0, 1]$ by,*

$$\mathcal{T}_{\Phi}^M(t) = \mathcal{T}(t)\mathbb{1}_{t \in [2\beta, 1]} + (\mathcal{T}(2\beta)/2\beta)t\mathbb{1}_{t \in [0, 2\beta]},$$

with $\mathcal{T}(t)$ defined in Definition 3. Suppose that \mathcal{H} satisfies the condition of Lemma 1 and $\tilde{\mathcal{T}}_{\Phi}^M$ is any lower bound of \mathcal{T}_{Φ}^M such that $\tilde{\mathcal{T}}_{\Phi}^M \leq \mathcal{T}_{\Phi}^M$. If $\tilde{\mathcal{T}}_{\Phi}^M$ is convex with $\tilde{\mathcal{T}}_{\Phi}^M(0) = 0$, then, for any hypothesis $h \in \mathcal{H}$ and any distribution under Massart's noise condition with β ,

$$\tilde{\mathcal{T}}_{\Phi}^M(\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}}^* + \mathcal{M}_{\ell_{0-1}, \mathcal{H}}) \leq \mathcal{R}_{\Phi}(h) - \mathcal{R}_{\Phi, \mathcal{H}}^* + \mathcal{M}_{\Phi, \mathcal{H}}.$$

Proof. Note the condition (13) in Theorem 8 is symmetric about $\Delta\eta(x) = 0$. Thus, condition (13) uniformly holds for all distributions is equivalent to the following holds for any $t \in [1/2 + \beta, 1]$:

$$\Psi(\langle 2t - 1 \rangle_{\epsilon}) \leq \inf_{x \in \mathcal{X}, h \in \mathcal{H}: h(x) < 0} \Delta \mathcal{C}_{\Phi, \mathcal{H}}(h, x, t), \quad (62)$$

It is clear that any lower bound $\tilde{\mathcal{T}}_{\Phi}^M$ of the modified \mathcal{H} -estimation error transformation verified condition (62). Then by Theorem 8, the proof is completed. \square

M.1. Quadratic Loss

For the quadratic loss $\Phi_{\text{quad}}(\alpha) := (1 - \alpha)^2 \mathbb{1}_{\alpha \leq 1}$, for all $h \in \mathcal{H}_{\text{all}}$ and $x \in \mathcal{X}$:

$$\begin{aligned} \mathcal{C}_{\Phi_{\text{quad}}}(h, x, t) &= t\Phi_{\text{quad}}(h(x)) + (1-t)\Phi_{\text{quad}}(-h(x)) \\ &= t(1-h(x))^2 \mathbb{1}_{h(x) \leq 1} + (1-t)(1+h(x))^2 \mathbb{1}_{h(x) \geq -1}. \\ \inf_{h \in \mathcal{H}_{\text{all}}} \mathcal{C}_{\Phi_{\text{quad}}}(h, x, t) &= 4t(1-t) \\ \mathcal{M}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{all}}} &= \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{all}}}^* - \mathbb{E}_X \left[\inf_{h \in \mathcal{H}_{\text{all}}} \mathcal{C}_{\Phi_{\text{quad}}}(h, x, \eta(x)) \right] \\ &= \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{all}}}^* - \mathbb{E}_X [4\eta(x)(1-\eta(x))] \\ &= 0 \end{aligned}$$

Thus, for $\frac{1}{2} < t \leq 1$, we have

$$\begin{aligned} \inf_{h \in \mathcal{H}_{\text{all}}: h(x) < 0} \mathcal{C}_{\Phi_{\text{quad}}}(h, x, t) &= t + (1-t) \\ &= 1 \\ \inf_{x \in \mathcal{X}} \inf_{h \in \mathcal{H}_{\text{all}}: h(x) < 0} \Delta \mathcal{C}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{all}}}(h, x, t) &= \inf_{x \in \mathcal{X}} \left(\inf_{h \in \mathcal{H}_{\text{all}}: h(x) < 0} \mathcal{C}_{\Phi_{\text{quad}}}(h, x, t) - \inf_{h \in \mathcal{H}_{\text{all}}} \mathcal{C}_{\Phi_{\text{quad}}}(h, x, t) \right) \\ &= \inf_{x \in \mathcal{X}} (1 - 4t(1-t)) \\ &= 1 - 4t(1-t) \\ &= \mathcal{T}(2t-1) \end{aligned}$$

where \mathcal{T} is the increasing and convex function on $[0, 1]$ defined by

$$\forall t \in [0, 1], \quad \mathcal{T}(t) = t^2.$$

By Proposition 1, for $\epsilon = 0$, the modified \mathcal{H}_{all} -estimation error transformation of the quadratic loss under Massart's noise condition with β is as follows:

$$\mathcal{T}_{\Phi_{\text{quad}}}^M(t) = \begin{cases} 2\beta t, & t \in [0, 2\beta], \\ t^2, & t \in [2\beta, 1]. \end{cases}$$

Therefore, $\mathcal{T}_{\Phi_{\text{quad}}}^M$ is convex, non-decreasing, invertible and satisfies that $\mathcal{T}_{\Phi_{\text{quad}}}^M(0) = 0$. By Proposition 1, we obtain the \mathcal{H}_{all} -consistency bound for the quadratic loss, valid for all $h \in \mathcal{H}_{\text{all}}$ such that $\mathcal{R}_{\Phi_{\text{quad}}}(h) - \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{all}}}^* \leq \mathcal{T}(2\beta) = 4\beta^2$ and distributions \mathcal{D} satisfies Massart's noise condition with β :

$$\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^* \leq \frac{\mathcal{R}_{\Phi_{\text{quad}}}(h) - \mathcal{R}_{\Phi_{\text{quad}}, \mathcal{H}_{\text{all}}}^*}{2\beta} \quad (63)$$

M.2. Logistic Loss

For the logistic loss $\Phi_{\log}(\alpha) := \log_2(1 + e^{-\alpha})$, for all $h \in \mathcal{H}_{\text{all}}$ and $x \in \mathcal{X}$:

$$\begin{aligned} \mathcal{C}_{\Phi_{\log}}(h, x, t) &= t\Phi_{\log}(h(x)) + (1-t)\Phi_{\log}(-h(x)), \\ &= t\log_2(1 + e^{-h(x)}) + (1-t)\log_2(1 + e^{h(x)}). \\ \inf_{h \in \mathcal{H}_{\text{all}}} \mathcal{C}_{\Phi_{\log}}(h, x, t) &= -t\log_2(t) - (1-t)\log_2(1-t) \\ \mathcal{M}_{\Phi_{\log}, \mathcal{H}_{\text{all}}} &= \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{all}}}^* - \mathbb{E}_X \left[\inf_{h \in \mathcal{H}_{\text{all}}} \mathcal{C}_{\Phi_{\log}}(h, x, \eta(x)) \right] \\ &= \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{all}}}^* - \mathbb{E}_X [-\eta(x)\log_2(\eta(x)) - (1-\eta(x))\log_2(1-\eta(x))] \\ &= 0 \end{aligned}$$

Thus, for $\frac{1}{2} < t \leq 1$, we have

$$\begin{aligned} \inf_{h \in \mathcal{H}_{\text{all}}: h(x) < 0} \mathcal{C}_{\Phi_{\log}}(h, x, t) &= t\log_2(1 + e^{-0}) + (1-t)\log_2(1 + e^0) \\ &= 1, \\ \inf_{x \in \mathcal{X}} \inf_{h \in \mathcal{H}_{\text{all}}: h(x) < 0} \Delta \mathcal{C}_{\Phi_{\log}, \mathcal{H}_{\text{all}}}(h, x, t) &= \inf_{x \in \mathcal{X}} \left(\inf_{h \in \mathcal{H}_{\text{all}}: h(x) < 0} \mathcal{C}_{\Phi_{\log}}(h, x, t) - \inf_{h \in \mathcal{H}_{\text{all}}} \mathcal{C}_{\Phi_{\log}}(h, x, t) \right) \\ &= \inf_{x \in \mathcal{X}} (1 + t\log_2(t) + (1-t)\log_2(1-t)) \\ &= 1 + t\log_2(t) + (1-t)\log_2(1-t) \\ &= \mathcal{T}(2t-1), \end{aligned}$$

where \mathcal{T} is the increasing and convex function on $[0, 1]$ defined by

$$\forall t \in [0, 1], \quad \mathcal{T}(t) = \frac{t+1}{2} \log_2(t+1) + \frac{1-t}{2} \log_2(1-t)$$

By Proposition 1, for $\epsilon = 0$, the modified \mathcal{H}_{all} -estimation error transformation of the logistic loss under Massart's noise condition with β is as follows:

$$\mathcal{T}_{\Phi_{\log}}^M = \begin{cases} \mathcal{T}(t), & t \in [2\beta, 1], \\ \frac{\mathcal{T}(2\beta)}{2\beta} t, & t \in [0, 2\beta]. \end{cases}$$

Therefore, $\mathcal{T}_{\Phi_{\log}}^M$ is convex, non-decreasing, invertible and satisfies that $\mathcal{T}_{\Phi_{\log}}^M(0) = 0$. By Proposition 1, we obtain the \mathcal{H}_{all} -consistency bound for the logistic loss, valid for all $h \in \mathcal{H}_{\text{all}}$ such that $\mathcal{R}_{\Phi_{\log}}(h) - \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{all}}}^* \leq \mathcal{T}(2\beta) = \frac{2\beta+1}{2} \log_2(2\beta+1) + \frac{1-2\beta}{2} \log_2(1-2\beta)$ and distributions \mathcal{D} satisfies Massart's noise condition with β :

$$\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^* \leq \frac{2\beta \left(\mathcal{R}_{\Phi_{\log}}(h) - \mathcal{R}_{\Phi_{\log}, \mathcal{H}_{\text{all}}}^* \right)}{\frac{2\beta+1}{2} \log_2(2\beta+1) + \frac{1-2\beta}{2} \log_2(1-2\beta)} \quad (64)$$

M.3. Exponential Loss

For the exponential loss $\Phi_{\text{exp}}(\alpha) := e^{-\alpha}$, for all $h \in \mathcal{H}_{\text{all}}$ and $x \in \mathcal{X}$:

$$\begin{aligned} \mathcal{C}_{\Phi_{\text{exp}}}(h, x, t) &= t\Phi_{\text{exp}}(h(x)) + (1-t)\Phi_{\text{exp}}(-h(x)) \\ &= te^{-h(x)} + (1-t)e^{h(x)}. \\ \inf_{h \in \mathcal{H}_{\text{all}}} \mathcal{C}_{\Phi_{\text{exp}}}(h, x, t) &= 2\sqrt{t(1-t)} \\ \mathcal{M}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{all}}} &= \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{all}}}^* - \mathbb{E}_X \left[\inf_{h \in \mathcal{H}_{\text{all}}} \mathcal{C}_{\Phi_{\text{exp}}}(h, x, \eta(x)) \right] \\ &= \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{all}}}^* - \mathbb{E}_X \left[2\sqrt{\eta(x)(1-\eta(x))} \right] \\ &= 0. \end{aligned}$$

Thus, for $\frac{1}{2} < t \leq 1$, we have

$$\begin{aligned} \inf_{h \in \mathcal{H}_{\text{all}}: h(x) < 0} \mathcal{C}_{\Phi_{\text{exp}}}(h, x, t) &= te^{-0} + (1-t)e^0 \\ &= 1. \\ \inf_{x \in \mathcal{X}} \inf_{h \in \mathcal{H}_{\text{all}}: h(x) < 0} \Delta \mathcal{C}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{all}}}(h, x) &= \inf_{x \in \mathcal{X}} \left(\inf_{h \in \mathcal{H}_{\text{all}}: h(x) < 0} \mathcal{C}_{\Phi_{\text{exp}}}(h, x) - \inf_{h \in \mathcal{H}_{\text{all}}} \mathcal{C}_{\Phi_{\text{exp}}}(h, x) \right) \\ &= \inf_{x \in \mathcal{X}} \left(1 - 2\sqrt{t(1-t)} \right) \\ &= 1 - 2\sqrt{t(1-t)} \\ &= \mathcal{T}(2t-1), \end{aligned}$$

where \mathcal{T} is the increasing and convex function on $[0, 1]$ defined by

$$\forall t \in [0, 1], \quad \mathcal{T}(t) = 1 - \sqrt{1-t^2}.$$

By Proposition 1, for $\epsilon = 0$, the modified \mathcal{H}_{all} -estimation error transformation of the exponential loss under Massart's noise condition with β is as follows:

$$\mathcal{T}_{\Phi_{\text{exp}}}^M = \begin{cases} \mathcal{T}(t), & t \in [2\beta, 1], \\ \frac{\mathcal{T}(2\beta)}{2\beta} t, & t \in [0, 2\beta]. \end{cases}$$

Therefore, $\mathcal{T}_{\Phi_{\text{exp}}}^M$ is convex, non-decreasing, invertible and satisfies that $\mathcal{T}_{\Phi_{\text{exp}}}^M(0) = 0$. By Proposition 1, we obtain the \mathcal{H}_{all} -consistency bound for the exponential loss, valid for all $h \in \mathcal{H}_{\text{all}}$ such that $\mathcal{R}_{\Phi_{\text{exp}}}(h) - \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{all}}}^* \leq \mathcal{T}(2\beta) = 1 - \sqrt{1-4\beta^2}$ and distributions \mathcal{D} satisfies Massart's noise condition with β :

$$\mathcal{R}_{\ell_{0-1}}(h) - \mathcal{R}_{\ell_{0-1}, \mathcal{H}_{\text{all}}}^* \leq \frac{2\beta \left(\mathcal{R}_{\Phi_{\text{exp}}}(h) - \mathcal{R}_{\Phi_{\text{exp}}, \mathcal{H}_{\text{all}}}^* \right)}{1 - \sqrt{1-4\beta^2}} \quad (65)$$

N. Derivation of Adversarial \mathcal{H} -Consistency Bounds under Massart's Noise Condition

As with the non-adversarial scenario in Section 5.5, we introduce a modified adversarial \mathcal{H} -estimation error transformation. We assume that $\epsilon = 0$ throughout this section.

Proposition 2. *Under Massart's noise condition with β , the modified adversarial \mathcal{H} -estimation error transformation of $\tilde{\Phi}$ for $\epsilon = 0$ is defined on $t \in [0, 1]$ by*

$$\mathcal{T}_{\tilde{\Phi}}^M(t) = \min \{ \mathcal{T}_1^M(t), \mathcal{T}_2^M(t) \},$$

where

$$\begin{aligned}\mathcal{T}_1^M(t) &:= \widehat{\mathcal{T}}_1(t) \mathbb{1}_{t \in [\frac{1}{2} + \beta, 1]} + 2/(1 + 2\beta) \widehat{\mathcal{T}}_1\left(\frac{1}{2} + \beta\right) t \mathbb{1}_{t \in [0, \frac{1}{2} + \beta)}, \\ \mathcal{T}_2^M(t) &:= \widehat{\mathcal{T}}_2(t) \mathbb{1}_{t \in [2\beta, 1]} + \frac{\widehat{\mathcal{T}}_2(2\beta)}{2\beta} t \mathbb{1}_{t \in [0, 2\beta)},\end{aligned}$$

with $\widehat{\mathcal{T}}_1(t)$ and $\widehat{\mathcal{T}}_2(t)$ defined in Definition 5. Suppose that \mathcal{H} is symmetric and $\widetilde{\mathcal{T}}_{\Phi}^M$ is any lower bound of \mathcal{T}_{Φ}^M such that $\widetilde{\mathcal{T}}_{\Phi}^M \leq \mathcal{T}_{\Phi}^M$. If $\widetilde{\mathcal{T}}_{\Phi}^M$ is convex with $\widetilde{\mathcal{T}}_{\Phi}^M(0) = 0$, then, for any hypothesis $h \in \mathcal{H}$ and any distribution under Massart's noise condition with β ,

$$\widetilde{\mathcal{T}}_{\Phi}^M(\mathcal{R}_{\ell_{\gamma}}(h) - \mathcal{R}_{\ell_{\gamma}, \mathcal{H}}^* + \mathcal{M}_{\ell_{\gamma}, \mathcal{H}}) \leq \mathcal{R}_{\Phi}(h) - \mathcal{R}_{\Phi, \mathcal{H}}^* + \mathcal{M}_{\Phi, \mathcal{H}}.$$

Proof. Note the condition (17) in Theorem 10 is symmetric about $\Delta\eta(x) = 0$. Thus, condition (17) uniformly holds for all distributions under Massart's noise condition with β is equivalent to the following holds for any $t \in [1/2 + \beta, 1]$:

$$\begin{aligned}\Psi(\langle t \rangle_{\epsilon}) &\leq \inf_{x \in \mathcal{X}, h \in \mathcal{H}_{\gamma}(x) \not\subseteq \mathcal{H}} \Delta \mathcal{C}_{\Phi, \mathcal{H}}(h, x, t), \\ \Psi(\langle 2t - 1 \rangle_{\epsilon}) &\leq \inf_{x \in \mathcal{X}, h \in \mathcal{H}: \bar{h}_{\gamma}(x) < 0} \Delta \mathcal{C}_{\Phi, \mathcal{H}}(h, x, t),\end{aligned}\tag{66}$$

It is clear that any lower bound $\widetilde{\mathcal{T}}_{\Phi}^M$ of the modified adversarial \mathcal{H} -estimation error transformation verified condition (66). Then by Theorem 10, the proof is completed. \square

N.1. Linear Hypotheses

By the definition of \mathcal{H}_{lin} , for any $x \in \mathcal{X}$,

$$\begin{aligned}\underline{h}_{\gamma}(x) &= w \cdot x - \gamma \|w\|_q + b \\ &\in \begin{cases} [-W \|x\|_p - \gamma W - B, W \|x\|_p - \gamma W + B] & \|x\|_p \geq \gamma \\ [-W \|x\|_p - \gamma W - B, B] & \|x\|_p < \gamma \end{cases}, \\ \bar{h}_{\gamma}(x) &= w \cdot x + \gamma \|w\|_q + b \\ &\in \begin{cases} [-W \|x\|_p + \gamma W - B, W \|x\|_p + \gamma W + B] & \|x\|_p \geq \gamma \\ [-B, W \|x\|_p + \gamma W + B] & \|x\|_p < \gamma \end{cases}.\end{aligned}$$

Note \mathcal{H}_{lin} is symmetric. For any $x \in \mathcal{X}$, there exist $w = 0$ and any $0 < b \leq B$ such that $w \cdot x - \gamma \|w\|_q + b > 0$. Thus by Lemma 2, for any $x \in \mathcal{X}$, $\mathcal{C}_{\ell_{\gamma}, \mathcal{H}_{\text{lin}}}^*(x) = \min\{\eta(x), 1 - \eta(x)\}$. The $(\ell_{\gamma}, \mathcal{H}_{\text{lin}})$ -minimizability gap can be expressed as follows:

$$\mathcal{M}_{\ell_{\gamma}, \mathcal{H}_{\text{lin}}} = \mathcal{R}_{\ell_{\gamma}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}].$$

N.1.1. SUPREMUM-BASED HINGE LOSS

For the supremum-based hinge loss

$$\widetilde{\Phi}_{\text{hinge}} := \sup_{x': \|x - x'\|_p \leq \gamma} \Phi_{\text{hinge}}(yh(x')), \quad \text{where } \Phi_{\text{hinge}}(\alpha) = \max\{0, 1 - \alpha\},$$

for all $h \in \mathcal{H}_{\text{lin}}$ and $x \in \mathcal{X}$:

$$\begin{aligned}
& \mathcal{C}_{\tilde{\Phi}_{\text{hinge}}}(h, x, t) \\
&= t\tilde{\Phi}_{\text{hinge}}(h(x)) + (1-t)\tilde{\Phi}_{\text{hinge}}(-h(x)) \\
&= t\Phi_{\text{hinge}}(\underline{h}_\gamma(x)) + (1-t)\Phi_{\text{hinge}}(-\bar{h}_\gamma(x)) \\
&= t \max\{0, 1 - \underline{h}_\gamma(x)\} + (1-t) \max\{0, 1 + \bar{h}_\gamma(x)\} \\
&\geq [t \max\{0, 1 - \bar{h}_\gamma(x)\} + (1-t) \max\{0, 1 + \bar{h}_\gamma(x)\}] \wedge [t \max\{0, 1 - \underline{h}_\gamma(x)\} + (1-t) \max\{0, 1 + \underline{h}_\gamma(x)\}] \\
&\inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\tilde{\Phi}_{\text{hinge}}}(h, x, t) \\
&\geq \inf_{h \in \mathcal{H}_{\text{lin}}} [t \max\{0, 1 - \bar{h}_\gamma(x)\} + (1-t) \max\{0, 1 + \bar{h}_\gamma(x)\}] \wedge \inf_{h \in \mathcal{H}_{\text{lin}}} [t \max\{0, 1 - \underline{h}_\gamma(x)\} + (1-t) \max\{0, 1 + \underline{h}_\gamma(x)\}] \\
&= 1 - |2t - 1| \min\{W \max\{\|x\|_p, \gamma\} - \gamma W + B, 1\} \\
&\inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\tilde{\Phi}_{\text{hinge}}}(h, x, t) \\
&= \inf_{h \in \mathcal{H}_{\text{lin}}} [t \max\{0, 1 - \underline{h}_\gamma(x)\} + (1-t) \max\{0, 1 + \bar{h}_\gamma(x)\}] \\
&= \inf_{h \in \mathcal{H}_{\text{lin}}} [t \max\{0, 1 - w \cdot x + \gamma \|w\|_q - b\} + (1-t) \max\{0, 1 + w \cdot x + \gamma \|w\|_q + b\}] \\
&\leq \inf_{b \in [-B, B]} [t \max\{0, 1 - b\} + (1-t) \max\{0, 1 + b\}] \\
&= 1 - |2t - 1| \min\{B, 1\} \\
&\mathcal{M}_{\tilde{\Phi}_{\text{hinge}}, \mathcal{H}_{\text{lin}}} \\
&= \mathcal{R}_{\tilde{\Phi}_{\text{hinge}}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E} \left[\inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\tilde{\Phi}_{\text{hinge}}}(h, x, \eta(x)) \right] \\
&\leq \mathcal{R}_{\tilde{\Phi}_{\text{hinge}}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E} [1 - |2\eta(x) - 1| \min\{W \max\{\|x\|_p, \gamma\} - \gamma W + B, 1\}]
\end{aligned}$$

Thus, for $\frac{1}{2} < t \leq 1$, we have

$$\begin{aligned}
& \inf_{h \in \mathcal{H}_{\text{lin}}: \underline{h}_\gamma(x) \leq 0 \leq \bar{h}_\gamma(x)} \mathcal{C}_{\tilde{\Phi}_{\text{hinge}}}(h, x, t) = t + (1-t) \\
&= 1 \\
& \inf_{x \in \mathcal{X}} \inf_{h \in \mathcal{H}_{\text{lin}}: \underline{h}_\gamma(x) \leq 0 \leq \bar{h}_\gamma(x)} \Delta \mathcal{C}_{\tilde{\Phi}_{\text{hinge}}, \mathcal{H}_{\text{lin}}}(h, x, t) = \inf_{x \in \mathcal{X}} \left\{ 1 - \inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\tilde{\Phi}_{\text{hinge}}}(h, x, t) \right\} \\
&\geq \inf_{x \in \mathcal{X}} (2t - 1) \min\{B, 1\} \\
&= (2t - 1) \min\{B, 1\} \\
&= \mathcal{T}_1(t),
\end{aligned}$$

where \mathcal{T}_1 is the increasing and convex function on $[0, 1]$ defined by

$$\mathcal{T}_1(t) = \begin{cases} \min\{B, 1\} (2t - 1), & t \in [1/2 + \beta, 1], \\ \min\{B, 1\} \frac{4\beta}{1+2\beta} t, & t \in [0, 1/2 + \beta]. \end{cases}$$

$$\begin{aligned}
& \inf_{h \in \mathcal{H}_{\text{lin}}: \bar{h}_\gamma(x) < 0} \mathcal{C}_{\tilde{\Phi}_{\text{hinge}}}(h, x, t) \geq \inf_{h \in \mathcal{H}_{\text{lin}}: \bar{h}_\gamma(x) < 0} [t \max\{0, 1 - \bar{h}_\gamma(x)\} + (1-t) \max\{0, 1 + \bar{h}_\gamma(x)\}] \\
&= t \max\{0, 1 - 0\} + (1-t) \max\{0, 1 + 0\} \\
&= 1
\end{aligned}$$

$$\begin{aligned}
& \inf_{x \in \mathcal{X}} \inf_{h \in \mathcal{H}_{\text{lin}}: \bar{h}_\gamma(x) < 0} \Delta \mathcal{C}_{\tilde{\Phi}_{\text{hinge}}, \mathcal{H}_{\text{lin}}}(h, x, t) = \inf_{x \in \mathcal{X}} \left\{ \inf_{h \in \mathcal{H}_{\text{lin}}: \bar{h}_\gamma(x) < 0} \mathcal{C}_{\tilde{\Phi}_{\text{hinge}}}(h, x, t) - \inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\tilde{\Phi}_{\text{hinge}}}(h, x, t) \right\} \\
&\geq \inf_{x \in \mathcal{X}} (2t - 1) \min\{B, 1\} \\
&= (2t - 1) \min\{B, 1\} \\
&= \mathcal{T}_2(2t - 1),
\end{aligned}$$

where \mathcal{T}_2 is the increasing and convex function on $[0, 1]$ defined by

$$\forall t \in [0, 1], \quad \mathcal{T}_2(t) = \min\{B, 1\} t.$$

By Proposition 2, for $\epsilon = 0$, the modified adversarial \mathcal{H}_{lin} -estimation error transformation of the supremum-based hinge loss under Massart's noise condition with β is lower bounded as follows:

$$\mathcal{J}_{\tilde{\Phi}_{\text{hinge}}}^M \geq \tilde{\mathcal{J}}_{\tilde{\Phi}_{\text{hinge}}}^M := \min\{\mathcal{T}_1, \mathcal{T}_2\} = \begin{cases} \min\{B, 1\} (2t - 1), & t \in [1/2 + \beta, 1], \\ \min\{B, 1\} \frac{4\beta}{1+2\beta} t, & t \in [0, 1/2 + \beta). \end{cases}$$

Note $\tilde{\mathcal{J}}_{\tilde{\Phi}_{\text{hinge}}}^M$ is convex, non-decreasing, invertible and satisfies that $\tilde{\mathcal{J}}_{\tilde{\Phi}_{\text{hinge}}}^M(0) = 0$. By Proposition 2, using the fact that $\tilde{\mathcal{J}}_{\tilde{\Phi}_{\text{hinge}}}^M \geq \min\{B, 1\} \frac{4\beta}{1+2\beta} t$ yields the adversarial \mathcal{H}_{lin} -consistency bound for the supremum-based hinge loss, valid for all $h \in \mathcal{H}_{\text{lin}}$ and distributions \mathcal{D} satisfies Massart's noise condition with β :

$$\mathcal{R}_{\ell_\gamma}(h) - \mathcal{R}_{\ell_\gamma, \mathcal{H}_{\text{lin}}}^* \leq \frac{1+2\beta}{4\beta} \frac{\mathcal{R}_{\tilde{\Phi}_{\text{hinge}}}(h) - \mathcal{R}_{\tilde{\Phi}_{\text{hinge}}, \mathcal{H}_{\text{lin}}}^* + \mathcal{M}_{\tilde{\Phi}_{\text{hinge}}, \mathcal{H}_{\text{lin}}}}{\min\{B, 1\}} - \mathcal{M}_{\ell_\gamma, \mathcal{H}_{\text{lin}}} \quad (67)$$

Since

$$\begin{aligned} \mathcal{M}_{\ell_\gamma, \mathcal{H}_{\text{lin}}} &= \mathcal{R}_{\ell_\gamma, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}], \\ \mathcal{M}_{\tilde{\Phi}_{\text{hinge}}, \mathcal{H}_{\text{lin}}} &\leq \mathcal{R}_{\tilde{\Phi}_{\text{hinge}}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}[1 - |2\eta(x) - 1| \min\{W \max\{\|x\|_p, \gamma\} - \gamma W + B, 1\}], \end{aligned}$$

the inequality can be relaxed as follows:

$$\mathcal{R}_{\ell_\gamma}(h) \leq \frac{1+2\beta}{4\beta} \frac{\mathcal{R}_{\tilde{\Phi}_{\text{hinge}}}(h)}{\min\{B, 1\}} + \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}] - \frac{1+2\beta}{4\beta} \frac{\mathbb{E}[1 - |2\eta(x) - 1| \min\{W \max\{\|x\|_p, \gamma\} - \gamma W + B, 1\}]}{\min\{B, 1\}}$$

Note that: $\min\{W \max\{\|x\|_p, \gamma\} - \gamma W + B, 1\} \leq 1$ and $1 - |1 - 2\eta(x)| = 2 \min\{\eta(x), 1 - \eta(x)\}$. Thus the inequality can be further relaxed as follows:

$$\mathcal{R}_{\ell_\gamma}(h) \leq \frac{1+2\beta}{4\beta} \frac{\mathcal{R}_{\tilde{\Phi}_{\text{hinge}}}(h)}{\min\{B, 1\}} - \left(\frac{1+2\beta}{2\beta \min\{B, 1\}} - 1 \right) \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}].$$

When $B \geq 1$, it can be equivalently written as follows:

$$\mathcal{R}_{\ell_\gamma}(h) \leq \frac{1+2\beta}{4\beta} \mathcal{R}_{\tilde{\Phi}_{\text{hinge}}}(h) - \frac{1}{2\beta} \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}]. \quad (68)$$

N.1.2. SUPREMUM-BASED SIGMOID LOSS

For the supremum-based sigmoid loss

$$\tilde{\Phi}_{\text{sig}} := \sup_{x': \|x-x'\|_p \leq \gamma} \Phi_{\text{sig}}(yh(x')), \quad \text{where } \Phi_{\text{sig}}(\alpha) = 1 - \tanh(k\alpha), \quad k > 0,$$

for all $h \in \mathcal{H}_{\text{lin}}$ and $x \in \mathcal{X}$:

$$\begin{aligned}
\mathcal{C}_{\tilde{\Phi}_{\text{sig}}}(h, x, t) &= t\tilde{\Phi}_{\text{sig}}(h(x)) + (1-t)\tilde{\Phi}_{\text{sig}}(-h(x)) \\
&= t\Phi_{\text{sig}}(\underline{h}_\gamma(x)) + (1-t)\Phi_{\text{sig}}(-\bar{h}_\gamma(x)) \\
&= t(1 - \tanh(k\underline{h}_\gamma(x))) + (1-t)(1 + \tanh(k\bar{h}_\gamma(x))) \\
&\geq \max\{1 + (1-2t)\tanh(k\bar{h}_\gamma(x)), 1 + (1-2t)\tanh(k\underline{h}_\gamma(x))\} \\
\inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\tilde{\Phi}_{\text{sig}}}(h, x, t) &\geq \max\left\{\inf_{h \in \mathcal{H}_{\text{lin}}} [1 + (1-2t)\tanh(k\bar{h}_\gamma(x))], \inf_{h \in \mathcal{H}_{\text{lin}}} [1 + (1-2t)\tanh(k\underline{h}_\gamma(x))]\right\} \\
&= 1 - |1-2t|\tanh(k(W \max\{\|x\|_p, \gamma\} - \gamma W + B)) \\
\inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\tilde{\Phi}_{\text{sig}}}(h, x, t) &= \inf_{h \in \mathcal{H}_{\text{lin}}} [t(1 - \tanh(k(w \cdot x - \gamma\|w\|_q + b))) + (1-t)(1 + \tanh(k(w \cdot x + \gamma\|w\|_q + b)))] \\
&\leq \inf_{b \in [-B, B]} [t(1 - \tanh(kb)) + (1-t)(1 + \tanh(kb))] \\
&= \max\{t, 1-t\}(1 - \tanh(kB)) + \min\{t, 1-t\}(1 + \tanh(kB)) \\
&= 1 - |1-2t|\tanh(kB) \\
\mathcal{M}_{\tilde{\Phi}_{\text{sig}}, \mathcal{H}_{\text{lin}}} &= \mathcal{R}_{\tilde{\Phi}_{\text{sig}}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}\left[\inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\tilde{\Phi}_{\text{sig}}}(h, x, \eta(x))\right] \\
&\leq \mathcal{R}_{\tilde{\Phi}_{\text{sig}}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}[1 - |1-2\eta(x)|\tanh(k(W \max\{\|x\|_p, \gamma\} - \gamma W + B))]
\end{aligned}$$

Thus, for $\frac{1}{2} < t \leq 1$, we have

$$\begin{aligned}
\inf_{h \in \mathcal{H}_{\text{lin}}: \underline{h}_\gamma(x) \leq 0 \leq \bar{h}_\gamma(x)} \mathcal{C}_{\tilde{\Phi}_{\text{sig}}}(h, x, t) &= t + (1-t) \\
&= 1 \\
\inf_{x \in \mathcal{X}} \inf_{h \in \mathcal{H}_{\text{lin}}: \underline{h}_\gamma(x) \leq 0 \leq \bar{h}_\gamma(x)} \Delta \mathcal{C}_{\tilde{\Phi}_{\text{sig}}, \mathcal{H}_{\text{lin}}}(h, x, t) &= \inf_{x \in \mathcal{X}} \left\{1 - \inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\tilde{\Phi}_{\text{sig}}}(h, x, t)\right\} \\
&\geq \inf_{x \in \mathcal{X}} (2t - 1)\tanh(kB) \\
&= (2t - 1)\tanh(kB) \\
&= \mathcal{T}_1(t),
\end{aligned}$$

where \mathcal{T}_1 is the increasing and convex function on $[0, 1]$ defined by

$$\mathcal{T}_1(t) = \begin{cases} \tanh(kB) \frac{4\beta}{1+2\beta} t, & t \in [0, 1/2 + \beta], \\ \tanh(kB) (2t - 1), & t \in [1/2 + \beta, 1]. \end{cases}$$

$$\begin{aligned}
\inf_{h \in \mathcal{H}_{\text{lin}}: \bar{h}_\gamma(x) < 0} \mathcal{C}_{\tilde{\Phi}_{\text{sig}}}(h, x, t) &\geq \inf_{h \in \mathcal{H}_{\text{lin}}: \bar{h}_\gamma(x) < 0} [1 + (1-2t)\tanh(k\bar{h}_\gamma(x))] \\
&= 1
\end{aligned}$$

$$\begin{aligned}
\inf_{x \in \mathcal{X}} \inf_{h \in \mathcal{H}_{\text{lin}}: \bar{h}_\gamma(x) < 0} \Delta \mathcal{C}_{\tilde{\Phi}_{\text{sig}}, \mathcal{H}_{\text{lin}}}(h, x, t) &= \inf_{x \in \mathcal{X}} \left\{ \inf_{h \in \mathcal{H}_{\text{lin}}: \bar{h}_\gamma(x) < 0} \mathcal{C}_{\tilde{\Phi}_{\text{sig}}}(h, x, t) - \inf_{h \in \mathcal{H}_{\text{lin}}} \mathcal{C}_{\tilde{\Phi}_{\text{sig}}}(h, x, t) \right\} \\
&\geq \inf_{x \in \mathcal{X}} (2t - 1)\tanh(kB) \\
&= (2t - 1)\tanh(kB) \\
&= \mathcal{T}_2(2t - 1),
\end{aligned}$$

where \mathcal{T}_2 is the increasing and convex function on $[2\beta, 1]$ defined by

$$\forall t \in [0, 1], \quad \mathcal{T}_2(t) = \tanh(kB)t;$$

By Proposition 2, for $\epsilon = 0$, the modified adversarial \mathcal{H}_{lin} -estimation error transformation of the supremum-based sigmoid loss under Massart's noise condition with β is lower bounded as follows:

$$\mathcal{J}_{\tilde{\Phi}_{\text{sig}}}^M \geq \tilde{\mathcal{J}}_{\tilde{\Phi}_{\text{sig}}}^M = \min\{\mathcal{J}_1, \mathcal{J}_2\} = \begin{cases} \tanh(kB) \frac{4\beta}{1+2\beta} t, & t \in [0, 1/2 + \beta], \\ \tanh(kB) (2t - 1), & t \in [1/2 + \beta, 1]. \end{cases}$$

Note $\tilde{\mathcal{J}}_{\tilde{\Phi}_{\text{sig}}}^M$ is convex, non-decreasing, invertible and satisfies that $\tilde{\mathcal{J}}_{\tilde{\Phi}_{\text{sig}}}^M(0) = 0$. By Proposition 2, using the fact that $\tilde{\mathcal{J}}_{\tilde{\Phi}_{\text{sig}}}^M \geq \tanh(kB) \frac{4\beta}{1+2\beta} t$ yields the adversarial \mathcal{H}_{lin} -consistency bound for the supremum-based sigmoid loss, valid for all $h \in \mathcal{H}_{\text{lin}}$ and distributions \mathcal{D} satisfies Massart's noise condition with β :

$$\mathcal{R}_{\ell_\gamma}(h) - \mathcal{R}_{\ell_\gamma, \mathcal{H}_{\text{lin}}}^* \leq \frac{1 + 2\beta}{4\beta} \frac{\mathcal{R}_{\tilde{\Phi}_{\text{sig}}}(h) - \mathcal{R}_{\tilde{\Phi}_{\text{sig}}, \mathcal{H}_{\text{lin}}}^* + \mathcal{M}_{\tilde{\Phi}_{\text{sig}}, \mathcal{H}_{\text{lin}}}}{\tanh(kB)} - \mathcal{M}_{\ell_\gamma, \mathcal{H}_{\text{lin}}}. \quad (69)$$

Since

$$\begin{aligned} \mathcal{M}_{\ell_\gamma, \mathcal{H}_{\text{lin}}} &= \mathcal{R}_{\ell_\gamma, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}], \\ \mathcal{M}_{\tilde{\Phi}_{\text{sig}}, \mathcal{H}_{\text{lin}}} &\leq \mathcal{R}_{\tilde{\Phi}_{\text{sig}}, \mathcal{H}_{\text{lin}}}^* - \mathbb{E}[1 - |1 - 2\eta(x)| \tanh(k(W \max\{\|x\|_p, \gamma\} - \gamma W + B))], \end{aligned}$$

the inequality can be relaxed as follows:

$$\begin{aligned} \mathcal{R}_{\ell_\gamma}(h) &\leq \frac{1 + 2\beta}{4\beta} \frac{\mathcal{R}_{\tilde{\Phi}_{\text{sig}}}(h)}{\tanh(kB)} + \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}] - \frac{1 + 2\beta}{4\beta} \frac{\mathbb{E}[1 - |1 - 2\eta(x)| \tanh(k(W \max\{\|x\|_p, \gamma\} - \gamma W + B))]}{\tanh(kB)} \end{aligned}$$

Note that: $\tanh(k(W \max\{\|x\|_p, \gamma\} - \gamma W + B)) \leq 1$ and $1 - |1 - 2\eta(x)| = 2 \min\{\eta(x), 1 - \eta(x)\}$. Thus the inequality can be further relaxed as follows:

$$\mathcal{R}_{\ell_\gamma}(h) \leq \frac{1 + 2\beta}{4\beta} \frac{\mathcal{R}_{\tilde{\Phi}_{\text{sig}}}(h)}{\tanh(kB)} - \left(\frac{1 + 2\beta}{2\beta \tanh(kB)} - 1 \right) \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}].$$

When $B = +\infty$, it can be equivalently written as follows:

$$\mathcal{R}_{\ell_\gamma}(h) \leq \frac{1 + 2\beta}{4\beta} \mathcal{R}_{\tilde{\Phi}_{\text{sig}}}(h) - \frac{1}{2\beta} \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}]. \quad (70)$$

N.2. One-Hidden-Layer ReLU Neural Networks

By the definition of \mathcal{H}_{NN} , for any $x \in \mathcal{X}$,

$$\begin{aligned} \underline{h}_\gamma(x) &= \inf_{x': \|x-x'\|_p \leq \gamma} \sum_{j=1}^n u_j(w_j \cdot x' + b)_+ \\ \bar{h}_\gamma(x) &= \sup_{x': \|x-x'\|_p \leq \gamma} \sum_{j=1}^n u_j(w_j \cdot x' + b)_+ \end{aligned}$$

Note \mathcal{H}_{NN} is symmetric. For any $x \in \mathcal{X}$, there exist $u = (\frac{1}{\Lambda}, \dots, \frac{1}{\Lambda})$, $w = 0$ and any $0 < b \leq B$ satisfy that $\underline{h}_\gamma(x) > 0$. Thus by Lemma 2, for any $x \in \mathcal{X}$, $\mathcal{C}_{\ell_\gamma, \mathcal{H}_{\text{NN}}}^*(x) = \min\{\eta(x), 1 - \eta(x)\}$. The $(\ell_\gamma, \mathcal{H}_{\text{NN}})$ -minimizability gap can be expressed as follows:

$$\mathcal{M}_{\ell_\gamma, \mathcal{H}_{\text{NN}}} = \mathcal{R}_{\ell_\gamma, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}].$$

N.2.1. SUPREMUM-BASED HINGE LOSS

For the supremum-based hinge loss

$$\tilde{\Phi}_{\text{hinge}} := \sup_{x': \|x-x'\|_p \leq \gamma} \Phi_{\text{hinge}}(yh(x')), \quad \text{where } \Phi_{\text{hinge}}(\alpha) = \max\{0, 1 - \alpha\},$$

for all $h \in \mathcal{H}_{\text{NN}}$ and $x \in \mathcal{X}$:

$$\begin{aligned}
& \mathcal{C}_{\tilde{\Phi}_{\text{hinge}}}(h, x, t) \\
&= t\tilde{\Phi}_{\text{hinge}}(h(x)) + (1-t)\tilde{\Phi}_{\text{hinge}}(-h(x)) \\
&= t\Phi_{\text{hinge}}(\underline{h}_\gamma(x)) + (1-t)\Phi_{\text{hinge}}(-\bar{h}_\gamma(x)) \\
&= t\max\{0, 1 - \underline{h}_\gamma(x)\} + (1-t)\max\{0, 1 + \bar{h}_\gamma(x)\} \\
&\geq [t\max\{0, 1 - \bar{h}_\gamma(x)\} + (1-t)\max\{0, 1 + \bar{h}_\gamma(x)\}] \wedge [t\max\{0, 1 - \underline{h}_\gamma(x)\} + (1-t)\max\{0, 1 + \underline{h}_\gamma(x)\}] \\
&\inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\tilde{\Phi}_{\text{hinge}}}(h, x, t) \\
&\geq \inf_{h \in \mathcal{H}_{\text{NN}}} [t\max\{0, 1 - \bar{h}_\gamma(x)\} + (1-t)\max\{0, 1 + \bar{h}_\gamma(x)\}] \wedge \inf_{h \in \mathcal{H}_{\text{NN}}} [t\max\{0, 1 - \underline{h}_\gamma(x)\} + (1-t)\max\{0, 1 + \underline{h}_\gamma(x)\}] \\
&= 1 - |2t - 1| \min\left\{ \sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x), 1 \right\} \\
&\inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\tilde{\Phi}_{\text{hinge}}}(h, x, t) \\
&\leq \inf_{h \in \mathcal{H}_{\text{NN}}: w=0} \mathcal{C}_{\tilde{\Phi}_{\text{hinge}}}(h, x, t) \\
&= 1 - |2t - 1| \min\{\Lambda B, 1\} \\
&\mathcal{M}_{\tilde{\Phi}_{\text{hinge}}, \mathcal{H}_{\text{NN}}} \\
&= \mathcal{R}_{\tilde{\Phi}_{\text{hinge}}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}\left[\inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\tilde{\Phi}_{\text{hinge}}}(h, x, \eta(x)) \right] \\
&\leq \mathcal{R}_{\tilde{\Phi}_{\text{hinge}}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}\left[1 - |2\eta(x) - 1| \min\left\{ \sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x), 1 \right\} \right]
\end{aligned}$$

Thus, for $\frac{1}{2} < t \leq 1$, we have

$$\begin{aligned}
& \inf_{h \in \mathcal{H}_{\text{NN}}: \underline{h}_\gamma(x) \leq 0 \leq \bar{h}_\gamma(x)} \mathcal{C}_{\tilde{\Phi}_{\text{hinge}}}(h, x, t) = t + (1-t) \\
&= 1 \\
& \inf_{x \in \mathcal{X}} \inf_{h \in \mathcal{H}_{\text{NN}}: \underline{h}_\gamma(x) \leq 0 \leq \bar{h}_\gamma(x)} \Delta \mathcal{C}_{\tilde{\Phi}_{\text{hinge}}, \mathcal{H}_{\text{NN}}}(h, x, t) = \inf_{x \in \mathcal{X}} \left\{ 1 - \inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\tilde{\Phi}_{\text{hinge}}}(h, x, t) \right\} \\
&\geq \inf_{x \in \mathcal{X}} (2t - 1) \min\{\Lambda B, 1\} \\
&= (2t - 1) \min\{\Lambda B, 1\} \\
&= \mathcal{J}_1(t),
\end{aligned}$$

where \mathcal{J}_1 is the increasing and convex function on $[0, 1]$ defined by

$$\mathcal{J}_1(t) = \begin{cases} \min\{\Lambda B, 1\} \frac{4\beta}{1+2\beta} t, & t \in [0, 1/2 + \beta], \\ \min\{\Lambda B, 1\} (2t - 1), & t \in [1/2 + \beta, 1]. \end{cases}$$

$$\begin{aligned}
& \inf_{h \in \mathcal{H}_{\text{NN}}: \bar{h}_\gamma(x) < 0} \mathcal{C}_{\tilde{\Phi}_{\text{hinge}}}(h, x, t) \geq \inf_{h \in \mathcal{H}_{\text{NN}}: \bar{h}_\gamma(x) < 0} [t\max\{0, 1 - \bar{h}_\gamma(x)\}t + (1-t)\max\{0, 1 + \bar{h}_\gamma(x)\}] \\
&= t\max\{0, 1 - 0\} + (1-t)\max\{0, 1 + 0\} \\
&= 1
\end{aligned}$$

$$\begin{aligned}
& \inf_{x \in \mathcal{X}} \inf_{h \in \mathcal{H}_{\text{NN}}: \bar{h}_\gamma(x) < 0} \Delta \mathcal{C}_{\tilde{\Phi}_{\text{hinge}}, \mathcal{H}_{\text{NN}}}(h, x, t) = \inf_{x \in \mathcal{X}} \left\{ \inf_{h \in \mathcal{H}_{\text{NN}}: \bar{h}_\gamma(x) < 0} \mathcal{C}_{\tilde{\Phi}_{\text{hinge}}}(h, x, t) - \inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\tilde{\Phi}_{\text{hinge}}}(h, x, t) \right\} \\
&\geq \inf_{x \in \mathcal{X}} (2t - 1) \min\{\Lambda B, 1\} \\
&= (2t - 1) \min\{\Lambda B, 1\} \\
&= \mathcal{J}_2(2t - 1),
\end{aligned}$$

where \mathcal{T}_2 is the increasing and convex function on $[0, 1]$ defined by

$$\forall t \in [0, 1], \quad \mathcal{T}_2(t) = \min\{\Lambda B, 1\} t;$$

By Proposition 2, for $\epsilon = 0$, the modified adversarial \mathcal{H}_{NN} -estimation error transformation of the supremum-based hinge loss under Massart's noise condition with β is lower bounded as follows:

$$\mathcal{T}_{\tilde{\Phi}_{\text{hinge}}^M} \geq \tilde{\mathcal{T}}_{\tilde{\Phi}_{\text{hinge}}^M} := \min\{\mathcal{T}_1, \mathcal{T}_2\} = \begin{cases} \min\{\Lambda B, 1\} (2t - 1), & t \in [1/2 + \beta, 1], \\ \min\{\Lambda B, 1\} \frac{4\beta}{1+2\beta} t, & t \in [0, 1/2 + \beta]. \end{cases}$$

Note $\tilde{\mathcal{T}}_{\tilde{\Phi}_{\text{hinge}}^M}$ is convex, non-decreasing, invertible and satisfies that $\tilde{\mathcal{T}}_{\tilde{\Phi}_{\text{hinge}}^M}(0) = 0$. By Proposition 2, using the fact that $\tilde{\mathcal{T}}_{\tilde{\Phi}_{\text{hinge}}^M} \geq \min\{\Lambda B, 1\} \frac{4\beta}{1+2\beta} t$ yields the adversarial \mathcal{H}_{NN} -consistency bound for the supremum-based hinge loss, valid for all $h \in \mathcal{H}_{\text{NN}}$ and distributions \mathcal{D} satisfies Massart's noise condition with β

$$\mathcal{R}_{\ell_\gamma}(h) - \mathcal{R}_{\ell_\gamma, \mathcal{H}_{\text{NN}}}^* \leq \frac{1+2\beta}{4\beta} \frac{\mathcal{R}_{\tilde{\Phi}_{\text{hinge}}}(h) - \mathcal{R}_{\tilde{\Phi}_{\text{hinge}}, \mathcal{H}_{\text{NN}}}^* + \mathcal{M}_{\tilde{\Phi}_{\text{hinge}}, \mathcal{H}_{\text{NN}}}}{\min\{\Lambda B, 1\}} - \mathcal{M}_{\ell_\gamma, \mathcal{H}_{\text{NN}}} \quad (71)$$

Since

$$\begin{aligned} \mathcal{M}_{\ell_\gamma, \mathcal{H}_{\text{NN}}} &= \mathcal{R}_{\ell_\gamma, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}], \\ \mathcal{M}_{\tilde{\Phi}_{\text{hinge}}, \mathcal{H}_{\text{NN}}} &\leq \mathcal{R}_{\tilde{\Phi}_{\text{hinge}}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}\left[1 - |2\eta(x) - 1| \min\left\{\sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x), 1\right\}\right], \end{aligned}$$

the inequality can be relaxed as follows:

$$\mathcal{R}_{\ell_\gamma}(h) \leq \frac{1+2\beta}{4\beta} \frac{\mathcal{R}_{\tilde{\Phi}_{\text{hinge}}}(h)}{\min\{\Lambda B, 1\}} + \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}] - \frac{1+2\beta}{4\beta} \frac{\mathbb{E}[1 - |2\eta(x) - 1| \min\{\sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x), 1\}]}{\min\{\Lambda B, 1\}}$$

Observe that

$$\begin{aligned} \sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x) &= \sup_{\|u\|_1 \leq \Lambda, \|w_j\|_q \leq W, |b| \leq B} \inf_{x': \|x-x'\|_p \leq \gamma} \sum_{j=1}^n u_j (w_j \cdot x' + b)_+ \\ &\leq \inf_{x': \|x-x'\|_p \leq \gamma} \sup_{\|u\|_1 \leq \Lambda, \|w_j\|_q \leq W, |b| \leq B} \sum_{j=1}^n u_j (w_j \cdot x' + b)_+ \\ &= \inf_{x': \|x-x'\|_p \leq \gamma} \Lambda (W \|x'\|_p + B) \\ &= \begin{cases} \Lambda (W \|x\|_p - \gamma W + B) & \text{if } \|x\|_p \geq \gamma \\ \Lambda B & \text{if } \|x\|_p < \gamma \end{cases} \\ &= \Lambda (W \max\{\|x\|_p, \gamma\} - \gamma W + B). \end{aligned}$$

Thus, the inequality can be further relaxed as follows:

$$\begin{aligned} \mathcal{R}_{\ell_\gamma}(h) &\leq \frac{1+2\beta}{4\beta} \frac{\mathcal{R}_{\tilde{\Phi}_{\text{hinge}}}(h)}{\min\{\Lambda B, 1\}} + \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}] - \frac{1+2\beta}{4\beta} \frac{\mathbb{E}[1 - |2\eta(x) - 1| \min\{\Lambda (W \max\{\|x\|_p, \gamma\} - \gamma W + B), 1\}]}{\min\{\Lambda B, 1\}} \end{aligned}$$

Note that: $\min\{\Lambda (W \max\{\|x\|_p, \gamma\} - \gamma W + B), 1\} \leq 1$ and $1 - |1 - 2\eta(x)| = 2 \min\{\eta(x), 1 - \eta(x)\}$. Thus the inequality can be further relaxed as follows:

$$\mathcal{R}_{\ell_\gamma}(h) \leq \frac{1+2\beta}{4\beta} \frac{\mathcal{R}_{\tilde{\Phi}_{\text{hinge}}}(h)}{\min\{\Lambda B, 1\}} - \left(\frac{1+2\beta}{2\beta \min\{\Lambda B, 1\}} - 1\right) \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}]. \quad (72)$$

When $\Lambda B \geq 1$, it can be equivalently written as follows:

$$\mathcal{R}_{\ell_\gamma}(h) \leq \frac{1+2\beta}{4\beta} \mathcal{R}_{\tilde{\Phi}_{\text{hinge}}}(h) - \frac{1}{2\beta} \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}].$$

N.2.2. SUPREMUM-BASED SIGMOID LOSS

For the supremum-based sigmoid loss

$$\widetilde{\Phi}_{\text{sig}} := \sup_{x': \|x-x'\|_p \leq \gamma} \Phi_{\text{sig}}(yh(x')), \quad \text{where } \Phi_{\text{sig}}(\alpha) = 1 - \tanh(k\alpha), \quad k > 0,$$

for all $h \in \mathcal{H}_{\text{NN}}$ and $x \in \mathcal{X}$:

$$\begin{aligned} \mathcal{C}_{\widetilde{\Phi}_{\text{sig}}}(h, x, t) &= t\widetilde{\Phi}_{\text{sig}}(h(x)) + (1-t)\widetilde{\Phi}_{\text{sig}}(-h(x)) \\ &= t\Phi_{\text{sig}}(\underline{h}_\gamma(x)) + (1-t)\Phi_{\text{sig}}(-\bar{h}_\gamma(x)) \\ &= t(1 - \tanh(k\underline{h}_\gamma(x))) + (1-t)(1 + \tanh(k\bar{h}_\gamma(x))) \\ &\geq \max\{1 + (1-2t)\tanh(k\bar{h}_\gamma(x)), 1 + (1-2t)\tanh(k\underline{h}_\gamma(x))\} \\ \inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\widetilde{\Phi}_{\text{sig}}}(h, x, t) &\geq \max\left\{\inf_{h \in \mathcal{H}_{\text{NN}}} [1 + (1-2t)\tanh(k\bar{h}_\gamma(x))], \inf_{h \in \mathcal{H}_{\text{NN}}} [1 + (1-2t)\tanh(k\underline{h}_\gamma(x))]\right\} \\ &= 1 - |1-2t|\tanh\left(k \sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x)\right) \\ \inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\widetilde{\Phi}_{\text{sig}}}(h, x, t) &\leq \max\{t, 1-t\}(1 - \tanh(k\Lambda B)) + \min\{t, 1-t\}(1 + \tanh(k\Lambda B)) \\ &= 1 - |1-2t|\tanh(k\Lambda B) \\ \mathcal{M}_{\widetilde{\Phi}_{\text{sig}}, \mathcal{H}_{\text{NN}}} &= \mathcal{R}_{\widetilde{\Phi}_{\text{sig}}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}\left[\inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\widetilde{\Phi}_{\text{sig}}}(h, x, \eta(x))\right] \\ &\leq \mathcal{R}_{\widetilde{\Phi}_{\text{sig}}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}\left[1 - |1-2\eta(x)|\tanh\left(k \sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x)\right)\right] \end{aligned}$$

For $\frac{1}{2} < t \leq 1$, we have

$$\begin{aligned} \inf_{h \in \mathcal{H}_{\text{NN}}: \underline{h}_\gamma(x) \leq 0 \leq \bar{h}_\gamma(x)} \mathcal{C}_{\widetilde{\Phi}_{\text{sig}}}(h, x, t) &= t + (1-t) \\ &= 1 \\ \inf_{x \in \mathcal{X}} \inf_{h \in \mathcal{H}_{\text{NN}}: \underline{h}_\gamma(x) \leq 0 \leq \bar{h}_\gamma(x)} \Delta \mathcal{C}_{\widetilde{\Phi}_{\text{sig}}, \mathcal{H}_{\text{NN}}}(h, x, t) &= \inf_{x \in \mathcal{X}} \left\{1 - \inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\widetilde{\Phi}_{\text{sig}}}(h, x, t)\right\} \\ &\geq \inf_{x \in \mathcal{X}} (2t-1)\tanh(k\Lambda B) \\ &= (2t-1)\tanh(k\Lambda B) \\ &= \mathcal{J}_1(t), \end{aligned}$$

where \mathcal{J}_1 is the increasing and convex function on $[0, 1]$ defined by

$$\mathcal{J}_1(t) = \begin{cases} \tanh(k\Lambda B) \frac{4\beta}{1+2\beta} t, & t \in [0, 1/2 + \beta], \\ \tanh(k\Lambda B) (2t-1), & t \in [1/2 + \beta, 1]. \end{cases}$$

$$\begin{aligned} \inf_{h \in \mathcal{H}_{\text{NN}}: \bar{h}_\gamma(x) < 0} \mathcal{C}_{\widetilde{\Phi}_{\text{sig}}}(h, x, t) &\geq \inf_{h \in \mathcal{H}_{\text{NN}}: \bar{h}_\gamma(x) < 0} 1 + (1-2t)\tanh(k\bar{h}_\gamma(x)) \\ &= 1 \end{aligned}$$

$$\begin{aligned} \inf_{x \in \mathcal{X}} \inf_{h \in \mathcal{H}_{\text{NN}}: \bar{h}_\gamma(x) < 0} \Delta \mathcal{C}_{\widetilde{\Phi}_{\text{sig}}, \mathcal{H}_{\text{NN}}}(h, x) &= \inf_{x \in \mathcal{X}} \left\{ \inf_{h \in \mathcal{H}_{\text{NN}}: \bar{h}_\gamma(x) < 0} \mathcal{C}_{\widetilde{\Phi}_{\text{sig}}}(h, x, t) - \inf_{h \in \mathcal{H}_{\text{NN}}} \mathcal{C}_{\widetilde{\Phi}_{\text{sig}}}(h, x, t) \right\} \\ &\geq \inf_{x \in \mathcal{X}} (2t-1)\tanh(k\Lambda B) \\ &= (2t-1)\tanh(k\Lambda B) \\ &= \mathcal{J}_2(2t-1), \end{aligned}$$

where \mathcal{T}_2 is the increasing and convex function on $[0, 1]$ defined by

$$\forall t \in [0, 1], \quad \mathcal{T}_2(t) = \tanh(k\Lambda B) t;$$

By Proposition 2, for $\epsilon = 0$, the modified adversarial \mathcal{H}_{NN} -estimation error transformation of the supremum-based sigmoid loss under Massart's noise condition with β is lower bounded as follows:

$$\mathcal{T}_{\tilde{\Phi}_{\text{sig}}}^M \geq \tilde{\mathcal{T}}_{\tilde{\Phi}_{\text{sig}}}^M = \min\{\mathcal{T}_1, \mathcal{T}_2\} = \begin{cases} \tanh(k\Lambda B) \frac{4\beta}{1+2\beta} t, & t \in [0, 1/2 + \beta], \\ \tanh(k\Lambda B) (2t - 1), & t \in [1/2 + \beta, 1]. \end{cases}$$

Note $\tilde{\mathcal{T}}_{\tilde{\Phi}_{\text{sig}}}^M$ is convex, non-decreasing, invertible and satisfies that $\tilde{\mathcal{T}}_{\tilde{\Phi}_{\text{sig}}}^M(0) = 0$. By Proposition 2, using the fact that $\tilde{\mathcal{T}}_{\tilde{\Phi}_{\text{sig}}}^M \geq \tanh(k\Lambda B) \frac{4\beta}{1+2\beta} t$ yields the adversarial \mathcal{H}_{NN} -consistency bound for the supremum-based sigmoid loss, valid for all $h \in \mathcal{H}_{\text{NN}}$ and distributions \mathcal{D} satisfies Massart's noise condition with β :

$$\mathcal{R}_{\ell_\gamma}(h) - \mathcal{R}_{\ell_\gamma, \mathcal{H}_{\text{NN}}}^* \leq \frac{1+2\beta}{4\beta} \frac{\mathcal{R}_{\tilde{\Phi}_{\text{sig}}}(h) - \mathcal{R}_{\tilde{\Phi}_{\text{sig}}, \mathcal{H}_{\text{NN}}}^* + \mathcal{M}_{\tilde{\Phi}_{\text{sig}}, \mathcal{H}_{\text{NN}}} - \mathcal{M}_{\ell_\gamma, \mathcal{H}_{\text{NN}}}}{\tanh(k\Lambda B)} \quad (73)$$

Since

$$\begin{aligned} \mathcal{M}_{\ell_\gamma, \mathcal{H}_{\text{NN}}} &= \mathcal{R}_{\ell_\gamma, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}], \\ \mathcal{M}_{\tilde{\Phi}_{\text{sig}}, \mathcal{H}_{\text{NN}}} &\leq \mathcal{R}_{\tilde{\Phi}_{\text{sig}}, \mathcal{H}_{\text{NN}}}^* - \mathbb{E}\left[1 - |1 - 2\eta(x)| \tanh\left(k \sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x)\right)\right], \end{aligned}$$

the inequality can be relaxed as follows:

$$\mathcal{R}_{\ell_\gamma}(h) \leq \frac{1+2\beta}{4\beta} \frac{\mathcal{R}_{\tilde{\Phi}_{\text{sig}}}(h)}{\tanh(k\Lambda B)} + \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}] - \frac{1+2\beta}{4\beta} \frac{\mathbb{E}\left[1 - |1 - 2\eta(x)| \tanh\left(k \sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x)\right)\right]}{\tanh(k\Lambda B)}$$

Observe that

$$\begin{aligned} \sup_{h \in \mathcal{H}_{\text{NN}}} \underline{h}_\gamma(x) &= \sup_{\|u\|_1 \leq \Lambda, \|w_j\|_q \leq W, |b| \leq B} \inf_{x': \|x-x'\|_p \leq \gamma} \sum_{j=1}^n u_j (w_j \cdot x' + b)_+ \\ &\leq \inf_{x': \|x-x'\|_p \leq \gamma} \sup_{\|u\|_1 \leq \Lambda, \|w_j\|_q \leq W, |b| \leq B} \sum_{j=1}^n u_j (w_j \cdot x' + b)_+ \\ &= \inf_{x': \|x-x'\|_p \leq \gamma} \Lambda(W\|x'\|_p + B) \\ &= \begin{cases} \Lambda(W\|x\|_p - \gamma W + B) & \text{if } \|x\|_p \geq \gamma \\ \Lambda B & \text{if } \|x\|_p < \gamma \end{cases} \\ &= \Lambda(W \max\{\|x\|_p, \gamma\} - \gamma W + B). \end{aligned}$$

Thus, the inequality can be further relaxed as follows:

$$\begin{aligned} &\mathcal{R}_{\ell_\gamma}(h) \\ &\leq \frac{1+2\beta}{4\beta} \frac{\mathcal{R}_{\tilde{\Phi}_{\text{sig}}}(h)}{\tanh(k\Lambda B)} + \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}] - \frac{1+2\beta}{4\beta} \frac{\mathbb{E}\left[1 - |1 - 2\eta(x)| \tanh\left(k\Lambda(W \max\{\|x\|_p, \gamma\} - \gamma W + B)\right)\right]}{\tanh(k\Lambda B)} \end{aligned}$$

Note that: $\tanh\left(k\Lambda(W \max\{\|x\|_p, \gamma\} - \gamma W + B)\right) \leq 1$ and $1 - |1 - 2\eta(x)| = 2 \min\{\eta(x), 1 - \eta(x)\}$. Thus the inequality can be further relaxed as follows:

$$\mathcal{R}_{\ell_\gamma}(h) \leq \frac{1+2\beta}{4\beta} \frac{\mathcal{R}_{\tilde{\Phi}_{\text{sig}}}(h)}{\tanh(k\Lambda B)} - \left(\frac{1+2\beta}{2\beta \tanh(k\Lambda B)} - 1\right) \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}]. \quad (74)$$

When $\Lambda B = +\infty$, it can be equivalently written as follows:

$$\mathcal{R}_{\ell_\gamma}(h) \leq \frac{1+2\beta}{4\beta} \mathcal{R}_{\tilde{\Phi}_{\text{sig}}}(h) - \frac{1}{2\beta} \mathbb{E}_X[\min\{\eta(x), 1 - \eta(x)\}].$$