Predictor-Rejector Multi-Class Abstention: Theoretical Analysis and Algorithms

Anqi Mao Courant Institute of Mathematical Sciences, New York

Mehryar Mohri Google Research and Courant Institute of Mathematical Sciences, New York

Yutao Zhong Courant Institute of Mathematical Sciences, New York

Editors: Claire Vernade and Daniel Hsu

Abstract

We study the key framework of learning with abstention in the multi-class classification setting. In this setting, the learner can choose to abstain from making a prediction with some pre-defined cost. We present a series of new theoretical and algorithmic results for this learning problem in the predictor-rejector framework. We introduce several new families of surrogate losses for which we prove strong non-asymptotic and hypothesis set-specific consistency guarantees, thereby resolving positively two existing open questions. These guarantees provide upper bounds on the estimation error of the abstention loss function in terms of that of the surrogate loss. We analyze both a single-stage setting where the predictor and rejector are learned simultaneously and a two-stage setting crucial in applications, where the predictor is learned in a first stage using a standard surrogate loss such as cross-entropy. These guarantees suggest new multi-class abstention algorithms based on minimizing these surrogate losses. We also report the results of extensive experiments comparing these algorithms to the current state-of-the-art algorithms on CIFAR-10, CIFAR-100 and SVHN datasets. Our results demonstrate empirically the benefit of our new surrogate losses and show the remarkable performance of our broadly applicable two-stage abstention algorithm. **Keywords:** abstention, learning to abstain, consistency, learning theory

1. Introduction

The problem of learning with abstention has become increasingly crucial in various applications. In natural language generation or question-answering, it is not always possible to provide accurate or factual responses. Therefore, it is of utmost importance to develop the ability to abstain from generating a response in such cases to prevent the occurrence of misleading or incorrect information, often referred to as *hallucinations* (Wei et al., 2022; Filippova, 2020; Maynez et al., 2020). Abstention plays a critical role in several other application areas. For instance, in the field of autonomous vehicle control, an incorrect prediction can pose a significant threat to human lives. Similarly, in decision-making systems, incorrect choices can have severe ethical implications. These scenarios highlight the necessity of incorporating abstention mechanisms to ensure the safety, reliability, and ethical soundness of the learning process.

We can distinguish several broad methods for learning with abstention in the literature: *confidence-based methods*, which consist of abstaining when the score returned by a pre-trained model falls below some threshold (Chow, 1957, 1970; Bartlett and Wegkamp, 2008; Yuan and Wegkamp, 2010, 2011; Ramaswamy et al., 2018; Ni et al., 2019); *selective classification*, which analyzes a set-up with

AQMAO@CIMS.NYU.EDU

MOHRI@GOOGLE.COM

YUTAO@CIMS.NYU.EDU

a *predictor* and a *selector* and defines a selection risk or loss normalized by the expected selection or coverage (El-Yaniv et al., 2010; Wiener and El-Yaniv, 2011; El-Yaniv and Wiener, 2012; Wiener and El-Yaniv, 2015; Geifman and El-Yaniv, 2017, 2019); a *predictor-rejector formulation*, which is based on learning both a *predictor* and a *rejector*, each from a different family of functions, and that takes into account explicitly the abstention cost *c* (Cortes et al., 2016a,b, 2023; Mohri et al., 2024); and a more recent *score-based formulation* that consists of augmenting the multi-class categories with a rejection label and of abstaining when the score assigned to the rejection label is the highest (Mozannar and Sontag, 2020; Cao et al., 2022; Mao et al., 2024b). Another problem closely related to abstention is that of *deferring* to an alternative model, or even to a human in some instances. This can also be considered as a special case of the general abstention scenario and tackled in a similar way (Madras et al., 2018; Raghu et al., 2019a; Mozannar and Sontag, 2020; Okati et al., 2022; Narasimhan et al., 2022; Verma et al., 2023; Mao et al., 2023, 2024a).

We will be particularly interested in the predictor-rejector formulation, which explicitly models the cost of abstention. The selective classification of El-Yaniv et al. (2010) is also interesting, but it does not explicitly factor in the cost c and is based on a distinct objective. Confidence-based methods are also very natural and straightforward, but they may fail when the predictor is not calibrated, a property that often does not hold. Additionally, they have been shown to be suboptimal when the predictor differs from the Bayes classifier (Cortes et al., 2016a). The score-based formulation (Mozannar and Sontag, 2020) admits very common properties and also explicitly takes into account the rejection cost c. We will compare the predictor-rejector formulation with the score-based one. We will show via an example that the predictor-rejector is more natural in some instances and will also compare the two formulations in our experiments. We further elaborate on the difference between the two formulations in Appendix D.

How should the problem of multi-class classification with abstention be formulated and when is it appropriate to abstain? The extension of the results of Cortes et al. (2016a) to *multi-class classification* was found to be very challenging by Ni et al. (2019). In fact, these authors left the following as an open question: *can we define Bayes-consistent surrogate losses for the predictor-rejector abstention formulation in the multi-class setting?* This paper deals precisely with this topic: we present a series of new theoretical, algorithmic, and empirical results for multi-class learning with abstention in predictor-rejector formulation and, in particular, resolve this open question in a strongly positive way.

For the score-based formulation, a surrogate loss function based on cross-entropy was introduced by Mozannar and Sontag (2020), which was proven to be Bayes-consistent. Building upon this work, Cao et al. (2022) presented a more comprehensive collection of Bayes-consistent surrogate losses for the score-based formulation. These surrogate losses can be constructed using any consistent loss function for the standard multi-class classification problem. More recently, Mao et al. (2024b) gave an extensive analysis of surrogate losses for the score-based formulation supported by \mathcal{H} -consistency bounds. In a recent study, Mozannar et al. (2023) demonstrated that existing score-based surrogate losses for abstention are not *realizable consistent* with respect to the abstention loss, as defined by Long and Servedio (2013). Instead, they proposed a novel surrogate loss that achieves realizable (\mathcal{H}, \mathcal{R})-consistency, provided that the sets of predictors \mathcal{H} and rejectors \mathcal{R} are *closed under scaling*. However, the authors expressed uncertainty regarding the Bayes-consistency of their proposed surrogate losses and left open the question of *identifying abstention surrogate losses that are both consistent and realizable* (\mathcal{H}, \mathcal{R})-*consistent when* \mathcal{H} *and* \mathcal{R} *satisfy the scaling closure property*. We address this open question by demonstrating that our newly proposed surrogate losses benefit from both Bayes-consistency and realizable consistency. We give a more comprehensive discussion of related work in Appendix A.

We show in Section 3 that in some instances the optimal solution cannot be derived in the scorebased formulation, unless we resort to more complex scoring functions. In contrast, the solution can be straightforwardly derived in the predictor-rejector formulation. In Section 4, we present and analyze a new family of surrogate loss functions for multi-class abstention in the predictor-rejector formulation, first in the *single-stage setting*, where the predictor h and the rejector r are selected simultaneously, next in a *two-stage setting*, where first the predictor h is chosen and fixed and subsequently the rejector r is determined. The two-stage setting is crucial in many applications since the predictor h is often already learned after a costly training of several hours or days. Re-training to ensure a simultaneous learning of h and r is then inconceivable due to its prohibitive cost.

In the *single-stage setting* (Section 4.1), we first give a negative result, ruling out abstention surrogate losses that do not verify a technical condition. Next, we present several positive results for abstention surrogate losses verifying that condition, for which we prove non-asymptotic $(\mathcal{H}, \mathcal{R})$ *consistency bounds* (Awasthi et al., 2022a,b) that are stronger than Bayes-consistency. Next, in Section 4.2, we also prove $(\mathcal{H}, \mathcal{R})$ -*consistency bounds* for the *two-stage setting*. Minimizing these new surrogate losses directly result in new algorithms for multi-class abstention.

In Section 4.3, we prove realizable consistency guarantees for both single-stage and two-stage predictor-rejector surrogate losses. In Section 5, we empirically show that our two-stage predictor-rejector surrogate loss consistently outperforms state-of-the-art scored-based surrogate losses, while our single-stage one achieves comparable results. Our main contributions are summarized below:

- · Counterexample for score-based abstention formulation.
- Negative results for single-stage predictor-rejector surrogate losses.
- New families of single-stage predictor-rejector surrogate losses for which we prove strong nonasymptotic and hypothesis set-specific consistency guarantees, thereby resolving positively an open question mentioned by Ni et al. (2019).
- Two-stage predictor-rejector formulations and their H-consistency bounds guarantees.
- Realizable consistency guarantees for both single-stage and two-stage surrogate losses, which resolve positively the recent open question posed by Mozannar et al. (2023).
- Experiments on CIFAR-10, CIFAR-100 and SVHN datasets empirically demonstrating the usefulness of our proposed surrogate losses.

2. Preliminaries

We first introduce some preliminary concepts and definitions, including the description of the predictor-rejector formulation and background on \mathcal{H} -consistency bounds. We examine the standard multi-class classification scenario with an input space \mathcal{X} and a set of $n \ge 2$ classes or labels $\mathcal{Y} = \{1, \ldots, n\}$. We will denote by \mathcal{D} a distribution over $\mathcal{X} \times \mathcal{Y}$ and by $p(x, y) = \mathcal{D}(Y = y \mid X = x)$ the conditional probability of Y = y given X = x. We will also adopt the shorthand $p(x) = (p(x, 1), \ldots, p(x, n))$ to denote the vector of conditional probabilities, given $x \in \mathcal{X}$. We study more specifically the learning scenario of multi-class classification with abstention, within the predictor-rejector formulation introduced by Cortes et al. (2016b).

Predictor-rejector formulation. In this formulation of the abstention problem, we have a hypothesis set \mathcal{H} of prediction functions mapping from $\mathcal{X} \times \mathcal{Y}$ to \mathbb{R} , along with a family \mathcal{R} of abstention functions, or *rejectors*, which map from \mathcal{X} to \mathbb{R} . For an input $x \in \mathcal{X}$ and a hypothesis $h \in \mathcal{H}$, the predicted label h(x) is defined as the one with the highest score, $h(x) = \operatorname{argmax}_{y \in \mathcal{Y}} h(x, y)$, using an arbitrary yet fixed deterministic strategy to break ties. A rejector $r \in \mathcal{R}$ is used to abstain from predicting on input x if r(x) is non-positive, $r(x) \leq 0$, at a cost $c(x) \in [0, 1]$. The *predictor-rejector abstention loss* L_{abs} for this formulation is thus defined for any $(h, r) \in \mathcal{H} \times \mathcal{R}$ and $(x, y) \in \mathcal{X} \times \mathcal{Y}$ as

$$\mathcal{L}_{abs}(h, r, x, y) = \mathbb{1}_{\mathsf{h}(x) \neq y} \mathbb{1}_{r(x) > 0} + c(x) \mathbb{1}_{r(x) \le 0}.$$
(1)

When the learner does not abstain (r(x) > 0), it incurs the familiar zero-one classification loss. Otherwise, it abstains $(r(x) \le 0)$ at the expense of a cost c(x). The learning problem consists of using a finite sample of size m drawn i.i.d. from \mathcal{D} to select a predictor h and rejector r with small expected L_{abs} loss, $\mathbb{E}_{(x,y)\sim\mathcal{D}}[L_{abs}(h, r, x, y)]$. For simplification, in the following, we assume a constant cost function $c \in (0, 1)$. This assumption is not necessary however, and all $(\mathcal{H}, \mathcal{R})$ consistency bounds results in Sections 4.1 and 4.2 extend straightforwardly to general cost functions.

Optimizing the predictor-rejector abstention loss is intractable for most hypothesis sets. Thus, learning algorithms in this context must rely on a surrogate loss L for L_{abs}. In the subsequent sections, we study the consistency properties of several surrogate losses. Given a surrogate loss function L in the predictor-rejector framework, we denote by $\mathcal{E}_{L}(h,r) = \mathbb{E}_{(x,y)\sim \mathcal{D}}[L(h,r,x,y)]$ the L-expected loss of a pair $(h,r) \in \mathcal{H} \times \mathcal{R}$ and by $\mathcal{E}_{L}^{*}(\mathcal{H},\mathcal{R}) = \inf_{h \in \mathcal{H}, r \in \mathcal{R}} \mathcal{E}_{L}(h,r)$ its infimum over $\mathcal{H} \times \mathcal{R}$.

 $(\mathcal{H}, \mathcal{R})$ -consistency bounds. We will prove $(\mathcal{H}, \mathcal{R})$ -consistency bounds for several surrogate losses in the predictor-rejector framework, which extend to the predictor-rejector framework the \mathcal{H} -consistency bounds of Awasthi et al. (2021a,b, 2022a,b, 2023, 2024); Mao et al. (2023c,d,e); Zheng et al. (2023); Mao et al. (2023b,f). These are inequalities upper-bounding the predictor-rejector abstention estimation loss L_{abs} of a hypothesis $h \in \mathcal{H}$ and rejector $r \in \mathcal{R}$ with respect to their surrogate estimation loss. They admit the following form:

$$\mathcal{E}_{\mathsf{L}_{abs}}(h,r) - \mathcal{E}^*_{\mathsf{L}_{abs}}(\mathcal{H},\mathcal{R}) \le f(\mathcal{E}_{\mathsf{L}}(h,r) - \mathcal{E}^*_{\mathsf{L}}(\mathcal{H},\mathcal{R})),$$

where f is a non-decreasing function. Thus, the estimation error $(\mathcal{E}_{L_{abs}}(h,r) - \mathcal{E}^*_{L_{abs}}(\mathcal{H},\mathcal{R}))$ is then bounded by $f(\epsilon)$ if the surrogate estimation loss $(\mathcal{E}_{L}(h,r) - \mathcal{E}^*_{L}(\mathcal{H},\mathcal{R}))$ is reduced to ϵ . An important term that appears in these bounds is the *minimizability gap*, which is defined by $\mathcal{M}_{L}(\mathcal{H},\mathcal{R}) = \mathcal{E}^*_{L}(\mathcal{H},\mathcal{R}) - \mathbb{E}_x[\inf_{h \in \mathcal{H}, r \in \mathcal{R}} \mathbb{E}_y[L(h,r,X,y) | X = x]]$. When the loss function L depends solely on $h(x,\cdot)$ and r(x), which holds for most loss functions used in applications, and when \mathcal{H} and \mathcal{R} include all measurable functions, the minimizability gap is null (Steinwart, 2007, lemma 2.5). However, it is generally non-zero for restricted hypothesis sets \mathcal{H} and \mathcal{R} . The minimizability gap can be upperbounded by the approximation error $\mathcal{A}_{L}(\mathcal{H},\mathcal{R}) = \mathcal{E}^*_{L}(\mathcal{H},\mathcal{R}) - \mathbb{E}_x[\inf_{h,r} \mathbb{E}_y[L(h,r,X,y) | X = x]]$, where the infimum is taken over all measurable functions. But, the minimizability gap is a finer quantity and leads to more favorable guarantees.

3. Counterexample for score-based abstention losses

We first discuss a natural example here that can be tackled straightforwardly in the predictor-rejector formulation but for which the same solution cannot be derived in the score-based formulation setting, unless we resort to more complex functions. This natural example motivates our study of surrogate

losses for the predictor-rejector formulation in Section 4. We begin by discussing the score-based formulation and subsequently highlight its relationship with the predictor-rejector formulation.

Score-based abstention formulation. In this version of the abstention problem, the label set \mathcal{Y} is expanded by adding an extra category (n + 1), which represents abstention. We indicate the augmented set as $\widetilde{\mathcal{Y}} = \{1, \ldots, n, n + 1\}$ and consider a hypothesis set $\widetilde{\mathcal{H}}$ comprising functions that map from $\mathcal{X} \times \widetilde{\mathcal{Y}}$ to \mathbb{R} . The label assigned to an input $x \in \mathcal{X}$ by $\widetilde{h} \in \widetilde{\mathcal{H}}$ is denoted as $\widetilde{h}(x)$. It is defined as $\widetilde{h}(x) = n + 1$ if $\widetilde{h}(x, n + 1) \ge \max_{y \in \mathcal{Y}} \widetilde{h}(x, y)$; otherwise, $\widetilde{h}(x)$ is determined as an element in \mathcal{Y} with the highest score, $\widetilde{h}(x) = \arg \max_{y \in \mathcal{Y}} \widetilde{h}(x, y)$, using an arbitrary yet fixed deterministic strategy to break ties. When $\widetilde{h}(x) = n + 1$, the learner chooses to abstain from predicting for x and incurs a cost c. In contrast, it predicts the label $y = \widetilde{h}(x)$ if otherwise. The *score-based abstention loss* \widetilde{L}_{abs} for this formulation is defined for any $\widetilde{h} \in \widetilde{\mathcal{H}}$ and $(x, y) \in \mathfrak{X} \times \mathcal{Y}$ as follows:

$$\widetilde{\mathsf{L}}_{\mathrm{abs}}(\widetilde{h}, x, y) = \mathbb{1}_{\widetilde{\mathsf{h}}(x) \neq y} \mathbb{1}_{\widetilde{\mathsf{h}}(x) \neq n+1} + c \, \mathbb{1}_{\widetilde{\mathsf{h}}(x) = n+1}.$$
(2)

Thus, as in the predictor-rejector context, when the learner does not abstain $(\tilde{h}(x) \neq n+1)$, it reduces to the familiar zero-one classification loss. Conversely, when it abstains $(\tilde{h}(x) = n+1)$, it incurs the cost *c*. With a finite sample drawn i.i.d. from \mathcal{D} , the learning problem involves choosing a hypothesis \tilde{h} within $\tilde{\mathcal{H}}$ that yields a minimal expected score-based abstention loss, $\mathbb{E}_{(x,y)\sim\mathcal{D}}[\tilde{L}_{abs}(\tilde{h}, x, y)]$.

Relationship between the two formulations. In the score-based formulation, rejection is defined by the condition $\tilde{h}(x, n + 1) - \max_{y \in \mathcal{Y}} \tilde{h}(x, y) \ge 0$. Thus, a predictor-rejector formulation with $(h, r) \in \mathcal{H} \times \mathcal{R}$ can be equivalently formulated as a score-based problem with \tilde{h} defined by $\tilde{h}(x, y) = h(x, y)$ for $y \in \mathcal{Y}$ and $\tilde{h}(x, n+1) = \max_{y \in \mathcal{Y}} h(x, y) - r(x)$: $\tilde{L}_{abs}(\tilde{h}, x, y) = L_{abs}(h, r, x, y)$ for all $(x, y) \times \mathcal{X} \times \mathcal{Y}$. Note, however, that function $\tilde{h}(\cdot, n + 1)$ is in general more complex. As an example, while h and r may both be in a family of linear functions, in general, $\tilde{h}(\cdot, n + 1)$ defined in this way is no more linear. Thus, a score-based formulation might require working with more complex functions. We further elaborate on the difference between the two formulations in Appendix D.

Counterexample setting. Let $\mathcal{Y} = \{1,2\}$ and let x follow the uniform distribution on the unit ball $B_2(1) = \{x: \|x\|_2 \le 1\}$. We will consider the linear models $f \in \mathcal{F}_{\text{lin}} = \{x \to \mathbf{w} \cdot x + \mathbf{b} \mid \|\mathbf{w}\|_2 = 1\}$. We set the label of a point x as follows: fix two linear functions $f_{\text{abs}}(x) = \mathbf{w}_{\text{abs}} \cdot x + \mathbf{b}_{\text{abs}}$ and $f_{\text{pred}}(x) = \mathbf{w}_{\text{pred}} \cdot x + \mathbf{b}_{\text{pred}}$ in \mathcal{F}_{lin} , if $f_{\text{abs}}(x) \le 0$, then set y = 1 with probability $\frac{1}{2}$ and y = 2 with probability $\frac{1}{2}$; if $f_{\text{abs}}(x) > 0$ and $f_{\text{pred}}(x) > 0$, then set y = 1; if $f_{\text{abs}}(x) > 0$ and $f_{\text{pred}}(x) \le 0$, then set y = 2; see Figure 1.



We denote by \mathcal{H}_{lin} the hypothesis set of linear scoring functions $-i_0$ $-i_0$ -

In the predictor-rejector formulation, the learner seeks to select a hypothesis h in \mathcal{H}_{lin} and a rejector r in \mathcal{F}_{lin} with small expected predictor-rejector loss, $\mathbb{E}_{(x,y)\sim\mathcal{D}}[\mathsf{L}_{\text{abs}}(h,r,x,y)]$. In the score-based abstention formulation, the learner seeks to select a hypothesis \tilde{h} in \mathcal{H}_{lin} with small expected score-based abstention loss, $\mathbb{E}_{(x,y)\sim\mathcal{D}}[\tilde{\mathsf{L}}_{\text{abs}}(\tilde{h},x,y)]$. We will show that, in the predictor-rejector formulation, it is straightforward to find the Bayes solution but, in the score-based formulation, the same solution cannot be achieved, unless a more complex family of functions is adopted for $\tilde{h}(\cdot, 3)$.

Predictor-rejector formulation succeeds. For the predictor-rejector abstention loss L_{abs}, it is straightforward to see that for any cost $c \in [0, \frac{1}{2})$, the best-in-class predictor and rejector $h^*_{\mathcal{H}_{\text{lin}}}$ and $r^*_{\mathcal{F}_{\text{lin}}}$ can be expressed as follows: $h^*_{\mathcal{H}_{\text{lin}}}(\cdot, 1) = f_{\text{pred}}(\cdot), h^*_{\mathcal{H}_{\text{lin}}}(\cdot, 2) = -f_{\text{pred}}(\cdot), r^*_{\mathcal{H}_{\text{lin}}} = f_{\text{abs}}$. Moreover, it is clear that $h^*_{\mathcal{H}_{\text{lin}}}$ and $r^*_{\mathcal{F}_{\text{lin}}}$ match the Bayes solution.

Score-based abstention formulation fails. For the score-based abstention loss \widetilde{L}_{abs} , for any cost $c \in [0, \frac{1}{2})$, the best-in-class classifier $\widetilde{h}_{\widetilde{\mathcal{H}}_{lin}}^*$ has the following form: $\widetilde{h}_{\widetilde{\mathcal{H}}_{lin}}^*(\cdot, 1) = f_1(\cdot), \widetilde{h}_{\widetilde{\mathcal{H}}_{lin}}^*(\cdot, 2) = -f_1(\cdot), \widetilde{h}_{\widetilde{\mathcal{H}}_{lin}}^*(\cdot, 3) = f_2(\cdot)$, for some $f_1, f_2 \in \mathcal{F}_{lin}$. Thus, $\widetilde{h}_{\widetilde{\mathcal{H}}_{lin}}^*$ abstains from making a prediction on $\{x \in B_2(1) : f_2(x) \le |f_1(x)|\}$ and otherwise predicts according to the decision surface $f_1(x) = 0$. To match the Bayes solution, f_1 must equal f_{pred} and f_2 must satisfy the following condition: $f_2(x) \ge |f_{\text{pred}}(x)| \Leftrightarrow f_{\text{abs}}(x) \le 0$, which does not hold. Therefore, unless we resort to more complex functions for $\widetilde{h}(\cdot, 3)$, the score-based formulation cannot result in the Bayes solution.

4. Predictor-rejector surrogate losses

In this section, we present and analyze a new family of surrogate loss functions L for L_{abs}, first in the *single-stage setting*, where the predictor h and the rejector r are selected simultaneously, next in a *two-stage setting*, where first the predictor h is chosen and fixed and subsequently the rejector r is determined. We give $(\mathcal{H}, \mathcal{R})$ -consistency bounds and guarantees for both settings.

4.1. Single-stage predictor-rejector surrogate losses

In view of the expression of the predictor-rejector abstention loss $L_{abs}(h, r, x, y) = \mathbb{1}_{h(x)\neq y} \mathbb{1}_{r(x)>0} + c\mathbb{1}_{r(x)\leq 0}$, if ℓ is a surrogate loss for the zero-one multi-class classification loss over the set of labels \mathcal{Y} , then, L defined as follows is a natural surrogate loss for L_{abs} : for all $(x, y) \in \mathcal{X} \times \mathcal{Y}$,

$$\mathsf{L}(h, r, x, y) = \ell(h, x, y)\Phi(-\alpha r(x)) + \Psi(c)\Phi(\beta r(x)), \tag{3}$$

where Ψ is a non-decreasing function, Φ is a non-increasing auxiliary function upper bounding $t \mapsto \mathbb{1}_{t \leq 0}$ and α , β are positive constants. The formulation (3) of L is a multi-class generalization of the binary abstention surrogate loss proposed in (Cortes et al., 2016b,a), where the binary margin-based loss $\tilde{\Phi}(yh(x))$ and $\Psi(t) = t$ are used instead:

$$\mathsf{L}_{\mathrm{bin}}(h, r, x, y) = \Phi(yh(x))\Phi(-\alpha r(x)) + c\Phi(\beta r(x)). \tag{4}$$

Minimizing L_{bin} with a regularization term was shown to achieve state-of-the-art results in the binary case with margin-based losses $\tilde{\Phi}$ such as the exponential loss $\tilde{\Phi}_{\exp}(t) = \exp(-t)$ and the hinge loss $\tilde{\Phi}_{\text{hinge}}(t) = \max\{1 - t, 0\}$ (Cortes et al., 2016b,a). However, we will show below that its multi-class generalization L imposes a more stringent condition on the choice of the surrogate loss ℓ , which rules out for example the multi-class exponential loss. We will show, however, that several other loss functions do satisfy that condition, for example the multi-class hinge loss. In the following, for

simplicity, we consider $\Phi(t) = \exp(-t)$ as in (Cortes et al., 2016b) for the main analysis, though a similar analysis can be given for other functions Φ . We first present a negative result, ruling out surrogate losses L that are based on a loss ℓ that does not verify a certain condition.

As with (Awasthi et al., 2022b), we assume that the hypothesis sets are symmetric and complete. We say that a hypothesis set \mathcal{G} is symmetric if there exists a family \mathcal{F} of functions f mapping from \mathcal{X} to \mathbb{R} such that $\{[g(x,1),\ldots,g(x,n)]:g \in \mathcal{G}\} = \{[f_1(x),\ldots,f_n(x)]:f_1,\ldots,f_n \in \mathcal{F}\}$, for any $x \in \mathcal{X}$. We say that a hypothesis set \mathcal{H} is complete if the set of scores it generates spans \mathbb{R} , that is, $\{g(x,y):g \in \mathcal{G}\} = \mathbb{R}$, for any $(x,y) \in \mathcal{X} \times \mathcal{Y}$. The hypothesis sets widely used in practice including linear models and multilayer feedforward neural networks are all symmetric and complete.

Theorem 1 (Negative result for single-stage surrogates) Assume that \mathcal{H} is symmetric and complete, and that \mathcal{R} is complete. If there exists $x \in \mathcal{X}$ such that $\inf_{h \in \mathcal{H}} \mathbb{E}_y[\ell(h, X, y) | X = x] \neq \frac{\beta \Psi(1 - \max_{y \in \mathcal{Y}} p(x, y))}{\alpha}$, then, there does not exist a non-decreasing function $\Gamma: \mathbb{R}_+ \to \mathbb{R}_+$ with the property $\lim_{t\to 0^+} \Gamma(t) = 0$ such that the following $(\mathcal{H}, \mathcal{R})$ -consistency bound holds: for all $h \in \mathcal{H}, r \in \mathcal{R}$, and any distribution,

$$\mathcal{E}_{\mathsf{L}_{\mathrm{abs}}}(h,r) - \mathcal{E}^*_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) \leq \Gamma(\mathcal{E}_{\mathsf{L}}(h,r) - \mathcal{E}^*_{\mathsf{L}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}}(\mathcal{H},\mathcal{R})).$$

The proof (Appendix G.1) proceeds by contradiction. Assuming that the $(\mathcal{H}, \mathcal{R})$ -consistency bound is valid would entail that the pointwise best-in-class predictor and best-in-class rejector for the single-stage surrogate loss align with those of the abstention loss, which can be characterized by Lemma 20 in Appendix F. Incorporating those explicit forms into the analysis of the conditional risk of the surrogate loss leads to a contradiction upon examination of its derivatives.

In view of Theorem 1, to find a surrogate loss L that admits a meaningful $(\mathcal{H}, \mathcal{R})$ -consistency bound, we need to consider multi-class surrogate losses ℓ for which the following condition holds for any $x \in \mathcal{X}$, for some Ψ and pair $(\alpha, \beta) \in \mathbb{R}^2_+$:

$$\inf_{h \in \mathcal{H}} \mathbb{E}[\ell(h, X, y) \mid X = x] = \frac{\beta \Psi (1 - \max_{y \in \mathcal{Y}} p(x, y))}{\alpha}.$$

In the binary case, this condition is easily verifiable, as $\max_{y \in \mathcal{Y}} p(x, y)$ uniquely determines the other probabilities. However, in the multi-class scenario, a fixed $\max_{y \in \mathcal{Y}} p(x, y)$ still allows for variation in the other probabilities within $\inf_{h \in \mathcal{H}} \mathbb{E}_y \left[\ell(h, X, y) \mid X = x \right]$. This leads to the difficulties in extending the binary framework to the multi-class classification.

Nevertheless, we will show that this necessary condition is satisfied by three common multi-class surrogate losses ℓ . Furthermore, we will prove $(\mathcal{H}, \mathcal{R})$ -consistency bounds for the predictor-rejector surrogate losses L based on any of these three choices of ℓ defined for all $h \in \mathcal{H}$ and (x, y) as follows:

(i) The mean absolute error loss (Ghosh et al., 2017): $\ell_{\text{mae}}(h, x, y) = 1 - \frac{e^{h(x, y)}}{\sum_{y' \in y} e^{h(x, y')}};$

(ii) The constrained ρ -hinge loss: $\ell_{\rho-\text{hinge}}(h, x, y) = \sum_{y' \neq y} \Phi_{\rho-\text{hinge}}(-h(x, y')), \rho > 0$, with $\Phi_{\rho-\text{hinge}}(t) = \max\{0, 1 - \frac{t}{\rho}\}$ the ρ -hinge loss, and the constraint $\sum_{y \in \mathcal{Y}} h(x, y) = 0$.

(iii) The ρ -Margin loss: $\ell_{\rho}(h, x, y) = \Phi_{\rho}(\rho_h(x, y))$, with $\rho_h(x, y) = h(x, y) - \max_{y' \neq y} h(x, y')$ the confidence margin and $\Phi_{\rho}(t) = \min\{\max\{0, 1 - \frac{t}{\rho}\}, 1\}, \rho > 0$ the ρ -margin loss.

Theorem 2 ($(\mathcal{H}, \mathcal{R})$ -consistency bounds for single-stage surrogates) Assume that \mathcal{H} is symmetric and complete and \mathcal{R} is complete. Then, for $\alpha = \beta$, and $\ell = \ell_{mae}$, or $\ell = \ell_{\rho}$ with $\Psi(t) = t$, or

 $\ell = \ell_{\rho-\text{hinge}}$ with $\Psi(t) = nt$, the following $(\mathfrak{H}, \mathfrak{R})$ -consistency bound holds for all $h \in \mathfrak{H}$, $r \in \mathfrak{R}$ and any distribution:

$$\mathcal{E}_{\mathsf{L}_{\mathrm{abs}}}(h,r) - \mathcal{E}^{*}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) \leq \Gamma(\mathcal{E}_{\mathsf{L}}(h,r) - \mathcal{E}^{*}_{\mathsf{L}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}}(\mathcal{H},\mathcal{R})),$$

where $\Gamma(t) = \max\{2n\sqrt{t}, nt\}$ for $\ell = \ell_{\text{mae}}$; $\Gamma(t) = \max\{2\sqrt{t}, t\}$ for $\ell = \ell_{\rho}$; and $\Gamma(t) = \max\{2\sqrt{nt}, t\}$ for $\ell = \ell_{\rho-\text{hinge}}$.

The theorem provides strong guarantees for the predictor-rejector surrogate losses we described in the single-stage setting. The technique used in the proof (Appendix G.2) is novel and requires careful analysis of various cases involving the pointwise best-in-class predictor and rejector. This analysis is challenging and needs to take into account the *conditional risk* and *calibration gap* (see Appendix F) of specific loss functions. The approach is substantially different from the standard scenarios examined in (Awasthi et al., 2022b), due to the simultaneous minimization of both the predictor and rejector in the abstention setting. Discussions on Theorem 2 are given in Remark 14 in Appendix B. The following is a direct consequence of Theorem 2 when \mathcal{H} and \mathcal{R} include all measurable functions, since the minimizability gaps $\mathcal{M}_{L_{abs}}$ and \mathcal{M}_{L} are then zero.

Corollary 3 (Excess error bounds for single-stage surrogates) For $\alpha = \beta$, and $\ell = \ell_{\text{mae}}$, or $\ell = \ell_{\rho}$ with $\Psi(t) = t$, or $\ell = \ell_{\rho-\text{hinge}}$ with $\Psi(t) = nt$, the following excess error bound holds for all $h \in \mathcal{H}_{\text{all}}$, $r \in \mathcal{R}_{\text{all}}$ and any distribution:

$$\mathcal{E}_{\mathsf{L}_{abs}}(h,r) - \mathcal{E}^{*}_{\mathsf{L}_{abs}}(\mathcal{H}_{all},\mathcal{R}_{all}) \leq \Gamma(\mathcal{E}_{\mathsf{L}}(h,r) - \mathcal{E}^{*}_{\mathsf{L}}(\mathcal{H}_{all},\mathcal{R}_{all})),$$

where Γ has the same form as in Theorem 2.

The corollary resolves in a positive way the open question mentioned by Ni et al. (2019). In fact it provides a stronger result since it gives non-asymptotic excess error bounds for the three abstention surrogate losses previously described. These are stronger guarantees than Bayes-consistency of these loss functions, which follow immediately by taking the limit.

It should be noted that our novel single-stage predictor-rejector surrogate losses might present some challenges for optimization. This is due to several factors: the difficulty of optimizing the mean absolute error loss (Zhang and Sabuncu, 2018) (also see Section 5), the restriction imposed by the constrained hinge loss, which is incompatible with the standard use of the softmax function in neural network hypotheses, and the non-convexity of the ρ -margin loss. However, our primary objective has been a theoretical analysis and the significance of these surrogate losses lies in their novelty and strong guarantees. As shown in Corollary 3, they are the first Bayes-consistent surrogate losses within the predictor-rejector formulation for multi-class abstention, addressing an open question in the literature (Ni et al., 2019).

4.2. Two-stage predictor-rejector surrogate losses

Here, we explore a two-stage algorithmic approach, for which we introduce surrogate losses with more flexible choices of ℓ that admit better optimization properties, and establish $(\mathcal{H}, \mathcal{R})$ -consistency bounds for them. This is a key scenario since, in practice, often a large pre-trained prediction model is already available (first stage), and retraining it would be prohibitively expensive. The problem

then consists of leaving the first stage prediction model unchanged and of subsequently learning a useful rejection model (second stage).

Let ℓ be a surrogate loss for standard multi-class classification and Φ a function like the exponential function that determines a margin-based loss $\Phi(yr(x))$ in binary classification, with $y \in \{-1, +1\}$ for a function r. We propose a two-stage algorithmic approach and a surrogate loss to minimize in the second stage: first, find a predictor h by minimizing ℓ ; second, with h fixed, find r by minimizing $\ell_{\Phi,h}$, a surrogate loss function of r for all (x, y), defined by

$$\ell_{\Phi,h}(r,x,y) = \mathbb{1}_{h(x)\neq y} \Phi(-r(x)) + c\Phi(r(x)),$$
(5)

where $t \mapsto \Phi(t)$ is a non-increasing auxiliary function upper bounding $\mathbb{1}_{t\leq 0}$. This algorithmic approach is straightforward since the first stage involves the familiar task of finding a predictor using a standard surrogate loss, such as logistic loss (or cross-entropy with softmax). The second stage is also relatively simple as h remains fixed and the form of $\ell_{\Phi,h}$ is uncomplicated, with Φ possibly being the logistic or exponential loss. It is important to underscore that the judicious selection of the indicator function in the initial term of (5) plays a crucial role in guaranteeing that the two-stage surrogate loss benefits from $(\mathcal{H}, \mathcal{R})$ -consistency bounds. If a surrogate loss is used in the first stage, this may not necessarily hold.

Note that in (5), *h* is fixed and only *r* is learned by minimizing the surrogate loss corresponding to that *h*, while in contrast both *h* and *r* are jointly learned in the abstention loss (1). We denote by $\ell_{abs,h}$ the two-stage version of the abstention loss (1) with a fixed predictor *h*, defined as: for any $r \in \mathcal{R}, x \in \mathcal{X}$ and $y \in \mathcal{Y}$:

$$\ell_{\mathrm{abs},h}(r,x,y) = \mathbb{1}_{\mathsf{h}(x)\neq y} \mathbb{1}_{r(x)>0} + c \mathbb{1}_{r(x)\leq 0}.$$
(6)

In other words, both $\ell_{\Phi,h}$ and $\ell_{abs,h}$ are loss functions of the abstention function r, while L_{abs} is a loss function of the pair $(h, r) \in (\mathcal{H}, \mathcal{R})$.

Define the binary zero-one classification loss as $\ell_{0-1}^{\text{binary}}(r, x, y) = \mathbb{1}_{y \neq \text{sign}(r(x))}$, where $\text{sign}(t) = \mathbb{1}_{t>0} - \mathbb{1}_{t\leq 0}$. As with the single-stage surrogate losses, the two-stage surrogate losses benefit from strong consistency guarantees as well. We first show that in the second stage where a predictor h is fixed, the surrogate loss function $\ell_{\Phi,h}$ benefits from \mathcal{R} -consistency bounds with respect to $\ell_{\text{abs},h}$ if, Φ admits an \mathcal{R} -consistency bound with respect to the binary zero-one loss $\ell_{0-1}^{\text{binary}}$.

Theorem 4 (\Re -consistency bounds for second-stage surrogates) Fix a predictor h. Assume that Φ admits an \Re -consistency bound with respect to $\ell_{0-1}^{\text{binary}}$. Thus, there exists a non-decreasing concave function Γ such that, for all $r \in \Re$,

$$\mathcal{E}_{\ell_{0-1}^{\mathrm{binary}}}(r) - \mathcal{E}_{\ell_{0-1}}^{*}(\mathcal{R}) + \mathcal{M}_{\ell_{0-1}^{\mathrm{binary}}}(\mathcal{R}) \leq \Gamma(\mathcal{E}_{\Phi}(r) - \mathcal{E}_{\Phi}^{*}(\mathcal{R}) + \mathcal{M}_{\Phi}(\mathcal{R})).$$

Then, the following \mathbb{R} *-consistency bound holds for all* $r \in \mathbb{R}$ *and any distribution:*

$$\mathcal{E}_{\ell_{\mathrm{abs},h}}(r) - \mathcal{E}^*_{\ell_{\mathrm{abs},h}}(\mathcal{R}) + \mathcal{M}_{\ell_{\mathrm{abs},h}}(\mathcal{R}) \leq \Gamma\Big(\Big(\mathcal{E}_{\ell_{\Phi,h}}(r) - \mathcal{E}^*_{\ell_{\Phi,h}}(\mathcal{R}) + \mathcal{M}_{\ell_{\Phi,h}}(\mathcal{R})\Big)/c\Big).$$

The proof (Appendix G.3) consists of analyzing the calibration gap of the abstention loss and second-stage surrogate loss, for a fixed predictor h. The calibration gap here is more complex than that in the standard setting as it takes into account the conditional probability, the error of that fixed predictor and the cost, and thus requires a completely different analysis. To establish \mathcal{R} -consistency bounds, we need to upper bound the calibration gap of the abstention loss by that of the surrogate

loss. However, directly working with them is rather difficult due to their complex forms. Instead, a key observation is that both forms share structural similarities with the calibration gaps in the standard classification. Motivated by the above observation, we construct an appropriate conditional distribution to transform the two calibration gaps into standard ones. Then, by applying Lemma 21 in Appendix F, we manage to leverage the \mathcal{R} -consistency bound of Φ with respect to the binary zero-one classification loss to upper bound the target calibration gap by that of the surrogate calibration gap.

We further discuss Theorem 4 in Remark 15 (Appendix B). In the special case where \mathcal{H} and \mathcal{R} are the family of all measurable functions, all the minimizability gap terms in Theorem 4 vanish. Thus, we obtain the following corollary.

Corollary 5 Fix a predictor h. Assume that Φ admits an excess error bound with respect to $\ell_{0-1}^{\text{binary}}$. Thus, there exists a non-decreasing concave functions Γ such that, for all $r \in \mathbb{R}_{all}$,

$$\mathcal{E}_{\ell_{0-1}^{\mathrm{binary}}}(r) - \mathcal{E}_{\ell_{0-1}^{\mathrm{binary}}}^*(\mathcal{R}_{\mathrm{all}}) \leq \Gamma(\mathcal{E}_{\Phi}(r) - \mathcal{E}_{\Phi}^*(\mathcal{R}_{\mathrm{all}})).$$

Then, the following excess error bound holds for all $r \in \mathbb{R}_{all}$ and any distribution:

$$\mathcal{E}_{\ell_{\mathrm{abs},h}}(r) - \mathcal{E}^*_{\ell_{\mathrm{abs},h}}(\mathcal{R}_{\mathrm{all}}) \leq \Gamma\Big(\Big(\mathcal{E}_{\ell_{\Phi,h}}(r) - \mathcal{E}^*_{\ell_{\Phi,h}}(\mathcal{R}_{\mathrm{all}})\Big)/c\Big).$$

See Remark 16 (Appendix B) for a brief discussion of Corollary 5.

We now present $(\mathcal{H}, \mathcal{R})$ -consistency bounds for the whole two-stage approach with respect to the abstention loss function L_{abs}. Let ℓ_{0-1} be the multi-class zero-one loss: $\ell_{0-1}(h, x, y) = 1_{h(x)\neq y}$. We will consider hypothesis sets \mathcal{R} that are *regular for abstention*, that is such that for any $x \in \mathcal{X}$, there exist $f, g \in \mathcal{R}$ with f(x) > 0 and $g(x) \le 0$. If \mathcal{R} is regular for abstention, then, for any x, there is an option to accept and an option to reject.

Theorem 6 ($(\mathcal{H}, \mathcal{R})$ -consistency bounds for two-stage approach) Suppose that \mathcal{R} is regular. Assume that ℓ admits an \mathcal{H} -consistency bound with respect to ℓ_{0-1} and that Φ admits an \mathcal{R} -consistency bound with respect to $\ell_{0-1}^{\text{binary}}$. Thus, there are non-decreasing concave functions Γ_1 and Γ_2 such that, for all $h \in \mathcal{H}$ and $r \in \mathcal{R}$,

$$\begin{aligned} & \mathcal{E}_{\ell_{0-1}}(h) - \mathcal{E}_{\ell_{0-1}}^{*}(\mathcal{H}) + \mathcal{M}_{\ell_{0-1}}(\mathcal{H}) \leq \Gamma_{1}(\mathcal{E}_{\ell}(h) - \mathcal{E}_{\ell}^{*}(\mathcal{H}) + \mathcal{M}_{\ell}(\mathcal{H})) \\ & \mathcal{E}_{\ell_{0-1}^{\mathrm{binary}}}(r) - \mathcal{E}_{\ell_{0-1}}^{*}(\mathcal{R}) + \mathcal{M}_{\ell_{0-1}^{\mathrm{binary}}}(\mathcal{R}) \leq \Gamma_{2}(\mathcal{E}_{\Phi}(r) - \mathcal{E}_{\Phi}^{*}(\mathcal{R}) + \mathcal{M}_{\Phi}(\mathcal{R})). \end{aligned}$$

Then, the following $(\mathcal{H}, \mathbb{R})$ *-consistency bound holds for all* $h \in \mathcal{H}$ *,* $r \in \mathbb{R}$ *and any distribution:*

$$\begin{split} \mathcal{E}_{\mathsf{L}_{\mathrm{abs}}}(h,r) &- \mathcal{E}^{*}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) \leq \Gamma_{1}(\mathcal{E}_{\ell}(h) - \mathcal{E}^{*}_{\ell}(\mathcal{H}) + \mathcal{M}_{\ell}(\mathcal{H})) \\ &+ (1+c)\Gamma_{2}\Big(\Big(\mathcal{E}_{\ell_{\Phi,h}}(r) - \mathcal{E}^{*}_{\ell_{\Phi,h}}(\mathcal{R}) + \mathcal{M}_{\ell_{\Phi,h}}(\mathcal{R})\Big)/c\Big), \end{split}$$

where the constant factors (1 + c) and $\frac{1}{c}$ can be removed when Γ_2 is linear.

In the proof (Appendix G.4), we express the pointwise estimation error term for the target abstention loss as the sum of two terms. The first term represents the pointwise estimation error of the abstention loss with a fixed h, while the second term denotes that with a fixed r^* . This proof is entirely novel and distinct from the approach used for a standard loss without abstention in (Awasthi et al., 2022b). Discussions on Theorem 6 are given in Remark 17 in Appendix B. As before, when \mathcal{H} and \mathcal{R} are the family of all measurable functions, the following result on excess error bounds holds.

Corollary 7 Assume that ℓ admits an excess error bound with respect to ℓ_{0-1} and that Φ admits an excess error bound with respect to $\ell_{0-1}^{\text{binary}}$. Thus, there are non-decreasing concave functions Γ_1 and Γ_2 such that, for all $h \in \mathcal{H}_{\text{all}}$ and $r \in \mathcal{R}_{\text{all}}$,

$$\begin{aligned} & \mathcal{E}_{\ell_{0-1}}(h) - \mathcal{E}^*_{\ell_{0-1}}(\mathcal{H}_{\text{all}}) \leq \Gamma_1(\mathcal{E}_{\ell}(h) - \mathcal{E}^*_{\ell}(\mathcal{H}_{\text{all}})) \\ & \mathcal{E}_{\ell_{0-1}^{\text{binary}}}(r) - \mathcal{E}^*_{\ell_{0-1}^{\text{binary}}}(\mathcal{R}_{\text{all}}) \leq \Gamma_2(\mathcal{E}_{\Phi}(r) - \mathcal{E}^*_{\Phi}(\mathcal{R}_{\text{all}})) \end{aligned}$$

Then, the following excess error bound holds for all $h \in \mathcal{H}_{all}$ and $r \in \mathcal{R}_{all}$ and any distribution:

$$\mathcal{E}_{\mathsf{L}_{\mathrm{abs}}}(h,r) - \mathcal{E}^{*}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H}_{\mathrm{all}},\mathcal{R}_{\mathrm{all}}) \leq \Gamma_{1}(\mathcal{E}_{\ell}(h) - \mathcal{E}^{*}_{\ell}(\mathcal{H}_{\mathrm{all}})) + (1+c)\Gamma_{2}\Big(\Big(\mathcal{E}_{\ell_{\Phi,h}}(r) - \mathcal{E}^{*}_{\ell_{\Phi,h}}(\mathcal{R}_{\mathrm{all}})\Big)/c\Big),$$

where the constant factors (1 + c) and $\frac{1}{c}$ can be removed when Γ_2 is linear.

See Remark 18 (Appendix B) for a brief discussion of Corollary 7.

These results provide a strong guarantee for surrogate losses in the two-stage setting. Additionally, while the choice of ℓ in the single-stage setting was subject to certain conditions, here, the multi-class surrogate loss ℓ can be chosen more flexibly. Specifically, it can be selected as the logistic loss (or cross-entropy with softmax), which is not only easier to optimize but is also better tailored for complex neural networks. In the second stage, the formulation is straightforward, and the choice of function Φ is flexible, resulting in a simple smooth convex optimization problem with respect to the rejector function r. Moreover, the second stage simplifies the process as h remains constant and only the rejector is optimized. This approach can enhance optimization efficiency. In Appendix C, we further highlight the significance of our findings regarding the two-stage formulation in comparison with single-stage surrogate losses.

4.3. Other advantages of the predictor-rejector formulation

In this section, we prove another advantage of our predictor-rejector surrogate losses, that is *realizable consistency* (Long and Servedio, 2013; Zhang and Agarwal, 2020). The property involves $(\mathcal{H}, \mathcal{R})$ -*realizable* distributions, under which there exist $h^* \in \mathcal{H}$ and $r^* \in \mathcal{R}$ such that $\mathcal{E}_{L_{abs}}(h^*, r^*) = 0$. The realizable distribution allows the optimal solution to abstain on points where the cost c(x) is zero.

Definition 8 L is realizable $(\mathcal{H}, \mathcal{R})$ -consistent with respect to L_{abs} if, for any $(\mathcal{H}, \mathcal{R})$ -realizable distribution, $\lim_{n \to +\infty} \mathcal{E}_{\mathsf{L}}(h_n, r_n) = \mathcal{E}^*_{\mathsf{L}}(\mathcal{H}, \mathcal{R}) \implies \lim_{n \to +\infty} \mathcal{E}_{\mathsf{L}_{abs}}(h_n, r_n) = \mathcal{E}^*_{\mathsf{L}_{abs}}(\mathcal{H}, \mathcal{R}).$

In the following, we will establish that our predictor-rejector surrogate losses, both the single-stage and two-stage variants, are realizable $(\mathcal{H}, \mathcal{R})$ -consistent when \mathcal{H} and \mathcal{R} are *closed under scaling*. A hypothesis set \mathcal{G} is closed under scaling if, $g \in \mathcal{G}$ implies that νg is also in \mathcal{G} for any $\nu \in \mathbb{R}$. We will adopt the following mild assumption for the auxiliary function Φ in the predictor-rejector surrogate losses.

Assumption 1 For any $t \in \mathbb{R}$, $\Phi(t) \ge \mathbb{1}_{t \le 0}$ and $\lim_{t \to +\infty} \Phi(t) = 0$.

In other words, Φ upper bounds the indicator function and approaches zero as t goes to infinity. We first prove a general result showing that single-stage predictor-rejector surrogate losses are realizable (\mathcal{H}, \mathcal{R})-consistent with respect to L_{abs} if the adopted ℓ satisfies Assumption 2.

Assumption 2 When $h(x, y) - \max_{y' \neq y} h(x, y') > 0$, $\lim_{\nu \to +\infty} \ell(\nu h, x, y) = 0$ and $\ell \ge \ell_{0-1}$.

The assumption implies that for a sample (x, y) for which a predictor h achieves zero error, the infimum value of the loss function ℓ is zero for any hypothesis set \mathcal{H} that includes the predictor h, provided that \mathcal{H} is closed under scaling.

Theorem 9 Assume that \mathcal{H} and \mathcal{R} are closed under scaling. Let $\Psi(0) = 0$ and Φ satisfy Assumption 1. Then, for any ℓ that satisfies Assumption 2, the following $(\mathcal{H}, \mathcal{R})$ -consistency bound holds for any $(\mathcal{H}, \mathcal{R})$ -realizable distribution, $h \in \mathcal{H}$ and $r \in \mathcal{R}$:

$$\mathcal{E}_{\mathsf{L}_{abs}}(h,r) - \mathcal{E}^*_{\mathsf{L}_{abs}}(\mathcal{H},\mathcal{R}) \leq \mathcal{E}_{\mathsf{L}}(h,r) - \mathcal{E}^*_{\mathsf{L}}(\mathcal{H},\mathcal{R})$$

The proof is included in Appendix G.5. We first establish upper bounds for $\mathcal{E}_{L}^{*}(\mathcal{H}, \mathcal{R})$ using the optimal predictor and rejector for the abstention loss, subsequently expressing the upper bound as the sum of two terms. By applying the Lebesgue dominated convergence theorem, we show that both terms vanish, and thus $\mathcal{E}_{L}^{*}(\mathcal{H}, \mathcal{R}) = 0$. It is worth noting that $(\mathcal{H}, \mathcal{R})$ -consistency bounds in Theorem 2 imply the realizable-consistency for considered loss functions. This is because under the realizable assumption, the minimizability gaps vanish for these loss functions. Nevertheless, Theorem 9 proves that a more general family of loss functions can actually achieve realizable consistency.

According to Theorem 9, for any distribution that is $(\mathcal{H}, \mathcal{R})$ -realizable, minimizing the singlestage surrogate estimation loss $\mathcal{E}_{L}(h, r) - \mathcal{E}^*_{L}(\mathcal{H}, \mathcal{R})$ results in the minimization of the abstention estimation loss $\mathcal{E}_{L_{abs}}(h, r) - \mathcal{E}^*_{L_{abs}}(\mathcal{H}, \mathcal{R})$. This suggests that the single-stage predictor-rejector surrogate loss functions are realizable $(\mathcal{H}, \mathcal{R})$ -consistent. In particular, when ℓ is chosen as $2\ell_{mae}$, $\ell = \ell_{\rho}$ and $\ell = \ell_{\rho-hinge}$ as suggested in Section 4.1, Assumption 1 is satisfied. Thus, we obtain the following corollary.

Corollary 10 Under the same assumption as in Theorem 9, for $\ell = 2\ell_{\text{mae}}$, $\ell = \ell_{\rho}$ and $\ell = \ell_{\rho-\text{hinge}}$, the single-stage predictor-rejector surrogate L is realizable $(\mathcal{H}, \mathcal{R})$ -consistent with respect to L_{abs} .

Next, we prove a similar result showing that the two-stage predictor-rejector surrogate losses are realizable $(\mathcal{H}, \mathcal{R})$ -consistent with respect to L_{abs} if the multi-class surrogate loss ℓ is realizable \mathcal{H} -consistent with respect to the multi-class zero-one loss ℓ_{0-1} when \mathcal{H} is closed under scaling.

Definition 11 We say that ℓ is realizable \mathcal{H} -consistent with respect to ℓ_{0-1} if, for any distribution such that $\mathcal{E}^*_{\ell_{0-1}}(\mathcal{H}) = 0$, $\lim_{n \to +\infty} \mathcal{E}_{\ell}(h_n) = \mathcal{E}^*_{\ell}(\mathcal{H}) \implies \lim_{n \to +\infty} \mathcal{E}_{\ell_{0-1}}(h_n) = \mathcal{E}^*_{\ell_{0-1}}(\mathcal{H}) = 0$.

Theorem 12 Assume that \mathcal{H} and \mathcal{R} are closed under scaling. Let ℓ be any multi-class surrogate loss that is realizable \mathcal{H} -consistent with respect to ℓ_{0-1} when \mathcal{H} is closed under scaling and Φ satisfies Assumption 1. Let \hat{h} be the minimizer of \mathcal{E}_{ℓ} and \hat{r} be the minimizer of $\mathcal{E}_{\ell_{\Phi,\hat{h}}}$. Then, for any $(\mathcal{H}, \mathcal{R})$ -realizable distribution, $\mathcal{E}_{\mathsf{L}_{abs}}(\hat{h}, \hat{r}) = 0$.

The proof is included in Appendix G.6. We first establish the upper bound $\mathcal{E}_{\mathsf{L}_{abs}}(\hat{h}, \hat{r}) \leq \mathcal{E}_{\ell_{\Phi,\hat{h}}}(\hat{r})$. Next, we analyze two cases: whether abstention occurs or not. By applying the Lebesgue dominated convergence theorem, we show that $\mathcal{E}_{\ell_{\Phi,\hat{h}}}(\hat{r}) = 0$ in both cases, consequently leading to $\mathcal{E}_{\mathsf{L}_{abs}}(\hat{h}, \hat{r}) = 0$. By Theorem 12, under the realizability assumption, minimizing a two-stage predictor-rejector surrogate loss leads to zero abstention loss. This implies that the two-stage predictor-rejector surrogate loss functions are also realizable $(\mathcal{H}, \mathcal{R})$ -consistent. Kuznetsov et al. (2014) prove the realizable \mathcal{H} -consistency of a broad family of multi-class surrogate losses including the logistic loss commonly used in practice. Thus, we obtain the following corollary.

Dataset	Method	Abstention loss
SVHN	(Mozannar and Sontag, 2020) (Cao et al., 2022) single-stage predictor-rejector (ℓ_{mae}) two-stage predictor-rejector	$\begin{array}{l} 1.61\% \pm 0.06\% \\ 2.16\% \pm 0.04\% \\ 2.22\% \pm 0.01\% \\ \textbf{0.94\% \pm 0.02\%} \end{array}$
CIFAR-10	(Mozannar and Sontag, 2020) (Cao et al., 2022) single-stage predictor-rejector (ℓ_{mae}) two-stage predictor-rejector	$\begin{array}{l} 4.48\% \pm 0.10\% \\ 3.62\% \pm 0.07\% \\ 3.64\% \pm 0.05\% \\ \textbf{3.31\% \pm 0.02\%} \end{array}$
CIFAR-100	(Mozannar and Sontag, 2020) (Cao et al., 2022) single-stage predictor-rejector (ℓ_{mae}) two-stage predictor-rejector	$10.40\% \pm 0.10\%$ 14.99% ± 0.01% 14.99% ± 0.01% 9.23% ± 0.03%

Table 1: Abstention loss of our predictor-rejector surrogate losses against baselines: the state-of-theart score-based abstention surrogate losses in (Mozannar and Sontag, 2020; Cao et al., 2022).

Corollary 13 Under the same assumption as in Theorem 12, for ℓ being the logistic loss, the two-stage predictor-rejector surrogate loss is realizable $(\mathcal{H}, \mathcal{R})$ -consistent with respect to L_{abs} .

Note that existing score-based abstention surrogate losses were shown to be not realizable consistent in (Mozannar et al., 2023). Recall that in Section 4.1 and Section 4.2, the $(\mathcal{H}, \mathcal{R})$ -consistency bounds guarantees (applicable to all distributions without any assumptions) indicate that both our single-stage and two-stage predictor-rejector surrogate losses are also Bayes-consistent, while it is unknown if the surrogate loss proposed by Mozannar et al. (2023) is. By combining the results from Section 4.1, Section 4.2 and Section 4.3, we demonstrate the advantages of the predictor-rejector formulation. As a by-product of our results, we address two open questions in the literature (Ni et al., 2019) and (Mozannar et al., 2023) indicated in Section 1.

5. Experiments

In this section, we present experimental results for the single-stage and two-stage predictor-rejector surrogate losses, as well as for the state-of-the-art score-based abstention surrogate losses (Mozannar and Sontag, 2020; Cao et al., 2022) on three popular datasets: SVHN (Netzer et al., 2011), CIFAR-10 and CIFAR-100 (Krizhevsky, 2009). Note that the basic confidence-based approach has already been shown in (Cao et al., 2022) to be empirically inferior to state-of-the-art score-based abstention surrogate losses. More details on the experiments are included in Appendix E.

Results. In Table 1, we report the mean and standard deviation of the abstention loss over three runs for our algorithms and the baselines. Table 1 shows that our two-stage predictor-rejector surrogate loss consistently outperforms the state-of-the-art score-based abstention surrogate losses in (Mozannar and Sontag, 2020; Cao et al., 2022) across all cases. The single-stage predictorrejector surrogate loss with ℓ set as the mean absolute error loss achieves comparable results. Our predictor-rejector surrogate losses, both the single-stage and two-stage variants, benefit from $(\mathcal{H}, \mathcal{R})$ consistency bounds and realizable $(\mathcal{H}, \mathcal{R})$ -consistency guarantees. While the optimization of mean absolute error loss is known to be challenging, as highlighted in the study by Zhang et al. (2018), our two-stage algorithm sidesteps this hurdle since it can use the more tractable logistic loss.

6. Conclusion

We presented a series of theoretical, algorithmic, and empirical results for multi-class classification with abstention. Our theoretical analysis, including proofs of $(\mathcal{H}, \mathcal{R})$ -consistency bounds and realizable $(\mathcal{H}, \mathcal{R})$ -consistency, covers both single-stage and two-stage predictor-rejector surrogate losses. These results further provide valuable tools applicable to the analysis of other loss functions in learning with abstention.

Our two-stage algorithmic approach provides practical and efficient solutions for multi-class abstention across various tasks. This approach proves particularly advantageous in scenarios where a large pre-trained prediction model is readily available, and the expense associated with retraining is prohibitive. Our empirical findings corroborate the efficacy of these algorithms, further reinforcing their practical usefulness. Additionally, our work reveals some limitations of the score-based abstention formulation, such as its inability to consistently yield optimal solutions in certain cases. In contrast, we present a collection of positive outcomes for various families of predictor-rejector surrogate loss functions. Importantly, our findings also provide resolutions to two open questions within the literature.

We believe that our analysis and the novel loss functions we introduced can guide the design of algorithms across a broad spectrum of scenarios beyond classification with abstention.

References

- Pranjal Awasthi, Natalie Frank, Anqi Mao, Mehryar Mohri, and Yutao Zhong. Calibration and consistency of adversarial surrogate losses. In Advances in Neural Information Processing Systems, 2021a.
- Pranjal Awasthi, Anqi Mao, Mehryar Mohri, and Yutao Zhong. A finer calibration analysis for adversarial robustness. *arXiv preprint arXiv:2105.01550*, 2021b.
- Pranjal Awasthi, Anqi Mao, Mehryar Mohri, and Yutao Zhong. H-consistency bounds for surrogate loss minimizers. In *International Conference on Machine Learning*, 2022a.
- Pranjal Awasthi, Anqi Mao, Mehryar Mohri, and Yutao Zhong. Multi-class H-consistency bounds. In *Advances in neural information processing systems*, 2022b.
- Pranjal Awasthi, Anqi Mao, Mehryar Mohri, and Yutao Zhong. Theoretically grounded loss functions and algorithms for adversarial robustness. In *International Conference on Artificial Intelligence and Statistics*, pages 10077–10094, 2023.
- Pranjal Awasthi, Anqi Mao, Mehryar Mohri, and Yutao Zhong. DC-programming for neural network optimizations. *Journal of Global Optimization*, pages 1–17, 2024.
- Gagan Bansal, Besmira Nushi, Ece Kamar, Eric Horvitz, and Daniel S Weld. Is the most accurate ai the best teammate? optimizing ai for teamwork. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 11405–11414, 2021.

- Peter L Bartlett and Marten H Wegkamp. Classification with a reject option using a hinge loss. *Journal of Machine Learning Research*, 9(8), 2008.
- Adbenour Bounsiar, Edith Grall, and Pierre Beauseroy. Kernel based rejection method for supervised classification. In WASET, 2007.
- Yuzhou Cao, Tianchi Cai, Lei Feng, Lihong Gu, Jinjie Gu, Bo An, Gang Niu, and Masashi Sugiyama. Generalizing consistent multi-class classification with rejection to be compatible with arbitrary losses. In *Advances in neural information processing systems*, 2022.
- Nontawat Charoenphakdee, Zhenghang Cui, Yivan Zhang, and Masashi Sugiyama. Classification with rejection based on cost-sensitive classification. In *International Conference on Machine Learning*, pages 1507–1517, 2021.
- C Chow. On optimum recognition error and reject tradeoff. *IEEE Transactions on information theory*, 16(1):41–46, 1970.
- C.K. Chow. An optimum character recognition system using decision function. IEEE T. C., 1957.
- Evgenii Chzhen, Christophe Denis, Mohamed Hebiri, and Titouan Lorieul. Set-valued classification– overview via a unified framework. *arXiv preprint arXiv:2102.12318*, 2021.
- Corinna Cortes, Giulia DeSalvo, and Mehryar Mohri. Learning with rejection. In *International Conference on Algorithmic Learning Theory*, pages 67–82, 2016a.
- Corinna Cortes, Giulia DeSalvo, and Mehryar Mohri. Boosting with abstention. In Advances in Neural Information Processing Systems, pages 1660–1668, 2016b.
- Corinna Cortes, Giulia DeSalvo, and Mehryar Mohri. Theory and algorithms for learning with rejection in binary classification. *Annals of Mathematics and Artificial Intelligence*, pages 1–39, 2023.
- Christophe Denis, Mohamed Hebiri, Boris Ndjia Njike, and Xavier Siebert. Active learning algorithm through the lens of rejection arguments. *arXiv preprint arXiv:2208.14682*, 2022.
- Bernard Dubuisson and Mylene Masson. A statistical decision rule with incomplete knowledge about classes. *Pattern recognition*, 26(1):155–165, 1993.
- Ran El-Yaniv and Yair Wiener. Active learning via perfect selective classification. Journal of Machine Learning Research, 13(2), 2012.
- Ran El-Yaniv et al. On the foundations of noise-free selective classification. *Journal of Machine Learning Research*, 11(5), 2010.
- Charles Elkan. The foundations of cost-sensitive learning. In International joint conference on artificial intelligence, pages 973–978, 2001.
- Katja Filippova. Controlled hallucinations:learning to generate faithfully from noisy data. In *Findings* of *EMNLP 2020*, 2020.

- Giorgio Fumera and Fabio Roli. Support vector machines with embedded reject option. In *ICPR*, 2002.
- Giorgio Fumera, Fabio Roli, and Giorgio Giacinto. Multiple reject thresholds for improving classification reliability. In *ICAPR*, 2000.
- Aditya Gangrade, Anil Kag, and Venkatesh Saligrama. Selective classification via one-sided prediction. In *International Conference on Artificial Intelligence and Statistics*, pages 2179–2187, 2021.
- Yonatan Geifman and Ran El-Yaniv. Selective classification for deep neural networks. In Advances in neural information processing systems, 2017.
- Yonatan Geifman and Ran El-Yaniv. Selectivenet: A deep neural network with an integrated reject option. In *International conference on machine learning*, pages 2151–2159, 2019.
- Aritra Ghosh, Himanshu Kumar, and P Shanti Sastry. Robust loss functions under label noise for deep neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, 2017.
- Yves Grandvalet, Joseph Keshet, Alain Rakotomamonjy, and Stephane Canu. Support vector machines with a reject option. In *NIPS*, 2008.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- Radu Herbei and Marten Wegkamp. Classification with reject option. Can. J. Stat., 2005.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, Toronto University, 2009.
- Vitaly Kuznetsov, Mehryar Mohri, and Umar Syed. Multi-class deep boosting. In Advances in Neural Information Processing Systems, pages 2501–2509, 2014.
- Thomas Landgrebe, David Tax, Pavel Paclik, and Robert Duin. Interaction between classification and reject performance for distance-based reject-option classifiers. *PRL*, 2005.
- Hoel Le Capitaine and Carl Frelicot. An optimum class-rejective decision rule and its evaluation. In *International Conference on Pattern Recognition*, pages 3312–3315, 2010.
- Yoonkyung Lee, Yi Lin, and Grace Wahba. Multicategory support vector machines: Theory and application to the classification of microarray data and satellite radiance data. *Journal of the American Statistical Association*, 99(465):67–81, 2004.
- Lihong Li, Michael L Littman, and Thomas J Walsh. Knows what it knows: a framework for self-aware learning. In *International conference on Machine learning*, pages 568–575, 2008.
- Phil Long and Rocco Servedio. Consistency versus realizable H-consistency for multiclass classification. In *International Conference on Machine Learning*, pages 801–809, 2013.

- Ilya Loshchilov and Frank Hutter. SGDR: Stochastic gradient descent with warm restarts. *arXiv* preprint arXiv:1608.03983, 2016.
- David Madras, Toni Pitassi, and Richard Zemel. Predict responsibly: improving fairness and accuracy by learning to defer. In *Advances in Neural Information Processing Systems*, 2018.
- Anqi Mao, Christopher Mohri, Mehryar Mohri, and Yutao Zhong. Two-stage learning to defer with multiple experts. In *Advances in Neural Information Processing Systems*, 2023a.
- Anqi Mao, Mehryar Mohri, and Yutao Zhong. H-consistency bounds: Characterization and extensions. In Advances in Neural Information Processing Systems, 2023b.
- Anqi Mao, Mehryar Mohri, and Yutao Zhong. Cross-entropy loss functions: Theoretical analysis and applications. In *International conference on Machine learning*, 2023c.
- Anqi Mao, Mehryar Mohri, and Yutao Zhong. H-consistency bounds for pairwise misranking loss surrogates. In *International conference on Machine learning*, 2023d.
- Anqi Mao, Mehryar Mohri, and Yutao Zhong. Ranking with abstention. In *ICML 2023 Workshop The Many Facets of Preference-Based Learning*, 2023e.
- Anqi Mao, Mehryar Mohri, and Yutao Zhong. Structured prediction with stronger consistency guarantees. In Advances in Neural Information Processing Systems, 2023f.
- Anqi Mao, Mehryar Mohri, and Yutao Zhong. Principled approaches for learning to defer with multiple experts. In *International Symposium on Artificial Intelligence and Mathematics*, 2024a.
- Anqi Mao, Mehryar Mohri, and Yutao Zhong. Theoretically grounded loss functions and algorithms for score-based multi-class abstention. In *International Conference on Artificial Intelligence and Statistics*, 2024b.
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. On faithfulness and factuality in abstractive summarization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1906–1919, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.173.
- Iain Melvin, Jason Weston, Christina S. Leslie, and William S. Noble. Combining classifiers for improved classification of proteins from sequence or structure. *BMCB*, 2008.
- Christopher Mohri, Daniel Andor, Eunsol Choi, Michael Collins, Anqi Mao, and Yutao Zhong. Learning to reject with a fixed predictor: Application to decontextualization. In *International Conference on Learning Representations*, 2024.
- Hussein Mozannar and David Sontag. Consistent estimators for learning to defer to an expert. In *International Conference on Machine Learning*, pages 7076–7087, 2020.
- Hussein Mozannar, Hunter Lang, Dennis Wei, Prasanna Sattigeri, Subhro Das, and David Sontag. Who should predict? exact algorithms for learning to defer to humans. In *International Conference* on Artificial Intelligence and Statistics, pages 10520–10545, 2023.

- Harikrishna Narasimhan, Wittawat Jitkrittum, Aditya Krishna Menon, Ankit Singh Rawat, and Sanjiv Kumar. Post-hoc estimators for learning to defer to an expert. In *Advances in Neural Information Processing Systems*, 2022.
- Harikrishna Narasimhan, Aditya Krishna Menon, Wittawat Jitkrittum, and Sanjiv Kumar. Learning to reject meets ood detection: Are all abstentions created equal? *arXiv preprint arXiv:2301.12386*, 2023.
- Yurii E Nesterov. A method for solving the convex programming problem with convergence rate $o(1/k^2)$. *Dokl. akad. nauk Sssr*, 269:543–547, 1983.
- Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. In Advances in Neural Information Processing Systems, 2011.
- Chenri Ni, Nontawat Charoenphakdee, Junya Honda, and Masashi Sugiyama. On the calibration of multiclass classification with rejection. In Advances in Neural Information Processing Systems, pages 2582–2592, 2019.
- Nastaran Okati, Abir De, and Manuel Rodriguez. Differentiable learning under triage. Advances in Neural Information Processing Systems, 34:9140–9151, 2021.
- Carla S Pereira and Ana Pires. On optimal reject rules and ROC curves. PRL, 2005.
- Tadeusz Pietraszek. Optimizing abstaining classifiers using ROC. In ICML, 2005.
- Nikita Puchkin and Nikita Zhivotovskiy. Exponential savings in agnostic active learning through abstention. In *Conference on Learning Theory*, pages 3806–3832, 2021.
- Maithra Raghu, Katy Blumer, Greg Corrado, Jon Kleinberg, Ziad Obermeyer, and Sendhil Mullainathan. The algorithmic automation problem: Prediction, triage, and human effort. *arXiv* preprint arXiv:1903.12220, 2019a.
- Maithra Raghu, Katy Blumer, Rory Sayres, Ziad Obermeyer, Bobby Kleinberg, Sendhil Mullainathan, and Jon Kleinberg. Direct uncertainty prediction for medical second opinions. In *International Conference on Machine Learning*, pages 5281–5290, 2019b.
- Harish G Ramaswamy, Ambuj Tewari, and Shivani Agarwal. Consistent algorithms for multiclass classification with an abstain option. *Electronic Journal of Statistics*, 12(1):530–554, 2018.
- Mark D Reid and Robert C Williamson. Composite binary losses. The Journal of Machine Learning Research, 11:2387–2422, 2010.
- Nicolas Schreuder and Evgenii Chzhen. Classification with abstention but without disparities. In Uncertainty in Artificial Intelligence, pages 1227–1236. PMLR, 2021.
- Ingo Steinwart. How to compare different loss functions and their risks. *Constructive Approximation*, 26(2):225–287, 2007.
- David MJ Tax and Robert PW Duin. Growing a multi-class classifier with a reject option. *Pattern Recognition Letters*, 29(10):1565–1570, 2008.

Francesco Tortorella. An optimal reject rule for binary classifiers. In ICAPR, 2001.

- Rajeev Verma and Eric Nalisnick. Calibrated learning to defer with one-vs-all classifiers. In *International Conference on Machine Learning*, pages 22184–22202, 2022.
- Rajeev Verma, Daniel Barrejón, and Eric Nalisnick. Learning to defer to multiple experts: Consistent surrogate losses, confidence calibration, and conformal ensembles. In *International Conference* on Artificial Intelligence and Statistics, pages 11415–11434, 2023.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. Emergent abilities of large language models. *CoRR*, abs/2206.07682, 2022.
- Yair Wiener and Ran El-Yaniv. Agnostic selective classification. In Advances in neural information processing systems, 2011.
- Yair Wiener and Ran El-Yaniv. Agnostic pointwise-competitive selective classification. *Journal of* Artificial Intelligence Research, 52:171–201, 2015.
- Yair Wiener, Steve Hanneke, and Ran El-Yaniv. A compression technique for analyzing disagreementbased active learning. J. Mach. Learn. Res., 16:713–745, 2015.
- Bryan Wilder, Eric Horvitz, and Ece Kamar. Learning to complement humans. In *International Joint Conferences on Artificial Intelligence*, pages 1526–1533, 2021.
- Ming Yuan and Marten Wegkamp. Classification methods with reject option based on convex risk minimization. *Journal of Machine Learning Research*, 11(1), 2010.
- Ming Yuan and Marten Wegkamp. SVMs with a reject option. In Bernoulli, 2011.
- Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.
- Chicheng Zhang and Kamalika Chaudhuri. The extended littlestone's dimension for learning with mistakes and abstentions. In *Conference on Learning Theory*, pages 1584–1616, 2016.
- Mingyuan Zhang and Shivani Agarwal. Bayes consistency vs. H-consistency: The interplay between surrogate loss functions and the scoring function class. In *Advances in Neural Information Processing Systems*, 2020.
- Zhilu Zhang and Mert Sabuncu. Generalized cross entropy loss for training deep neural networks with noisy labels. In *Advances in neural information processing systems*, 2018.
- Chenyu Zheng, Guoqiang Wu, Fan Bao, Yue Cao, Chongxuan Li, and Jun Zhu. Revisiting discriminative vs. generative classifiers: Theory and implications. In *International Conference on Machine Learning*, 2023.
- Yinglun Zhu and Robert Nowak. Efficient active learning with abstention. *arXiv preprint arXiv:2204.00043*, 2022.

Liu Ziyin, Zhikang Wang, Paul Pu Liang, Ruslan Salakhutdinov, Louis-Philippe Morency, and Masahito Ueda. Deep gamblers: Learning to abstain with portfolio theory. *arXiv preprint arXiv:1907.00208*, 2019.

Contents of Appendix

A	Related work	21	
B	Remarks on some key results	23	
С	C Significance of two-stage formulation compared with single-stage losses		
D	Difference between predictor-rejector and score-based formulations		
E	Experimental details	25	
F	Useful lemmas		
	F.1 Lemma 20 and proof	26	
	F.2 Lemma 21 and proof	27	
G	Proofs of main theorems	28	
	G.1 Proof of negative result for single-stage surrogates (Theorem 1)	28	
	G.2 Proof of $(\mathcal{H}, \mathcal{R})$ -consistency bounds for single-stage surrogates (Theorem 2)	29	
	G.3 Proof of R-consistency bounds for second-stage surrogates (Theorem 4)	40	
	G.4 Proof of $(\mathcal{H}, \mathcal{R})$ -consistency bounds for two-stage surrogates (Theorem 6)	41	
	G.5 Proof of realizable $(\mathcal{H}, \mathcal{R})$ -consistency bounds for single-stage surrogates (Theorem 9)	43	
	G.6 Proof of realizable $(\mathcal{H}, \mathcal{R})$ -consistency for two-stage surrogates (Theorem 12)	44	

Appendix A. Related work

Several broad methods for learning with abstention can be distinguished in the literature: confidencebased methods, which consist of abstaining when the score returned by a pre-trained model falls below some threshold (Chow, 1957, 1970; Bartlett and Wegkamp, 2008; Yuan and Wegkamp, 2010, 2011; Ramaswamy et al., 2018; Ni et al., 2019); selective classification, which analyzes a set-up with a *predictor* and a *selector* and defines a selection risk or loss normalized by the expected selection or coverage (El-Yaniv et al., 2010; Wiener and El-Yaniv, 2011; El-Yaniv and Wiener, 2012; Wiener and El-Yaniv, 2015; Geifman and El-Yaniv, 2017, 2019); a predictor-rejector formulation, which is based on learning both a *predictor* and a *rejector*, each from a different family of functions, and that takes into account explicitly the cost of abstention (Cortes et al., 2016a,b, 2023; Mohri et al., 2024); and a more recent score-based formulation that consists of augmenting the multi-class categories with a rejection label and of abstaining when the score assigned to the rejection label is the highest (Mozannar and Sontag, 2020; Cao et al., 2022; Mao et al., 2024b). Another problem closely related to abstention is that of *deferring* to an alternative model, or even to a human in some instances. This can also be considered as a special case of the general abstention scenario and tackled in a similar way (Madras et al., 2018; Raghu et al., 2019a; Mozannar and Sontag, 2020; Okati et al., 2021; Wilder et al., 2021; Verma and Nalisnick, 2022; Narasimhan et al., 2022; Verma et al., 2023; Mao et al., 2023a, 2024a).

The study of confidence-based methods was initiated by Chow (1957, 1970) who explored the trade-off between error rate and rejection rate, and also presented an analysis of the Bayes optimal decision in this context. Later, Fumera et al. (2000) proposed a multiple thresholds rule for situations where a posteriori probabilities were impacted by errors. Tortorella (2001) introduced an optimal rejection rule for binary classifiers, relying on the Receiver Operating Characteristic (ROC) curve. Additionally, Pereira and Pires (2005) compared their methodology with that of Chow (1970). Numerous publications have proposed various rejection techniques to reduce the misclassification rate, though without theoretical analysis (Fumera and Roli, 2002; Pietraszek, 2005; Bounsiar et al., 2007; Landgrebe et al., 2005; Melvin et al., 2008). Herbei and Wegkamp (2005) examined classification with a rejection option involving a cost and provided excess error bounds for these ternary functions. Bartlett and Wegkamp (2008) developed a loss function for this scenario that takes into account the abstention cost c. They proposed learning a predictor using a *double hinge* loss and demonstrated its consistency benefits. This approach has been further explored in several subsequent publications (Grandvalet et al., 2008; Yuan and Wegkamp, 2010, 2011). Ramaswamy et al. (2018) further examined confidence-based abstention in multi-class classification, showing that certain multi-class hinge loss formulations and a newly constructed polyhedral binary-encoded predictions (BEP) surrogate loss are Bayes-consistent. Charoenphakdee et al. (2021) suggested a cost-sensitive approach for multi-class abstention by breaking down the multi-class problem into multiple binary cost-sensitive classification problems (Elkan, 2001). They introduced a family of cost-sensitive one-versus-all surrogate losses, which are Bayes-consistent in that context. Narasimhan et al. (2023) investigated the connection between learning with abstention and out-of-distribution detection. They developed a plug-in method aimed at approximating the Bayes-optimal classifier, and demonstrated its application in the context of learning an out-of-distribution (OOD) aware classifier.

Selective classification methods were introduced by El-Yaniv et al. (2010) who investigated the trade-off between classifier coverage and accuracy. In a follow-up study, Wiener and El-Yaniv (2011) developed a strategy for learning a specific kind of selective classification called weakly optimal,

which has a diminishing rejection rate under certain Bernstein-type conditions. Many successful connections to selective classification have been established, including active learning (El-Yaniv and Wiener, 2012; Wiener et al., 2015; Wiener and El-Yaniv, 2015; Puchkin and Zhivotovskiy, 2021; Denis et al., 2022; Zhu and Nowak, 2022), multi-class rejection (Tax and Duin, 2008; Dubuisson and Masson, 1993; Le Capitaine and Frelicot, 2010), reinforcement learning (Li et al., 2008), online learning (Zhang and Chaudhuri, 2016), modern confidence-based rejection methods (Geifman and El-Yaniv, 2017), neural network architectures (Geifman and El-Yaniv, 2019), loss functions based on gambling's doubling rate (Ziyin et al., 2019), disparity-free approaches (Schreuder and Chzhen, 2021), and the abstention problem in a "confidence set" framework (Gangrade et al., 2021; Chzhen et al., 2021).

The predictor-rejector formulation was advocated by Cortes, DeSalvo, and Mohri (2016a) who contended that confidence-based abstention is generally suboptimal, unless the learned predictor is the Bayes classifier. They demonstrated that, in most cases, no threshold-based abstention can achieve the desired outcome. They proposed a new abstention framework that involves learning both a predictor h and a rejector r simultaneously, which can generally differ from a threshold-based function. They defined a predictor-rejector formulation loss function for the pair (h, r), considering the abstention cost c. The authors provided Rademacher complexity-based generalization bounds for this learning problem and proposed various surrogate loss functions for the binary classification abstention loss. They demonstrated that these surrogate losses offered consistency guarantees and developed algorithms based on these surrogate losses, which empirically outperformed confidence-based abstention benchmarks. This work led to several follow-up studies, including a theoretical and algorithmic investigation of boosting with abstention (Cortes et al., 2016b) and an analysis of extending the results to a multi-class setting (Ni et al., 2019). These authors acknowledged the difficulty in designing calibrated or Bayes-consistent surrogate losses based on the predictor-rejector abstention by Cortes et al. (2016a) and left it as an open question.

Mozannar and Sontag (2020) introduced an alternative *score-based formulation* for multi-class abstention. In this approach, besides the standard scoring functions associated with each label, a new scoring function is linked to a new rejection label. Rejection occurs when the score assigned to the rejection label exceeds other scores, implicitly defining the rejector through this specific rule. The authors proposed a surrogate loss for their method based on cross-entropy (logistic loss with softmax applied to neural network outputs), which they demonstrated to be Bayes-consistent. Building upon this work, Cao et al. (2022) presented a more comprehensive collection of Bayes-consistent surrogate losses for the score-based formulation. These surrogate losses can be constructed using any consistent loss function for the standard multi-class classification problem. More recently, Mao et al. (2024b) gave an extensive analysis of surrogate losses for the score-based formulation supported by \mathcal{H} -consistency bounds.

The problem of learning to defer is closely related to our study and can be seen as a specific instance of learning with abstention. Several recent publications have investigated this problem, including (Madras et al., 2018; Raghu et al., 2019a,b; Mozannar and Sontag, 2020; Okati et al., 2021; Wilder et al., 2021; Bansal et al., 2021; Verma and Nalisnick, 2022; Narasimhan et al., 2022; Verma et al., 2023). Confidence-based methods for deferral decisions were examined by Raghu et al. (2019b); Wilder et al. (2021); Bansal et al. (2021), but these methods may not be optimal for low capital models (Cortes et al., 2016a). To address this issue, Mozannar and Sontag (2020) proposed a cost-sensitive logistic loss and Verma and Nalisnick (2022) proposed a cost-sensitive one-versus-all proper composite loss (Reid and Williamson, 2010), both in the score-based formulation. Verma

et al. (2023) generalized the surrogate loss in (Verma and Nalisnick, 2022) to the deferral setting with multiple experts. Recently, Narasimhan et al. (2022) pointed out that the existing surrogate losses for learning to defer (Mozannar and Sontag, 2020; Verma and Nalisnick, 2022) may underfit in some practical settings and proposed a post-hoc correction for these loss functions.

We will be particularly interested in the predictor-rejector formulation, which explicitly models the cost of abstention. The selective classification of El-Yaniv et al. (2010) is also interesting, but it does not explicitly factor in the cost c and is based on a distinct objective. Confidence-based methods are also very natural and straightforward, but they may fail when the predictor is not calibrated, a property that often does not hold. Additionally, they have been shown to be suboptimal when the predictor differs from the Bayes classifier (Cortes et al., 2016a). The score-based formulation (Mozannar and Sontag, 2020) admits very common properties and also explicitly takes into account the rejection cost c. We will compare the predictor-rejector formulation with the score-based one. We will show via an example that the predictor-rejector is more natural in some instances and will also compare the two formulations in our experiments. We further elaborate on the difference between the two formulations in Appendix D.

Appendix B. Remarks on some key results

Remark 14 When the best-in-class error coincides with the Bayes error $\mathcal{E}_{L}^{*}(\mathcal{H}, \mathcal{R}) = \mathcal{E}_{L}^{*}(\mathcal{H}_{all}, \mathcal{R}_{all})$, the minimizability gaps $\mathcal{M}_{L}(\mathcal{H}, \mathcal{R})$ vanish. In those cases, the $(\mathcal{H}, \mathcal{R})$ -consistency bound in Theorem 2 guarantees that when the surrogate estimation error $\mathcal{E}_{L}(h, r) - \mathcal{E}_{L}^{*}(\mathcal{H}, \mathcal{R})$ is optimized up to ϵ , the estimation error of the abstention loss $\mathcal{E}_{L_{abs}}(h, r) - \mathcal{E}_{L_{abs}}^{*}(\mathcal{H}, \mathcal{R})$ is upper bounded by $\Gamma(\epsilon)$. For all the three loss functions, when ϵ is sufficiently small, the dependence of Γ on ϵ exhibits a square root relationship. However, if this is not the case, the dependence becomes linear. Note that the dependence is subject to the number of classes n for $\ell = \ell_{mae}$ and $\ell = \ell_{\rho-hinge}$.

Remark 15 When the best-in-class error coincides with the Bayes error $\mathcal{E}_{\ell}^{*}(\mathcal{R}) = \mathcal{E}_{\ell}^{*}(\mathcal{R}_{all})$ for $\ell = \ell_{\Phi,h}$ and $\ell = \ell_{abs,h}$, the minimizability gaps $\mathcal{M}_{\ell_{\Phi,h}}(\mathcal{R})$ and $\mathcal{M}_{\ell_{abs,h}}(\mathcal{R})$ vanish. In those cases, the \mathcal{R} -consistency bound in Theorem 4 guarantees that when the surrogate estimation error $\mathcal{E}_{\ell_{\Phi,h}}(r) - \mathcal{E}_{\ell_{abs,h}}^{*}(\mathcal{R})$ is optimized up to ϵ , the target estimation error $\mathcal{E}_{\ell_{abs,h}}(r) - \mathcal{E}_{\ell_{abs,h}}^{*}(\mathcal{R})$ is upper bounded by $\Gamma(\frac{\epsilon}{\epsilon})$.

Remark 16 Corollary 5 shows that $\ell_{\Phi,h}$ admits an excess error bound with respect to $\ell_{abs,h}$ with functional form $\Gamma(\frac{.}{c})$ when Φ admits an excess error bound with respect to ℓ_{0-1}^{binary} with functional form $\Gamma(\cdot)$.

Remark 17 Note that the minimizability gaps vanish when \mathcal{H} and \mathcal{R} are families of all measurable functions or when they include the Bayes predictor and rejector. In their absence, Theorem 6 shows that if the estimation prediction loss $(\mathcal{E}_{\ell}(h) - \mathcal{E}^*_{\ell}(\mathcal{H}))$ is reduced to ϵ_1 and the estimation rejection loss $(\mathcal{E}_{\ell_{\Phi,h}}(r) - \mathcal{E}^*_{\ell_{\Phi,h}}(\mathcal{R}))$ to ϵ_2 , then the abstention estimation loss $(\mathcal{E}_{L_{abs}}(h, r) - \mathcal{E}^*_{L_{abs}}(\mathcal{H}, \mathcal{R}))$ is, up to constant factors, bounded by $\Gamma_1(\epsilon_1) + \Gamma_2(\epsilon_2)$.

Remark 18 Corollary 7 shows that the proposed two-stage approach admits an excess error bound with respect to L_{abs} with functional form $\Gamma_1(\cdot) + (1+c)\Gamma_2(\frac{\cdot}{c})$ when ℓ admits an excess error bound with respect to ℓ_{0-1} with functional form $\Gamma_1(\cdot)$ and Φ admits an excess error bound with respect to ℓ_{0-1}^{binary} with functional form $\Gamma_2(\cdot)$.

Appendix C. Significance of two-stage formulation compared with single-stage losses

Here, we wish to further highlight the significance of our findings regarding the two-stage formulation. The $(\mathcal{H}, \mathcal{R})$ -consistency bounds we established for this scenario directly motivate an algorithm for a crucial scenario. As already indicated, in applications, often a prediction function is already available and has been trained using a standard loss function such as cross-entropy. Training may take days or months for some large models. The cost of a one-stage approach, which involves "retraining" to find a pair (h, r) with a new h, can thus be prohibitive. Instead, we demonstrate that a rejector r can be learned using a suitable surrogate loss function based on the existing predictor h and that the solution formed by the existing h and this rejector r benefits from $(\mathcal{H}, \mathcal{R})$ -consistency bounds.

Both our one-stage and two-stage solutions using our surrogate losses benefit from strong $(\mathcal{H}, \mathcal{R})$ consistency bounds: in the limit of large samples, both methods, one-stage and two-stage converge to the same joint minimizer of the target abstention loss. However, as already emphasized, the two-stage approach is advantageous in some scenarios where a predictor h is already available. Moreover, the two-stage solution is more beneficial from the optimization point of view: the first-stage optimization can be standard and based on say cross-entropy, and the second stage is based on a loss function (5) that is straightforward to minimize. In contrast, the one-stage minimization with the MAE loss is known to be more difficult, see (Zhang and Sabuncu, 2018). In Section 5, our empirical results show a more favorable performance for the two-stage solution, which we believe reflects this difference in optimization.

Appendix D. Difference between predictor-rejector and score-based formulations

Here, we hope to further emphasize our contributions by pointing out that the score-based formulation does not offer a direct loss function applicable to the predictor-rejector formulation. It is important to emphasize that the hypothesis set used in the score-based setting constitutes a subset of real-valued functions defined over $\mathfrak{X} \times \tilde{\mathfrak{Y}}$, where $\tilde{\mathfrak{Y}}$ is the original label set \mathfrak{Y} augmented with an additional label corresponding to rejection. In contrast, the predictor-rejector function involves selecting a predictor *h* out of a collection of real-valued functions defined over $\mathfrak{X} \times \mathfrak{Y}$ and a rejector *r* from a sub-family of real-valued functions defined over \mathfrak{X} . Thus, the hypothesis sets in these two frameworks differ entirely. Consequently, a score-based loss function cannot be directly applied to the hypothesis set of the predictor-rejector formulation.

One can instead, given the hypothesis sets \mathcal{H} and \mathcal{R} for the predictor and rejector functions in the predictor-rejector formulation, define a distinct hypothesis set $\tilde{\mathcal{H}}$ of real-valued functions defined over $\mathfrak{X} \times \tilde{\mathcal{Y}}$. Functions $\tilde{h} \in \tilde{\mathcal{H}}$ are defined from a pair $(h, r) \in \mathcal{H} \times \mathcal{R}$. The score-based loss function for $\tilde{h} \in \tilde{\mathcal{H}}$ then coincides with the predictor-rejector loss of (h, r). However, the family $\tilde{\mathcal{H}}$ is complex. As pointed out in Section 3, for instance, when \mathcal{H} and \mathcal{R} are families of linear functions, $\tilde{\mathcal{H}}$ is not linear and is more complex.

Moreover, there is a non-trivial coupling relating the scoring functions defined for the rejection label (n + 1) and other scoring functions, while such a coupling is not present in the standard score-based formulation. This makes it more difficult to minimize the loss function $\tilde{\ell}(\tilde{h}, x, y)$. Indeed, the minimization problem requires that the constraint $\tilde{h}(x, n+1) = \max_{y \in \mathcal{Y}} \tilde{h}(x, y) - r(x)$ be satisfied. This constraint relates the first *n* scoring functions $\tilde{h}(\cdot, y)$, $y \in \mathcal{Y}$, to the last scoring function $\tilde{h}(\cdot, n+1)$, via a maximum operator (and the function *r*). The constraint is non-differentiable and non-convex, which makes the minimization problem more challenging. This augmented complexity and coupling fundamentally differentiate the two formulations. Our counterexample in Section 3 underscores this intrinsic distinction between these two formulations: While the predictor-rejector formulation can easily handle certain instances, the score-based framework falls short to tackle them unless a more complex hypothesis set is adopted. This key difference between the two formulations is also the underlying reason for the historical difficulty in devising a consistent surrogate loss function for the predictor-rejector formulation within the standard multiclass setting, while the task has been comparatively more straightforward within the score-based formulation.

It is important to highlight that our novel families of predictor-rejector surrogate losses, alongside similar variants, establish the first Bayes-consistent and realizable consistent surrogate losses within the predictor-rejector formulation and they address two previously open questions in the literature (Ni et al., 2019) and (Mozannar et al., 2023) (see Section 4). Moreover, they outperform the state-of-the-art surrogate losses found in the score-based formulation (see Section 5). This underscores both the innovative nature and the significant contribution of our work.

In the following section, we will further showcase the advantages of the predictor-rejector formulation through empirical evidence.

Appendix E. Experimental details

Setup. We adopt ResNet-34 (He et al., 2016), a residual network with 34 convolutional layers, for SVHN and CIFAR-10, and WRN-28-10 (Zagoruyko and Komodakis, 2016), a residual network with 28 convolutional layers and a widening factor of 10, for CIFAR-100 both with ReLU activations (Zagoruyko and Komodakis, 2016). We train for 200 epochs using Stochastic Gradient Descent (SGD) with Nesterov momentum (Nesterov, 1983) following the cosine decay learning rate schedule (Loshchilov and Hutter, 2016) of an initial learning rate 0.1. During the training, the batch size is set to 1,024 and the weight decay is 1×10^{-4} . Except for SVHN, we adopt the standard data augmentation: a four pixel padding with 32×32 random crops and random horizontal flips.

We compare with a score-based surrogate loss proposed in (Mozannar and Sontag, 2020) based on cross-entropy and a score-based surrogate loss used in (Cao et al., 2022) based on generalized cross-entropy (Zhang and Sabuncu, 2018). For our single-stage predictor-rejector surrogate loss, we set ℓ to be the mean absolute error loss ℓ_{mae} since the constrained hinge loss imposes a restriction incompatible with the standard use of the softmax function with neural network hypotheses, and the ρ -margin loss is non-convex. For our two-stage predictor-rejector surrogate loss, we first use standard training with the logistic loss to learn a predictor h^* , and then in the second stage, optimize the loss function ℓ_{Φ,h^*} with $\Phi(t) = \exp(-t)$ to learn a rejector. We set the cost c to 0.03 for SVHN, 0.05 for CIFAR-10 and 0.15 for CIFAR-100. We observe that the performance remains close for other neighboring values of c. We highlight this particular choice of cost because a cost value that is not too far from the best-in-class zero-one classification loss encourages in practice a reasonable amount of input instances to be abstained.

Metrics. We use as evaluation metrics the average abstention loss, L_{abs} for predictor-rejector surrogate losses and \tilde{L}_{abs} for score-based abstention surrogate losses, which share the same semantic meaning. It's important to emphasize that the two abstention losses, L_{abs} and \tilde{L}_{abs} are indeed the same metric, albeit tailored for two distinct formulations. Consequently, their average numerical values can be directly compared. Note that both L_{abs} and \tilde{L}_{abs} account for the zero-one misclassification error when the sample is accepted, and the cost when the sample is rejected. The reason they are

Table 2: Abstention loss, zero-one misclassification error on the accepted data and rejection ratio of our predictor-rejector surrogate losses against baselines: the state-of-the-art score-based abstention surrogate losses in (Mozannar and Sontag, 2020; Cao et al., 2022) on CIFAR-10.

Method	Abstention loss	Misclassification error	Rejection ratio
(Mozannar and Sontag, 2020)	$4.48\% \pm 0.10\%$	$4.30\% \pm 0.14\%$	$25.99\% \pm 0.41\%$
(Cao et al., 2022)	$3.62\% \pm 0.07\%$	$3.08\% \pm 0.10\%$	$28.27\% \pm 0.18\%$
single-stage predictor-rejector (ℓ_{mae})	$3.64\% \pm 0.05\%$	$3.54\% \pm 0.05\%$	$17.21\%\pm0.22\%$
two-stage predictor-rejector	$3.31\%\pm0.02\%$	$2.69\%~\pm~0.05\%$	$22.83\% \pm 0.21\%$

adapted to the two formulations is due to the difference in the rejection method: $r(x) \le 0$ in the predictor-rejector formulation and $\tilde{h}(x) = n + 1$ in the score-based abstention formulation. It should also be noted that the abstention loss serves as a comprehensive metric that integrates the rejection ratio and zero-one misclassification error on the accepted data, thereby providing a singular, fair ground for comparison in Table 1. Nevertheless, we include all three metrics in Table 2 as a detailed comparison.

Appendix F. Useful lemmas

We first introduce some notation before presenting a lemma that will be used in our proofs. Recall that we denote by $p(x,y) = \mathcal{D}(Y = y | X = x)$ the conditional probability of Y = y given X = x. Thus, the generalization error for a general abstention surrogate loss can be rewritten as $\mathcal{E}_{L}(h,r) = \mathbb{E}_{X}[\mathcal{C}_{L}(h,r,x)]$, where $\mathcal{C}_{L}(h,r,x)$ is the conditional risk of L, defined by

$$\mathcal{C}_{\mathsf{L}}(h,r,x) = \sum_{y \in \mathcal{Y}} p(x,y) \mathsf{L}(h,r,x,y).$$

We denote by $C^*_L(\mathcal{H}, \mathcal{R}, x) = \inf_{h \in \mathcal{H}, r \in \mathcal{R}} C_L(h, r, x)$ the best-in-class conditional risk of L. Then, the minimizability gap can be rewritten as follows:

$$\mathcal{M}_{\mathsf{L}}(\mathcal{H},\mathcal{R}) = \mathcal{E}_{\mathsf{I}}^{*}(\mathcal{H},\mathcal{R}) - \mathbb{E}_{X}[\mathcal{C}_{\mathsf{I}}^{*}(\mathcal{H},\mathcal{R},x)].$$

We further refer to $\mathcal{C}_{\mathsf{L}}(h, r, x) - \mathcal{C}_{\mathsf{L}}^{*}(\mathcal{H}, \mathcal{R}, x)$ as the calibration gap and denote it by $\Delta \mathcal{C}_{\mathsf{L},\mathcal{H},\mathcal{R}}(h, r, x)$. We first prove a lemma on the calibration gap of the general abstention loss. For any $x \in \mathcal{X}$, we define the set of labels generated by hypotheses in \mathcal{H} as $\mathsf{H}(x) := \{\mathsf{h}(x) : h \in \mathcal{H}\}$. We will consider hypothesis sets \mathcal{R} which are *regular for abstention*.

Definition 19 (Regularity for Abstention) We say that a hypothesis set \mathbb{R} is regular for abstention *if for any* $x \in \mathcal{X}$ *, there exist* $f, g \in \mathbb{R}$ such that f(x) > 0 and $g(x) \le 0$.

In other words, if \mathcal{R} is regular for abstention, then, for any instance x, there is an option to accept and an option to reject.

F.1. Lemma 20 and proof

The following lemma characterizes the calibration gap of the predictor-rejector abstention.

Lemma 20 Assume that \Re is regular for abstention. For any $x \in \chi$, the minimal conditional L_{abs} -risk and the calibration gap for L_{abs} can be expressed as follows:

$$\begin{aligned} & \mathcal{C}^*_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H}, \mathcal{R}, x) = 1 - \max\left\{ \max_{y \in \mathsf{H}(x)} p(x, y), 1 - c \right\}, \\ & \Delta \mathcal{C}_{\mathsf{L}_{\mathrm{abs}}, \mathcal{H}, \mathcal{R}}(h, r, x) = \begin{cases} \max\{\max_{y \in \mathsf{H}(x)} p(x, y), 1 - c\} - p(x, \mathsf{h}(x)) & r(x) > 0\\ \max\{\max_{y \in \mathsf{H}(x)} p(x, y) - 1 + c, 0\} & r(x) \le 0. \end{cases} \end{aligned}$$

Proof By the definition, the conditional L_{abs} -risk can be expressed as follows:

$$C_{\mathsf{L}_{abs}}(h,r,x) = \sum_{y \in \mathcal{Y}} p(x,y) \mathbb{1}_{\mathsf{h}(x) \neq y} \mathbb{1}_{r(x)>0} + c \mathbb{1}_{r(x)\leq 0} = \begin{cases} 1 - p(x,\mathsf{h}(x)) & \text{if } r(x) > 0\\ c & \text{if } r(x) \leq 0. \end{cases}$$
(7)

Since \mathcal{R} is regular for abstention, the minimal conditional L_{abs}-risk can be expressed as follows:

$$\mathcal{C}^*_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R},x) = 1 - \max\left\{\max_{y \in \mathsf{H}(x)} p(x,y), 1 - c\right\},\$$

which proves the first part of the lemma. By the definition of the calibration gap, we have

$$\begin{split} \Delta \mathcal{C}_{\mathsf{L}_{\mathrm{abs}},\mathcal{H},\mathcal{R}}(h,r,x) &= \mathcal{C}_{\mathsf{L}_{\mathrm{abs}}}(h,r,x) - \mathcal{C}_{\mathsf{L}_{\mathrm{abs}}}^{*}(\mathcal{H},\mathcal{R},x) \\ &= \begin{cases} \max\{\max_{y \in \mathsf{H}(x)} p(x,y), 1-c\} - p(x,\mathsf{h}(x)) & \text{if } r(x) > 0\\ \max\{\max_{y \in \mathsf{H}(x)} p(x,y) - 1 + c, 0\} & \text{if } r(x) \le 0, \end{cases} \end{split}$$

which completes the proof.

F.2. Lemma 21 and proof

The following lemma would be useful in the proofs for two-stage surrogate losses.

Lemma 21 Assume that the following \mathbb{R} -consistency bound holds for all $r \in \mathbb{R}$ and any distribution,

$$\mathcal{E}_{\ell_{0-1}}(r) - \mathcal{E}^*_{\ell_{0-1}}(\mathcal{R}) + \mathcal{M}_{\ell_{0-1}}(\mathcal{R}) \leq \Gamma(\mathcal{E}_{\Phi}(r) - \mathcal{E}^*_{\Phi}(\mathcal{R}) + \mathcal{M}_{\Phi}(\mathcal{R}))$$

Then, for any $p_1, p_2 \in [0, 1]$ such that $p_1 + p_2 = 1$ and $x \in \mathcal{X}$, we have

$$p_{1}1_{r(x)>0} + p_{2}1_{r(x)\leq0} - \inf_{r\in\mathcal{R}} (p_{1}1_{r(x)>0} + p_{2}1_{r(x)\leq0})$$

$$\leq \Gamma \Big(p_{1}\Phi(-r(x)) + p_{2}\Phi(r(x)) - \inf_{r\in\mathcal{R}} (p_{1}\Phi(-r(x)) + p_{2}\Phi(r(x))) \Big)$$

Proof For any $x \in \mathcal{X}$, consider a distribution δ_x that concentrates on that point. Let $p_1 = \mathbb{P}(y = -1 | x)$ and $p_2 = \mathbb{P}(y = +1 | x)$. Then, by definition, $\mathcal{E}_{\ell_{0-1}}(r) - \mathcal{E}^*_{\ell_{0-1}}(\mathcal{R}) + \mathcal{M}_{\ell_{0-1}}(\mathcal{R})$ can be expressed as

$$\mathcal{E}_{\ell_{0-1}}(r) - \mathcal{E}_{\ell_{0-1}}^*(\mathcal{R}) + \mathcal{M}_{\ell_{0-1}}(\mathcal{R}) = p_1 \mathbf{1}_{r(x)>0} + p_2 \mathbf{1}_{r(x)\leq 0} - \inf_{r \in \mathcal{R}} \left(p_1 \mathbf{1}_{r(x)>0} + \mathbf{1}_{r(x)\leq 0} \right).$$

Similarly, $\mathcal{E}_{\Phi}(r) - \mathcal{E}_{\Phi}^{*}(\mathcal{R}) + \mathcal{M}_{\Phi}(\mathcal{R})$ can be expressed as

$$\mathcal{E}_{\Phi}(r) - \mathcal{E}_{\Phi}^{*}(\mathcal{R}) + \mathcal{M}_{\Phi}(\mathcal{R}) = p_{1}\Phi(-r(x)) + p_{2}\Phi(r(x)) - \inf_{r \in \mathcal{R}}(p_{1}\Phi(-r(x)) + p_{2}\Phi(r(x))).$$

Since the \mathcal{R} -consistency bound holds by the assumption, we complete the proof.

Appendix G. Proofs of main theorems

G.1. Proof of negative result for single-stage surrogates (Theorem 1)

Theorem 1 (Negative result for single-stage surrogates) Assume that \mathcal{H} is symmetric and complete, and that \mathcal{R} is complete. If there exists $x \in \mathcal{X}$ such that $\inf_{h \in \mathcal{H}} \mathbb{E}_y[\ell(h, X, y) | X = x] \neq \frac{\beta \Psi(1 - \max_{y \in \mathcal{Y}} p(x, y))}{\alpha}$, then, there does not exist a non-decreasing function $\Gamma: \mathbb{R}_+ \to \mathbb{R}_+$ with the property $\lim_{t\to 0^+} \Gamma(t) = 0$ such that the following $(\mathcal{H}, \mathcal{R})$ -consistency bound holds: for all $h \in \mathcal{H}, r \in \mathcal{R}$, and any distribution,

$$\mathcal{E}_{\mathsf{L}_{\mathrm{abs}}}(h,r) - \mathcal{E}^{*}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) \leq \Gamma(\mathcal{E}_{\mathsf{L}}(h,r) - \mathcal{E}^{*}_{\mathsf{L}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}}(\mathcal{H},\mathcal{R})).$$

Proof We prove by contradiction. Assume that the bound holds with some non-decreasing function Γ with $\lim_{t\to 0^+} \Gamma(t) = 0$, then, for all $h \in \mathcal{H}$, $r \in \mathcal{R}$, and any distribution, $\mathcal{E}_{L}(h, r) - \mathcal{E}_{L}^{*}(\mathcal{H}, \mathcal{R}) + \mathcal{M}_{L}(\mathcal{H}, \mathcal{R}) \to 0 \implies \mathcal{E}_{L_{abs}}(h, r) - \mathcal{E}_{L_{abs}}^{*}(\mathcal{H}, \mathcal{R}) + \mathcal{M}_{L_{abs}}(\mathcal{H}, \mathcal{R}) \to 0$. This further implies that for any $x \in \mathcal{X}$, the minimizer h^* and r^* of $\mathcal{C}_{L}(h, r, x)$ within \mathcal{H} and \mathcal{R} also achieves the minimum of $\mathcal{C}_{L_{abs}}(h, r, x)$ within \mathcal{H} and \mathcal{R} . When \mathcal{H} is symmetric and complete, we have $H(x) = \mathcal{Y}$. Since \mathcal{R} is complete, \mathcal{R} is regular for abstention. By Lemma 20, h^* and r^* need to satisfy the following conditions:

$$p(x,\mathsf{h}^*(x)) = \max_{y \in \mathcal{Y}} p(x,y), \quad \operatorname{sign}(r^*(x)) = \operatorname{sign}\left(\max_{y \in \mathcal{Y}} p(x,y) - (1-c)\right). \tag{8}$$

Next, we will show that (8) contradicts the assumption that there exists $x \in \mathcal{X}$ such that

$$\inf_{h \in \mathcal{H}} \mathbb{E}[\ell(h, X, y) \mid X = x] \neq \frac{\beta \Psi (1 - \max_{y \in \mathcal{Y}} p(x, y))}{\alpha}.$$

By definition, the conditional L-risk can be expressed as follows:

$$\mathcal{C}_{\mathsf{L}}(h,r,x) = \exp(\alpha r(x))\mathbb{E}_{y}[\ell(h,X,y) \mid X = x] + \Psi(c)\exp(-\beta r(x)).$$

Then, for any fixed $r \in \mathbb{R}$, $\alpha > 0$ and $\beta > 0$, we have

$$\inf_{h \in \mathcal{H}} \mathcal{C}_{\mathsf{L}}(h, r, x) = \exp(\alpha r(x)) \inf_{h \in \mathcal{H}} \mathbb{E}_{y}[\ell(h, X, y) \mid X = x] + \Psi(c) \exp(-\beta r(x)) \coloneqq \mathcal{F}(r(x)).$$

By taking the derivative, we obtain

$$\mathcal{F}'(r(x)) = \alpha \exp(\alpha r(x)) \inf_{h \in \mathcal{H}} \mathbb{E}_y[\ell(h, X, y) \mid X = x] - \beta \Psi(c) \exp(-\beta r(x)).$$

Since $\mathcal{F}(r(x))$ is convex with respect to r(x), we know that $\mathcal{F}'(r(x))$ is non-decreasing with respect to r(x). The equation (8) implies that $\mathcal{F}(r(x))$ is attained at $r^*(x)$ such that $\operatorname{sign}(r^*(x)) = \operatorname{sign}(\max_{y \in \mathcal{Y}} p(x, y) - (1 - c))$. Thus, we have

$$\max_{y \in \mathcal{Y}} p(x, y) \ge (1 - c) \implies r^*(x) \ge 0 \implies \mathcal{F}'(0) \le \mathcal{F}'(r^*(x)) = 0$$
$$\max_{y \in \mathcal{Y}} p(x, y) < (1 - c) \implies r^*(x) < 0 \implies \mathcal{F}'(0) \ge \mathcal{F}'(r^*(x)) = 0.$$

This implies that

$$\alpha \inf_{h \in \mathcal{H}} \mathbb{E}_{y}[\ell(h, X, y) \mid X = x] - \beta \Psi(c) \leq 0 \text{ whenever } \max_{y \in \mathcal{Y}} p(x, y) \geq (1 - c)$$

$$\alpha \inf_{h \in \mathcal{H}} \mathbb{E}_{y}[\ell(h, X, y) \mid X = x] - \beta \Psi(c) \geq 0 \text{ whenever } \max_{y \in \mathcal{Y}} p(x, y) \leq (1 - c),$$

which leads to

$$\alpha \inf_{h \in \mathcal{H}} \mathbb{E}_{y}[\ell(h, X, y) \mid X = x] - \beta \Psi(c) = 0 \text{ whenever } \max_{y \in \mathcal{Y}} p(x, y) = (1 - c).$$
(9)

It is clear that (9) contradicts the assumption that $\exists x \in \mathfrak{X}$ such that $\inf_{h \in \mathcal{H}} \mathbb{E}_y[\ell(h, X, y) \mid X = x] \neq \frac{\beta \Psi(1 - \max_{y \in \mathcal{Y}} p(x, y))}{2}$.

G.2. Proof of $(\mathcal{H}, \mathcal{R})$ -consistency bounds for single-stage surrogates (Theorem 2)

Theorem 2 ($(\mathcal{H}, \mathcal{R})$ -consistency bounds for single-stage surrogates) Assume that \mathcal{H} is symmetric and complete and \mathcal{R} is complete. Then, for $\alpha = \beta$, and $\ell = \ell_{\text{mae}}$, or $\ell = \ell_{\rho}$ with $\Psi(t) = t$, or $\ell = \ell_{\rho-\text{hinge}}$ with $\Psi(t) = nt$, the following $(\mathcal{H}, \mathcal{R})$ -consistency bound holds for all $h \in \mathcal{H}$, $r \in \mathcal{R}$ and any distribution:

$$\mathcal{E}_{\mathsf{L}_{\mathrm{abs}}}(h,r) - \mathcal{E}^*_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) \leq \Gamma(\mathcal{E}_{\mathsf{L}}(h,r) - \mathcal{E}^*_{\mathsf{L}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}}(\mathcal{H},\mathcal{R})),$$

where $\Gamma(t) = \max\{2n\sqrt{t}, nt\}$ for $\ell = \ell_{\text{mae}}$; $\Gamma(t) = \max\{2\sqrt{t}, t\}$ for $\ell = \ell_{\rho}$; and $\Gamma(t) = \max\{2\sqrt{nt}, t\}$ for $\ell = \ell_{\rho-\text{hinge}}$.

Proof When \mathcal{H} is symmetric and complete, $H(x) = \mathcal{Y}$. Since \mathcal{R} is complete, \mathcal{R} is regular for abstention. By Lemma 20,

$$\mathcal{C}^*_{\mathsf{L}_{abs}}(\mathcal{H}, \mathcal{R}, x) = 1 - \max\left\{\max_{y \in \mathcal{Y}} p(x, y), 1 - c\right\},
\Delta \mathcal{C}_{\mathsf{L}_{abs}, \mathcal{H}, \mathcal{R}}(h, r, x) = \left\{\max\{\max_{y \in \mathcal{Y}} p(x, y), 1 - c\} - p(x, \mathsf{h}(x)) \quad r(x) > 0 \\ \max\{\max_{y \in \mathcal{Y}} p(x, y) - 1 + c, 0\} \quad r(x) \le 0. \end{aligned} \right.$$
(10)

Then, the idea of proof for each loss ℓ is similar, which consists of the analysis in four cases depending on the sign of $\max_{y \in \mathcal{Y}} p(x, y) - (1 - c)$ and the sign of r(x), as shown below.

Mean absolute error loss: $\ell = \ell_{\text{mae}}$. When $\ell = \ell_{\text{mae}}$ with $\Psi(t) = t$, let $s_h(x, y) = \frac{e^{h(x,y)}}{\sum_{y' \in \mathcal{Y}} e^{h(x,y')}}$, by the assumption that \mathcal{H} is symmetric and complete and \mathcal{R} is complete, we obtain $\inf_{h \in \mathcal{H}} \mathbb{E}_y[\ell_{\text{mae}}(h, X, y) | X = x] = 1 - \max_{y \in \mathcal{Y}} p(x, y)$ and

$$\begin{aligned} & \mathcal{C}_{\mathsf{L}}(h,r,x) = \sum_{y \in \mathcal{Y}} p(x,y)(1-s_h(x,y))e^{\alpha r(x)} + ce^{-\alpha r(x)} \\ & \mathcal{C}_{\mathsf{L}}^*(\mathcal{H},\mathcal{R},x) = 2\sqrt{c\left(1-\max_{y \in \mathcal{Y}} p(x,y)\right)}. \end{aligned}$$

Note that for any $h \in \mathcal{H}$,

$$\begin{split} \sum_{y \in \mathcal{Y}} p(x,y)(1-s_h(x,y)) &- \left(1 - \max_{y \in \mathcal{Y}} p(x,y)\right) \\ &= \max_{y \in \mathcal{Y}} p(x,y) - \sum_{y \in \mathcal{Y}} p(x,y) s_h(x,y) \\ &\geq \max_{y \in \mathcal{Y}} p(x,y) - p(x,\mathsf{h}(x)s_h(x,\mathsf{h}(x)) - \max_{y \in \mathcal{Y}} p(x,y)(1-s_h(x,\mathsf{h}(x))) \\ &\qquad (p(x,y) \leq \max_{y \in \mathcal{Y}} p(x,y), \,\forall y \neq \mathsf{h}(x)) \\ &= s_h(x,\mathsf{h}(x)) \left(\max_{y \in \mathcal{Y}} p(x,y) - p(x,\mathsf{h}(x)\right) \right) \\ &\geq \frac{1}{n} \left(\max_{y \in \mathcal{Y}} p(x,y) - p(x,\mathsf{h}(x)\right). \qquad (s_h(x,\mathsf{h}(x)) = \max_{y \in \mathcal{Y}} s_h(x,y) \geq \frac{1}{n}) \end{split}$$

We will then analyze the following four cases.

(i) $\max_{y \in \mathcal{Y}} p(x, y) > (1 - c)$ and r(x) > 0. In this case, by (10), we have $\mathcal{C}^*_{\mathsf{L}_{abs}}(\mathcal{H}, \mathcal{R}, x) = 1 - \max_{y \in \mathcal{Y}} p(x, y)$ and $\Delta \mathcal{C}_{\mathsf{L}_{abs}}(\mathcal{H}, r, x) = \max_{y \in \mathcal{Y}} p(x, y) - p(x, \mathsf{h}(x))$. For the surrogate loss, we have

Therefore,

$$\begin{aligned} \mathcal{E}_{\mathsf{L}_{\mathrm{abs}}}(h,r) - \mathcal{E}_{\mathsf{L}_{\mathrm{abs}}}^{*}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) &= \mathbb{E}_{X} \Big[\Delta \mathcal{C}_{\mathsf{L}_{\mathrm{abs}},\mathcal{H},\mathcal{R}}(h,r,x) \Big] \\ &\leq \mathbb{E}_{X} \Big[\Gamma_{1} \Big(\Delta \mathcal{C}_{\mathsf{L},\mathcal{H},\mathcal{R}}(h,r,x) \Big) \Big] \\ &\leq \Gamma_{1} \Big(\mathbb{E}_{X} \Big[\Delta \mathcal{C}_{\mathsf{L},\mathcal{H},\mathcal{R}}(h,r,x) \Big] \Big) \qquad (\Gamma_{1} \text{ is concave}) \\ &= \Gamma_{1} \big(\mathcal{E}_{\mathsf{L}}(h,r) - \mathcal{E}_{\mathsf{L}}^{*}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}}(\mathcal{H},\mathcal{R})) \Big) \end{aligned}$$

where $\Gamma_1(t) = n t$.

(ii) $\max_{y \in \mathcal{Y}} p(x, y) \leq (1 - c)$ and r(x) > 0. In this case, by (10), we have $\mathcal{C}^*_{\mathsf{L}_{abs}}(\mathcal{H}, \mathcal{R}, x) = c$ and $\Delta \mathcal{C}_{\mathsf{L}_{abs}}(\mathcal{H}, r, x) = 1 - c - p(x, \mathsf{h}(x))$. For the surrogate loss, we have

Therefore,

$$\begin{aligned} \mathcal{E}_{\mathsf{L}_{\mathrm{abs}}}(h,r) - \mathcal{E}^{*}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) &= \mathbb{E}_{X} \Big[\Delta \mathcal{C}_{\mathsf{L}_{\mathrm{abs}},\mathcal{H},\mathcal{R}}(h,r,x) \Big] \\ &\leq \mathbb{E}_{X} \Big[\Gamma_{2} \Big(\Delta \mathcal{C}_{\mathsf{L},\mathcal{H},\mathcal{R}}(h,r,x) \Big) \Big] \\ &\leq \Gamma_{2} \Big(\mathbb{E}_{X} \Big[\Delta \mathcal{C}_{\mathsf{L},\mathcal{H},\mathcal{R}}(h,r,x) \Big] \Big) \qquad (\Gamma_{2} \text{ is concave}) \\ &= \Gamma_{2} \big(\mathcal{E}_{\mathsf{L}}(h,r) - \mathcal{E}^{*}_{\mathsf{L}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}}(\mathcal{H},\mathcal{R})) \end{aligned}$$

where $\Gamma_2(t) = 2n\sqrt{t}$.

(iii) $\max_{y \in \mathcal{Y}} p(x, y) \leq (1 - c) \text{ and } r(x) \leq 0.$ In this case, by (10), we have $\mathcal{C}^*_{\mathsf{L}_{abs}}(\mathcal{H}, \mathcal{R}, x) = c$ and $\Delta \mathcal{C}_{\mathsf{L}_{abs}}(\mathcal{H}, r, x) = 0$, which implies that $\mathcal{E}_{\mathsf{L}_{abs}}(h, r) - \mathcal{E}^*_{\mathsf{L}_{abs}}(\mathcal{H}, \mathcal{R}) + \mathcal{M}_{\mathsf{L}_{abs}}(\mathcal{H}, \mathcal{R}) = \mathbb{E}_X \Big[\Delta \mathcal{C}_{\mathsf{L}_{abs}}(\mathcal{H}, r, x) \Big] = 0 \leq \Gamma \Big(\mathcal{E}_{\mathsf{L}}(h, r) - \mathcal{E}^*_{\mathsf{L}}(\mathcal{H}, \mathcal{R}) + \mathcal{M}_{\mathsf{L}}(\mathcal{H}, \mathcal{R}) \Big) \text{ for any } \Gamma \geq 0.$

(iv) $\max_{y \in \mathcal{Y}} p(x, y) > (1 - c)$ and $r(x) \leq 0$. In this case, by (10), we have $\mathcal{C}^*_{\mathsf{L}_{abs}}(\mathcal{H}, \mathcal{R}, x) = 1 - \max_{y \in \mathcal{Y}} p(x, y)$ and $\Delta \mathcal{C}_{\mathsf{L}_{abs}, \mathcal{H}, \mathcal{R}}(h, r, x) = \max_{y \in \mathcal{Y}} p(x, y) - 1 + c$. For the surrogate loss, we have $\Delta \mathcal{C}_{\mathsf{L}_{abs}, \mathcal{H}, \mathcal{R}}(h, r, x)$

$$= \sum_{y \in \mathcal{Y}} p(x,y)(1 - s_h(x,y))e^{\alpha r(x)} + ce^{-\alpha r(x)} - 2\sqrt{c\left(1 - \max_{y \in \mathcal{Y}} p(x,y)\right)}$$

$$\geq \left(1 - \max_{y \in \mathcal{Y}} p(x,y)\right)e^{\alpha r(x)} + ce^{-\alpha r(x)} - 2\sqrt{c\left(1 - \max_{y \in \mathcal{Y}} p(x,y)\right)}$$

$$(\sum_{y \in \mathcal{Y}} p(x,y)(1 - s_h(x,y)) \geq 1 - \max_{y \in \mathcal{Y}} p(x,y))$$

$$\geq 1 - \max_{y \in \mathcal{Y}} p(x,y) + c - 2\sqrt{c\left(1 - \max_{y \in \mathcal{Y}} p(x,y)\right)} \quad (\text{decreasing for } r(x) \leq 0)$$

$$= \left(\sqrt{1 - \max_{y \in \mathcal{Y}} p(x,y)} - \sqrt{c}\right)^2$$

$$= \left(\frac{1 - \max_{y \in \mathcal{Y}} p(x,y) - c}{\sqrt{1 - \max_{y \in \mathcal{Y}} p(x,y)} + \sqrt{c}}\right)^2$$

$$\geq \left(\frac{\max_{y \in \mathcal{Y}} p(x,y) - 1 + c}{2}\right)^2 \quad (\sqrt{1 - \max_{y \in \mathcal{Y}} p(x,y)} + \sqrt{c} \leq 2)$$

$$= \frac{\Delta \mathfrak{C}_{\mathsf{L}_{abs}, \mathcal{H}, \mathcal{H}}(h, r, x)^2}{4}.$$

Therefore,

$$\begin{aligned} \mathcal{E}_{\mathsf{L}_{\mathrm{abs}}}(h,r) - \mathcal{E}^{*}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) &= \mathbb{E}_{X} \Big[\Delta \mathcal{C}_{\mathsf{L}_{\mathrm{abs}},\mathcal{H},\mathcal{R}}(h,r,x) \Big] \\ &\leq \mathbb{E}_{X} \Big[\Gamma_{3} \big(\Delta \mathcal{C}_{\mathsf{L},\mathcal{H},\mathcal{R}}(h,r,x) \big) \Big] \\ &\leq \Gamma_{3} \big(\mathbb{E}_{X} \Big[\Delta \mathcal{C}_{\mathsf{L},\mathcal{H},\mathcal{R}}(h,r,x) \Big] \big) \qquad (\Gamma_{3} \text{ is concave}) \\ &= \Gamma_{3} \big(\mathcal{E}_{\mathsf{L}}(h,r) - \mathcal{E}^{*}_{\mathsf{L}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}}(\mathcal{H},\mathcal{R})) \end{aligned}$$

where $\Gamma_3(t) = 2\sqrt{t}$.

Overall, we obtain

$$\mathcal{E}_{\mathsf{L}_{\mathrm{abs}}}(h,r) - \mathcal{E}^*_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) \leq \Gamma(\mathcal{E}_{\mathsf{L}}(h,r) - \mathcal{E}^*_{\mathsf{L}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}}(\mathcal{H},\mathcal{R}))$$

where $\Gamma(t) = \max{\{\Gamma_1(t), \Gamma_2(t), \Gamma_3(t)\}} = \max{\{2n\sqrt{t}, nt\}}$, which completes the proof.

 ρ -Margin loss: $\ell = \ell_{\rho}$. When $\ell = \ell_{\rho}$ with $\Psi(t) = t$, by the assumption that \mathcal{H} is symmetric and complete and \mathcal{R} is complete, we obtain $\inf_{h \in \mathcal{H}} \mathbb{E}_{y}[\ell_{\rho}(h, X, y) | X = x] = 1 - \max_{y \in \mathcal{Y}} p(x, y)$ and

$$\begin{split} \mathcal{C}_{\mathsf{L}}(h,r,x) &= \sum_{y \in \mathfrak{Y}} p(x,y) \min\left\{ \max\left\{0, 1 - \frac{\rho_h(x,y)}{\rho}\right\}, 1 \right\} e^{\alpha r(x)} + c e^{-\alpha r(x)} \\ &= \left(1 - \min\left\{1, \frac{\rho_h(x,\mathsf{h}(x))}{\rho}\right\} p(x,\mathsf{h}(x))\right) e^{\alpha r(x)} + c e^{-\alpha r(x)} \\ \mathcal{C}_{\mathsf{L}}^*(\mathcal{H},\mathcal{R},x) &= 2\sqrt{c \left(1 - \max_{y \in \mathfrak{Y}} p(x,y)\right)}. \end{split}$$

where $\rho_h(x, y) = h(x, y) - \max_{y' \neq y} h(x, y')$ is the margin. Note that for any $h \in \mathcal{H}$,

$$1 - \min\left\{1, \frac{\rho_h(x, \mathsf{h}(x))}{\rho}\right\} p(x, \mathsf{h}(x)) - \left(1 - \max_{y \in \mathcal{Y}} p(x, y)\right)$$
$$= \max_{y \in \mathcal{Y}} p(x, y) - \min\left\{1, \frac{\rho_h(x, \mathsf{h}(x))}{\rho}\right\} p(x, \mathsf{h}(x))$$
$$\geq \max_{y \in \mathcal{Y}} p(x, y) - p(x, \mathsf{h}(x). \qquad (\min\left\{1, \frac{\rho_h(x, \mathsf{h}(x))}{\rho}\right\} \le 1)$$

We will then analyze the following four cases.

(i) $\max_{y \in \mathcal{Y}} p(x, y) \leq (1 - c)$ and r(x) > 0. In this case, by (10), we have $\mathcal{C}^*_{\mathsf{L}_{abs}}(\mathcal{H}, \mathcal{R}, x) = c$ and $\Delta \mathcal{C}_{\mathsf{L}_{abs}, \mathcal{H}, \mathcal{R}}(h, r, x) = 1 - c - p(x, \mathsf{h}(x))$. For the surrogate loss, we have

$$\begin{split} \Delta \mathcal{C}_{\mathsf{L},\mathfrak{H},\mathfrak{K},\mathfrak{K}}(h,r,x) &= \left(1 - \min\left\{1,\frac{\rho_h(x,\mathsf{h}(x))}{\rho}\right\} p(x,\mathsf{h}(x))\right) e^{\alpha r(x)} + c e^{-\alpha r(x)} - 2\sqrt{c\left(1 - \max\left\{1,\frac{\rho_h(x,\mathsf{h}(x))}{\rho}\right\} p(x,\mathsf{h}(x))\right)} \\ &\geq \left(1 - \min\left\{1,\frac{\rho_h(x,\mathsf{h}(x))}{\rho}\right\} p(x,\mathsf{h}(x))\right) e^{\alpha r(x)} + c e^{-\alpha r(x)} - 2\sqrt{c\left(1 - \min\left\{1,\frac{\rho_h(x,\mathsf{h}(x))}{\rho}\right\} p(x,\mathsf{h}(x))\right)} \\ &\quad (1 - \min\left\{1,\frac{\rho_h(x,\mathsf{h}(x))}{\rho}\right\} p(x,\mathsf{h}(x)) + c - 2\sqrt{c\left(1 - \min\left\{1,\frac{\rho_h(x,\mathsf{h}(x))}{\rho}\right\} p(x,\mathsf{h}(x))\right)} \right) \\ &\geq 1 - \min\left\{1,\frac{\rho_h(x,\mathsf{h}(x))}{\rho}\right\} p(x,\mathsf{h}(x)) + c - 2\sqrt{c\left(1 - \min\left\{1,\frac{\rho_h(x,\mathsf{h}(x))}{\rho}\right\} p(x,\mathsf{h}(x))\right)} \\ &\quad (\text{increasing for } r(x) \ge 0) \\ &= \left(\sqrt{1 - \min\left\{1,\frac{\rho_h(x,\mathsf{h}(x))}{\rho}\right\} p(x,\mathsf{h}(x)) - \sqrt{c}}\right)^2 \\ &= \left(\frac{1 - \min\left\{1,\frac{\rho_h(x,\mathsf{h}(x))}{\rho}\right\} p(x,\mathsf{h}(x)) - (1 - \max_{y \in \mathcal{Y}} p(x,y)) + (1 - \max_{y \in \mathcal{Y}} p(x,y) - c)}{\sqrt{1 - \min\left\{1,\frac{\rho_h(x,\mathsf{h}(x))}{\rho}\right\}} p(x,\mathsf{h}(x)) - (1 - \max_{y \in \mathcal{Y}} p(x,y))\right)} \right)^2 \\ &= \left(\frac{\max_{y \in \mathcal{Y}} p(x,y) - p(x,\mathsf{h}(x) + (1 - \max_{y \in \mathcal{Y}} p(x,y) - c))}{2}\right)^2 \\ &= \left(\frac{\max_{y \in \mathcal{Y}} p(x,y) - p(x,\mathsf{h}(x) + (1 - \max_{y \in \mathcal{Y}} p(x,y) - c))}{2}\right)^2 \\ &= \left(1 - \min\left\{1,\frac{\rho_h(x,\mathsf{h}(x))}{\rho}\right\} p(x,\mathsf{h}(x)) - (1 - \max_{y \in \mathcal{Y}} p(x,y))\right) \ge \max_{y \in \mathcal{Y}} p(x,y) - p(x,\mathsf{h}(x)) \\ &= \frac{1}{4}(1 - c - p(x,\mathsf{h}(x)))^2 \\ &= \frac{\Delta \mathcal{C}_{\mathsf{L}_{abov},\mathcal{H},\mathcal{H}(x,\mathsf{h}(x))}{4} \end{aligned}\right)$$

Therefore,

$$\begin{aligned} \mathcal{E}_{\mathsf{L}_{\mathrm{abs}}}(h,r) - \mathcal{E}^{*}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) &= \mathbb{E}_{X} \Big[\Delta \mathcal{C}_{\mathsf{L}_{\mathrm{abs}},\mathcal{H},\mathcal{R}}(h,r,x) \Big] \\ &\leq \mathbb{E}_{X} \Big[\Gamma_{2} \Big(\Delta \mathcal{C}_{\mathsf{L},\mathcal{H},\mathcal{R}}(h,r,x) \Big) \Big] \\ &\leq \Gamma_{2} \Big(\mathbb{E}_{X} \Big[\Delta \mathcal{C}_{\mathsf{L},\mathcal{H},\mathcal{R}}(h,r,x) \Big] \Big) \qquad (\Gamma_{2} \text{ is concave}) \\ &= \Gamma_{2} \big(\mathcal{E}_{\mathsf{L}}(h,r) - \mathcal{E}^{*}_{\mathsf{L}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}}(\mathcal{H},\mathcal{R})) \end{aligned}$$

where $\Gamma_2(t) = 2\sqrt{t}$.

(ii) $\max_{y \in \mathcal{Y}} p(x, y) > (1 - c)$ and r(x) > 0. In this case, by (10), we have $\mathcal{C}^*_{\mathsf{L}_{abs}}(\mathcal{H}, \mathcal{R}, x) = 1 - \max_{y \in \mathcal{Y}} p(x, y)$ and $\Delta \mathcal{C}_{\mathsf{L}_{abs}}, \mathcal{H}, \mathcal{R}(h, r, x) = \max_{y \in \mathcal{Y}} p(x, y) - p(x, \mathsf{h}(x))$. For the surrogate loss, we have

Therefore, $\mathcal{E}_{\mathsf{L}_{abs}}(h, r) - \mathcal{E}^{*}_{\mathsf{L}_{abs}}(\mathcal{H}, \mathcal{R}) + \mathcal{M}_{\mathsf{L}_{abs}}(\mathcal{H}, \mathcal{R}) = \mathbb{E}_{X} \Big[\Delta \mathcal{C}_{\mathsf{L}_{abs}, \mathcal{H}, \mathcal{R}}(h, r, x) \Big] \leq \mathbb{E}_{X} \Big[\Gamma_{1} \Big(\Delta \mathcal{C}_{\mathsf{L}, \mathcal{H}, \mathcal{R}}(h, r, x) \Big) \Big] \leq \Gamma_{1} \Big(\mathbb{E}_{X} \Big[\Delta \mathcal{C}_{\mathsf{L}, \mathcal{H}, \mathcal{R}}(h, r, x) \Big] \Big) = \Gamma_{1} \big(\mathcal{E}_{\mathsf{L}}(h, r) - \mathcal{E}^{*}_{\mathsf{L}}(\mathcal{H}, \mathcal{R}) + \mathcal{M}_{\mathsf{L}}(\mathcal{H}, \mathcal{R})), \text{ where } \Gamma_{1}(t) = t \text{ is concave.} \Big)$

(iii) $\max_{y \in \mathcal{Y}} p(x, y) \leq (1 - c) \text{ and } r(x) \leq 0.$ In this case, by (10), we have $\mathcal{C}^*_{\mathsf{L}_{abs}}(\mathcal{H}, \mathcal{R}, x) = c$ and $\Delta \mathcal{C}_{\mathsf{L}_{abs}}(\mathcal{H}, r, x) = 0$, which implies that $\mathcal{E}_{\mathsf{L}_{abs}}(h, r) - \mathcal{E}^*_{\mathsf{L}_{abs}}(\mathcal{H}, \mathcal{R}) + \mathcal{M}_{\mathsf{L}_{abs}}(\mathcal{H}, \mathcal{R}) = \mathbb{E}_X \left[\Delta \mathcal{C}_{\mathsf{L}_{abs}}, \mathcal{H}, \mathcal{R}(h, r, x) \right] = 0 \leq \Gamma \left(\mathcal{E}_{\mathsf{L}}(h, r) - \mathcal{E}^*_{\mathsf{L}}(\mathcal{H}, \mathcal{R}) + \mathcal{M}_{\mathsf{L}}(\mathcal{H}, \mathcal{R}) \right)$ for any $\Gamma \geq 0$.

(iv) $\max_{y \in \mathcal{Y}} p(x, y) > (1 - c)$ and $r(x) \leq 0$. In this case, by (10), we have $\mathcal{C}^*_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H}, \mathcal{R}, x) = 1 - \max_{y \in \mathcal{Y}} p(x, y)$ and $\Delta \mathcal{C}_{\mathsf{L}_{\mathrm{abs}}, \mathcal{H}, \mathcal{R}}(h, r, x) = \max_{y \in \mathcal{Y}} p(x, y) - 1 + c$. For the surrogate loss, we have

$$\begin{split} \Delta \mathcal{C}_{\mathsf{L},\mathfrak{H},\mathfrak{R}}(h,r,x) &= \left(1 - \min\left\{1,\frac{\rho_{h}(x,\mathsf{h}(x))}{\rho}\right\} p(x,\mathsf{h}(x))\right) e^{\alpha r(x)} + c e^{-\alpha r(x)} - 2\sqrt{c\left(1 - \max_{y \in \mathfrak{Y}} p(x,y)\right)} \\ &\geq \left(1 - \max_{y \in \mathfrak{Y}} p(x,y)\right) e^{\alpha r(x)} + c e^{-\alpha r(x)} - 2\sqrt{c\left(1 - \max_{y \in \mathfrak{Y}} p(x,y)\right)} \\ &\quad (1 - \min\left\{1,\frac{\rho_{h}(x,\mathsf{h}(x))}{\rho}\right\} p(x,\mathsf{h}(x)) \geq 1 - \max_{y \in \mathfrak{Y}} p(x,y)) \\ &\geq 1 - \max_{y \in \mathfrak{Y}} p(x,y) + c - 2\sqrt{c\left(1 - \max_{y \in \mathfrak{Y}} p(x,y)\right)} \qquad (\text{decreasing for } r(x) \leq 0) \\ &= \left(\sqrt{1 - \max_{y \in \mathfrak{Y}} p(x,y)} - \sqrt{c}\right)^{2} \\ &= \left(\frac{1 - \max_{y \in \mathfrak{Y}} p(x,y) - 1 + c}{\sqrt{1 - \max_{y \in \mathfrak{Y}} p(x,y)} + \sqrt{c}}\right)^{2} \qquad (\sqrt{1 - \max_{y \in \mathfrak{Y}} p(x,y)} + \sqrt{c} \leq 2) \\ &= \frac{\Delta \mathcal{C}_{\mathsf{L}_{abs}}, \mathfrak{H}, \mathfrak{R}(h, r, x)^{2}}{4}. \end{split}$$

Therefore,

$$\begin{aligned} \mathcal{E}_{\mathsf{L}_{\mathrm{abs}}}(h,r) - \mathcal{E}^{*}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) &= \mathbb{E}_{X} \Big[\Delta \mathcal{C}_{\mathsf{L}_{\mathrm{abs}},\mathcal{H},\mathcal{R}}(h,r,x) \Big] \\ &\leq \mathbb{E}_{X} \Big[\Gamma_{3} \big(\Delta \mathcal{C}_{\mathsf{L},\mathcal{H},\mathcal{R}}(h,r,x) \big) \Big] \\ &\leq \Gamma_{3} \big(\mathbb{E}_{X} \Big[\Delta \mathcal{C}_{\mathsf{L},\mathcal{H},\mathcal{R}}(h,r,x) \Big] \big) \qquad (\Gamma_{3} \text{ is concave}) \\ &= \Gamma_{3} \big(\mathcal{E}_{\mathsf{L}}(h,r) - \mathcal{E}^{*}_{\mathsf{L}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}}(\mathcal{H},\mathcal{R})) \end{aligned}$$

where $\Gamma_3(t) = 2\sqrt{t}$.

Overall, we obtain

$$\mathcal{E}_{\mathsf{L}_{\mathrm{abs}}}(h,r) - \mathcal{E}^*_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) \leq \Gamma(\mathcal{E}_{\mathsf{L}}(h,r) - \mathcal{E}^*_{\mathsf{L}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}}(\mathcal{H},\mathcal{R}))$$

where $\Gamma(t) = \max{\{\Gamma_1(t), \Gamma_2(t), \Gamma_3(t)\}} = \max{\{2\sqrt{t}, t\}}$, which completes the proof.

Constrained ρ **-hinge loss:** $\ell = \ell_{\rho-\text{hinge}}$. When $\ell = \ell_{\rho-\text{hinge}}$ with $\Psi(t) = nt$, by the assumption that \mathcal{H} is symmetric and complete and \mathcal{R} is complete, we obtain $\inf_{h \in \mathcal{H}} \mathbb{E}_y[\ell_{\rho-\text{hinge}}(h, X, y) \mid X = x] =$

 $n(1 - \max_{y \in \mathcal{Y}} p(x, y))$ and with the constraint $\sum_{y \in \mathcal{Y}} h(x, y) = 0$,

$$\begin{split} \mathcal{C}_{\mathsf{L}}(h,r,x) &= \sum_{y \in \mathcal{Y}} p(x,y) \sum_{y' \neq y} \max\{0, 1 + \frac{h(x,y')}{\rho}\} e^{\alpha r(x)} + nce^{-\alpha r(x)} \\ &= \sum_{y \in \mathcal{Y}} (1 - p(x,y)) \max\left\{0, 1 + \frac{h(x,y)}{\rho}\right\} e^{\alpha r(x)} + nce^{-\alpha r(x)} \\ \mathcal{C}_{\mathsf{L}}^{*}(\mathcal{H},\mathcal{R},x) &= 2\sqrt{n^{2}c \left(1 - \max_{y \in \mathcal{Y}} p(x,y)\right)}. \end{split}$$

Take
$$h_{\rho} \in \mathcal{H}$$
 such that $h_{\rho}(x, y) = \begin{cases} h(x, y) & \text{if } y \notin \{y_{\max}, h(x)\} \\ -\rho & \text{if } y = h(x) \\ h(x, y_{\max}) + h(x, h(x)) + \rho & \text{if } y = y_{\max}. \end{cases}$
constraint $\sum_{y \in \mathcal{Y}} h_{\rho}(x, y) = 0$, where $y_{\max} = \operatorname{argmax}_{y \in \mathcal{Y}} p(x, y)$. Note that for any $h \in \mathcal{H}$,

$$\begin{split} &\sum_{y \in \mathcal{Y}} (1 - p(x, y)) \max\left\{0, 1 + \frac{h(x, y)}{\rho}\right\} - n\left(1 - \max_{y \in \mathcal{Y}} p(x, y)\right) \\ &\geq \sum_{y \in \mathcal{Y}} (1 - p(x, y)) \min\left\{n, \max\left\{0, 1 + \frac{h(x, y)}{\rho}\right\}\right\} - n\left(1 - \max_{y \in \mathcal{Y}} p(x, y)\right) \\ &\geq \sum_{y \in \mathcal{Y}} (1 - p(x, y)) \min\left\{n, \max\left\{0, 1 + \frac{h(x, y)}{\rho}\right\}\right\} - \sum_{y \in \mathcal{Y}} (1 - p(x, y)) \min\left\{n, \max\left\{0, 1 + \frac{h_{\rho}(x, y)}{\rho}\right\}\right\} \\ &\geq \min\left\{n, 1 + \frac{h(x, h(x))}{\rho}\right\} \left(\max_{y \in \mathcal{Y}} p(x, y) - p(x, h(x))\right) \qquad (\text{plug in } h_{\rho}(x, y)) \\ &\geq \max_{y \in \mathcal{Y}} p(x, y) - p(x, h(x). \qquad (h(x, h(x)) \ge 0) \end{split}$$

We will then analyze the following four cases.

(i) $\max_{y \in \mathcal{Y}} p(x, y) > (1 - c)$ and r(x) > 0. In this case, by (10), we have $\mathcal{C}^*_{\mathsf{L}_{abs}}(\mathcal{H}, \mathcal{R}, x) = 1 - \max_{y \in \mathcal{Y}} p(x, y)$ and $\Delta \mathcal{C}_{\mathsf{L}_{abs}}(\mathcal{H}, r, x) = \max_{y \in \mathcal{Y}} p(x, y) - p(x, \mathsf{h}(x))$. For the surrogate loss,

Therefore,

$$\begin{aligned} \mathcal{E}_{\mathsf{L}_{\mathrm{abs}}}(h,r) - \mathcal{E}^{*}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) &= \mathbb{E}_{X} \Big[\Delta \mathcal{C}_{\mathsf{L}_{\mathrm{abs}},\mathcal{H},\mathcal{R}}(h,r,x) \Big] \\ &\leq \mathbb{E}_{X} \Big[\Gamma_{1} \Big(\Delta \mathcal{C}_{\mathsf{L},\mathcal{H},\mathcal{R}}(h,r,x) \Big) \Big] \\ &\leq \Gamma_{1} \Big(\mathbb{E}_{X} \Big[\Delta \mathcal{C}_{\mathsf{L},\mathcal{H},\mathcal{R}}(h,r,x) \Big] \Big) \qquad (\Gamma_{1} \text{ is concave}) \\ &= \Gamma_{1} \big(\mathcal{E}_{\mathsf{L}}(h,r) - \mathcal{E}^{*}_{\mathsf{L}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}}(\mathcal{H},\mathcal{R}) \big) \end{aligned}$$

where $\Gamma_1(t) = t$.

$$\begin{aligned} \mathcal{E}_{\mathsf{L}_{\mathrm{abs}}}(h,r) - \mathcal{E}^{*}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) &= \mathbb{E}_{X} \Big[\Delta \mathcal{C}_{\mathsf{L}_{\mathrm{abs}},\mathcal{H},\mathcal{R}}(h,r,x) \Big] \\ &\leq \mathbb{E}_{X} \Big[\Gamma_{2} \Big(\Delta \mathcal{C}_{\mathsf{L},\mathcal{H},\mathcal{R}}(h,r,x) \Big) \Big] \\ &\leq \Gamma_{2} \Big(\mathbb{E}_{X} \Big[\Delta \mathcal{C}_{\mathsf{L},\mathcal{H},\mathcal{R}}(h,r,x) \Big] \Big) \qquad (\Gamma_{2} \text{ is concave}) \\ &= \Gamma_{2} \big(\mathcal{E}_{\mathsf{L}}(h,r) - \mathcal{E}^{*}_{\mathsf{L}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}}(\mathcal{H},\mathcal{R})) \end{aligned}$$

where $\Gamma_2(t) = 2\sqrt{nt}$.

(iii) $\max_{y \in \mathcal{Y}} p(x, y) \leq (1 - c) \text{ and } r(x) \leq 0.$ In this case, by (10), we have $\mathcal{C}^*_{\mathsf{L}_{abs}}(\mathcal{H}, \mathcal{R}, x) = c$ and $\Delta \mathcal{C}_{\mathsf{L}_{abs},\mathcal{H},\mathcal{R}}(h, r, x) = 0$, which implies that $\mathcal{E}_{\mathsf{L}_{abs}}(h, r) - \mathcal{E}^*_{\mathsf{L}_{abs}}(\mathcal{H}, \mathcal{R}) + \mathcal{M}_{\mathsf{L}_{abs}}(\mathcal{H}, \mathcal{R}) = \mathbb{E}_X \left[\Delta \mathcal{C}_{\mathsf{L}_{abs},\mathcal{H},\mathcal{R}}(h, r, x) \right] = 0 \leq \Gamma \left(\mathcal{E}_{\mathsf{L}}(h, r) - \mathcal{E}^*_{\mathsf{L}}(\mathcal{H}, \mathcal{R}) + \mathcal{M}_{\mathsf{L}}(\mathcal{H}, \mathcal{R}) \right)$ for any $\Gamma \geq 0$.

(iv) $\max_{y \in \mathcal{Y}} p(x, y) > (1 - c)$ and $r(x) \leq 0$. In this case, by (10), we have $\mathcal{C}^*_{\mathsf{L}_{abs}}(\mathcal{H}, \mathcal{R}, x) = 1 - \max_{y \in \mathcal{Y}} p(x, y)$ and $\Delta \mathcal{C}_{\mathsf{L}_{abs}}(\mathcal{H}, r, x) = \max_{y \in \mathcal{Y}} p(x, y) - 1 + c$. For the surrogate loss, we have

Therefore,

$$\begin{aligned} \mathcal{E}_{\mathsf{L}_{\mathrm{abs}}}(h,r) - \mathcal{E}^{*}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) &= \mathbb{E}_{X} \Big[\Delta \mathcal{C}_{\mathsf{L}_{\mathrm{abs}},\mathcal{H},\mathcal{R}}(h,r,x) \Big] \\ &\leq \mathbb{E}_{X} \Big[\Gamma_{3} \big(\Delta \mathcal{C}_{\mathsf{L},\mathcal{H},\mathcal{R}}(h,r,x) \big) \Big] \\ &\leq \Gamma_{3} \big(\mathbb{E}_{X} \big[\Delta \mathcal{C}_{\mathsf{L},\mathcal{H},\mathcal{R}}(h,r,x) \big] \big) \qquad (\Gamma_{3} \text{ is concave}) \\ &= \Gamma_{3} \big(\mathcal{E}_{\mathsf{L}}(h,r) - \mathcal{E}^{*}_{\mathsf{L}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}}(\mathcal{H},\mathcal{R})) \end{aligned}$$

where $\Gamma_3(t) = 2\sqrt{t/n}$.

Overall, we obtain

$$\mathcal{E}_{\mathsf{L}_{\mathrm{abs}}}(h,r) - \mathcal{E}^*_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) \leq \Gamma(\mathcal{E}_{\mathsf{L}}(h,r) - \mathcal{E}^*_{\mathsf{L}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}}(\mathcal{H},\mathcal{R}))$$

where $\Gamma(t) = \max{\{\Gamma_1(t), \Gamma_2(t), \Gamma_3(t)\}} = \max{\{2\sqrt{nt}, t\}}$, which completes the proof.

G.3. Proof of \mathcal{R} -consistency bounds for second-stage surrogates (Theorem 4)

Theorem 4 (\Re -consistency bounds for second-stage surrogates) Fix a predictor h. Assume that Φ admits an \Re -consistency bound with respect to $\ell_{0-1}^{\text{binary}}$. Thus, there exists a non-decreasing concave function Γ such that, for all $r \in \Re$,

$$\mathcal{E}_{\ell_{0-1}^{\text{binary}}}(r) - \mathcal{E}_{\ell_{0-1}^{\text{binary}}}^{*}(\mathcal{R}) + \mathcal{M}_{\ell_{0-1}^{\text{binary}}}(\mathcal{R}) \leq \Gamma(\mathcal{E}_{\Phi}(r) - \mathcal{E}_{\Phi}^{*}(\mathcal{R}) + \mathcal{M}_{\Phi}(\mathcal{R})).$$

Then, the following \mathbb{R} *-consistency bound holds for all* $r \in \mathbb{R}$ *and any distribution:*

$$\mathcal{E}_{\ell_{\mathrm{abs},h}}(r) - \mathcal{E}^*_{\ell_{\mathrm{abs},h}}(\mathcal{R}) + \mathcal{M}_{\ell_{\mathrm{abs},h}}(\mathcal{R}) \leq \Gamma\Big(\Big(\mathcal{E}_{\ell_{\Phi,h}}(r) - \mathcal{E}^*_{\ell_{\Phi,h}}(\mathcal{R}) + \mathcal{M}_{\ell_{\Phi,h}}(\mathcal{R})\Big)/c\Big).$$

Proof Given any fixed predictor h. For any $r \in \mathbb{R}$, $x \in \mathcal{X}$ and $y \in \mathcal{Y}$, the conditional risk of $\ell_{abs,h}$ and $\ell_{\Phi,h}$ can be written as

$$C_{\ell_{abs,h}}(r,x) = \sum_{y \in \mathcal{Y}} p(x,y) \mathbb{1}_{h(x) \neq y} \mathbb{1}_{r(x) > 0} + c \mathbb{1}_{r(x) \le 0}$$

$$C_{\ell_{\Phi,h}}(r,x) = \sum_{y \in \mathcal{Y}} p(x,y) \mathbb{1}_{h(x) \neq y} \Phi(-r(x)) + c \Phi(r(x)).$$
(11)

Thus, the best-in class conditional risk of $\ell_{abs,h}$ and $\ell_{\Phi,h}$ can be expressed as

$$\mathcal{C}^{*}_{\ell_{\mathrm{abs},h}}(\mathcal{R},x) = \inf_{r \in \mathcal{R}} \left(\sum_{y \in \mathcal{Y}} p(x,y) \mathbb{1}_{\mathsf{h}(x) \neq y} \mathbb{1}_{r(x) > 0} + c \mathbb{1}_{r(x) \leq 0} \right) \\
\mathcal{C}^{*}_{\ell_{\Phi,h}}(\mathcal{R},x) = \inf_{r \in \mathcal{R}} \left(\sum_{y \in \mathcal{Y}} p(x,y) \mathbb{1}_{\mathsf{h}(x) \neq y} \Phi(-r(x)) + c \Phi(r(x)) \right).$$
(12)

Let $p_1 = \frac{\sum_{y \in \mathcal{Y}} p(x,y) \mathbb{1}_{h(x)\neq y}}{\sum_{y \in \mathcal{Y}} p(x,y) \mathbb{1}_{h(x)\neq y}+c}$ and $p_2 = \frac{c}{\sum_{y \in \mathcal{Y}} p(x,y) \mathbb{1}_{h(x)\neq y}+c}$. Then, the calibration gap of $\ell_{abs,h}$ can be written as

$$\mathcal{C}_{\ell_{\Phi,h}}(r,x) - \mathcal{C}_{\ell_{\Phi,h}}^{*}(\mathcal{R},x) = \left(\sum_{y \in \mathcal{Y}} p(x,y) \mathbb{1}_{\mathsf{h}(x)\neq y} + c\right) \left[p_{1}\Phi(-r(x)) + p_{2}\Phi(r(x)) - \inf_{r \in \mathcal{R}} (p_{1}\Phi(-r(x)) + p_{2}\Phi(r(x))) \right].$$

By Lemma 21, we have

$$\begin{split} & \mathcal{C}_{\ell_{\mathrm{abs},h}}(r,x) - \mathcal{C}^{*}_{\ell_{\mathrm{abs},h}}(\mathcal{R},x) \\ &= p_{1}\mathbf{1}_{r(x)>0} + p_{2}\mathbf{1}_{r(x)\leq 0} - \inf_{r\in\mathcal{R}} \Big(p_{1}\mathbf{1}_{r(x)>0} + p_{2}\mathbf{1}_{r(x)\leq 0} \Big) \\ &\leq \Gamma \Big(p_{1}\Phi(-r(x)) + p_{2}\Phi(r(x)) - \inf_{r\in\mathcal{R}} (p_{1}\Phi(-r(x)) + p_{2}\Phi(r(x))) \Big) \\ &= \Gamma \Big(\frac{\mathcal{C}_{\ell_{\Phi,h}}(r,x) - \inf_{r\in\mathcal{R}} \mathcal{C}_{\ell_{\Phi,h}}(r,x)}{\sum_{y\in\mathcal{Y}} p(x,y)\mathbf{1}_{\mathsf{h}(x)\neq y} + c} \Big) \\ &\leq \Gamma \bigg(\frac{\mathcal{C}_{\ell_{\Phi,h}}(r,x) - \inf_{r\in\mathcal{R}} \mathcal{C}_{\ell_{\Phi,h}}(r,x)}{c} \bigg), \end{split}$$

where we use the fact that Γ is non-decreasing and $\sum_{y \in \mathcal{Y}} p(x, y) \mathbb{1}_{h(x)\neq y} + c \ge c$ in the last inequality. Since Γ is concave, taking the expectation on both sides and using Jensen's inequality, we obtain

$$\mathbb{E}_{X}\left[\mathcal{C}_{\ell_{\mathrm{abs},h}}(r,x) - \mathcal{C}^{*}_{\ell_{\mathrm{abs},h}}(\mathcal{R},x)\right] \leq \Gamma\left(\frac{\mathbb{E}_{X}\left[\mathcal{C}_{\ell_{\Phi,h}}(r,x) - \inf_{r \in \mathcal{R}} \mathcal{C}_{\ell_{\Phi,h}}(r,x)\right]}{c}\right).$$

Since the term $\mathbb{E}_X \Big[\mathbb{C}_{\ell_{\text{abs},h}}(r,x) - \mathbb{C}^*_{\ell_{\text{abs},h}}(\mathcal{R},x) \Big]$ and $\mathbb{E}_X \Big[\mathbb{C}_{\ell_{\Phi,h}}(r,x) - \inf_{r \in \mathcal{R}} \mathbb{C}_{\ell_{\Phi,h}}(r,x) \Big]$ can be expressed as

$$\begin{split} & \mathbb{E}_{X} \Big[\mathbb{C}_{\ell_{\mathrm{abs},h}}(r,x) - \mathbb{C}^{*}_{\ell_{\mathrm{abs},h}}(\mathcal{R},x) \Big] = \mathcal{E}_{\ell_{\mathrm{abs},h}}(r) - \mathcal{E}^{*}_{\ell_{\mathrm{abs},h}}(\mathcal{R}) + \mathcal{M}_{\ell_{\mathrm{abs},h}}(\mathcal{R}) \\ & \mathbb{E}_{X} \Big[\mathbb{C}_{\ell_{\Phi,h}}(r,x) - \inf_{r \in \mathcal{R}} \mathbb{C}_{\ell_{\Phi,h}}(r,x) \Big] = \mathcal{E}_{\ell_{\Phi,h}}(r) - \mathcal{E}^{*}_{\ell_{\Phi,h}}(\mathcal{R}) + \mathcal{M}_{\ell_{\Phi,h}}(\mathcal{R}), \end{split}$$

we have

$$\mathcal{E}_{\ell_{\mathrm{abs},h}}(r) - \mathcal{E}^*_{\ell_{\mathrm{abs},h}}(\mathcal{R}) + \mathcal{M}_{\ell_{\mathrm{abs},h}}(\mathcal{R}) \leq \Gamma\left(\frac{\mathcal{E}_{\ell_{\Phi,h}}(r) - \mathcal{E}^*_{\ell_{\Phi,h}}(\mathcal{R}) + \mathcal{M}_{\ell_{\Phi,h}}(\mathcal{R})}{c}\right),$$

which completes the proof.

G.4. Proof of $(\mathcal{H}, \mathcal{R})$ -consistency bounds for two-stage surrogates (Theorem 6)

Theorem 6 ($(\mathcal{H}, \mathcal{R})$ -consistency bounds for two-stage approach) Suppose that \mathcal{R} is regular. Assume that ℓ admits an \mathcal{H} -consistency bound with respect to ℓ_{0-1} and that Φ admits an \mathcal{R} -consistency bound with respect to $\ell_{0-1}^{\text{binary}}$. Thus, there are non-decreasing concave functions Γ_1 and Γ_2 such that, for all $h \in \mathcal{H}$ and $r \in \mathcal{R}$,

$$\begin{aligned} & \mathcal{E}_{\ell_{0-1}}(h) - \mathcal{E}^*_{\ell_{0-1}}(\mathcal{H}) + \mathcal{M}_{\ell_{0-1}}(\mathcal{H}) \leq \Gamma_1(\mathcal{E}_{\ell}(h) - \mathcal{E}^*_{\ell}(\mathcal{H}) + \mathcal{M}_{\ell}(\mathcal{H})) \\ & \mathcal{E}_{\ell_{0-1}^{\text{binary}}}(r) - \mathcal{E}^*_{\ell_{0-1}^{\text{binary}}}(\mathcal{R}) + \mathcal{M}_{\ell_{0-1}^{\text{binary}}}(\mathcal{R}) \leq \Gamma_2(\mathcal{E}_{\Phi}(r) - \mathcal{E}^*_{\Phi}(\mathcal{R}) + \mathcal{M}_{\Phi}(\mathcal{R})) \end{aligned}$$

Then, the following $(\mathcal{H}, \mathbb{R})$ *-consistency bound holds for all* $h \in \mathcal{H}$ *,* $r \in \mathbb{R}$ *and any distribution:*

$$\begin{aligned} \mathcal{E}_{\mathsf{L}_{\mathrm{abs}}}(h,r) - \mathcal{E}^{*}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) &\leq \Gamma_{1}(\mathcal{E}_{\ell}(h) - \mathcal{E}^{*}_{\ell}(\mathcal{H}) + \mathcal{M}_{\ell}(\mathcal{H})) \\ &+ (1+c)\Gamma_{2}\Big(\Big(\mathcal{E}_{\ell_{\Phi,h}}(r) - \mathcal{E}^{*}_{\ell_{\Phi,h}}(\mathcal{R}) + \mathcal{M}_{\ell_{\Phi,h}}(\mathcal{R})\Big)/c\Big), \end{aligned}$$

where the constant factors (1 + c) and $\frac{1}{c}$ can be removed when Γ_2 is linear.

Proof Since \mathcal{R} is regular, the conditional risk and the best-in-class conditional risk of the abstention loss L_{abs} can be expressed as

$$C_{\mathsf{L}_{abs}}(h, r, x) = \sum_{y \in \mathcal{Y}} p(x, y) \mathbb{1}_{\mathsf{h}(x) \neq y} \mathbb{1}_{r(x) > 0} + c \mathbb{1}_{r(x) \leq 0}$$

$$C^{*}_{\mathsf{L}_{abs}}(\mathcal{H}, \mathcal{R}, x) = \min\left\{ \inf_{h \in \mathcal{H}} \sum_{y \in \mathcal{Y}} p(x, y) \mathbb{1}_{\mathsf{h}(x) \neq y}, c \right\}.$$
(13)

Thus, by introducing the term $\min\{\sum_{y\in\mathcal{Y}} p(x,y)\mathbb{1}_{h(x)\neq y}, c\}$ and subsequently subtracting it after rearranging, the calibration gap of the abstention loss Labs can be written as follows

$$\begin{aligned} &\mathcal{C}_{\mathsf{L}_{abs}}(h,r,x) - \mathcal{C}_{\mathsf{L}_{abs}}^{*}(\mathfrak{H},\mathfrak{R},x) \\ &= \sum_{y \in \mathcal{Y}} p(x,y) \mathbb{1}_{\mathsf{h}(x) \neq y} \mathbb{1}_{r(x) > 0} + c \mathbb{1}_{r(x) \leq 0} - \min\left\{ \inf_{h \in \mathcal{H}} \sum_{y \in \mathcal{Y}} p(x,y) \mathbb{1}_{\mathsf{h}(x) \neq y}, c \right\} \\ &= \sum_{y \in \mathcal{Y}} p(x,y) \mathbb{1}_{\mathsf{h}(x) \neq y} \mathbb{1}_{r(x) > 0} + c \mathbb{1}_{r(x) \leq 0} - \min\left\{ \sum_{y \in \mathcal{Y}} p(x,y) \mathbb{1}_{\mathsf{h}(x) \neq y}, c \right\} \\ &+ \min\left\{ \sum_{y \in \mathcal{Y}} p(x,y) \mathbb{1}_{\mathsf{h}(x) \neq y}, c \right\} - \min\left\{ \inf_{h \in \mathcal{H}} \sum_{y \in \mathcal{Y}} p(x,y) \mathbb{1}_{\mathsf{h}(x) \neq y}, c \right\}. \end{aligned}$$
(14)

Note that by the property of the minimum, the second term can be upper bounded as

$$\begin{split} \min & \left\{ \sum_{y \in \mathcal{Y}} p(x, y) \mathbb{1}_{\mathsf{h}(x) \neq y}, c \right\} - \min \left\{ \inf_{h \in \mathcal{H}} \sum_{y \in \mathcal{Y}} p(x, y) \mathbb{1}_{\mathsf{h}(x) \neq y}, c \right\} \\ & \leq \sum_{y \in \mathcal{Y}} p(x, y) \mathbb{1}_{\mathsf{h}(x) \neq y} - \inf_{h \in \mathcal{H}} \sum_{y \in \mathcal{Y}} p(x, y) \mathbb{1}_{\mathsf{h}(x) \neq y} \\ & = \mathcal{C}_{\ell_{0-1}}(h, x) - \mathcal{C}^*_{\ell_{0-1}}(\mathcal{H}, x) \\ & \leq \Gamma_1(\mathcal{C}_{\ell}(h, x) - \mathcal{C}^*_{\ell}(\mathcal{H}, x)), \end{split}$$

where we use the \mathcal{H} -consistency bound of ℓ on the pointwise distribution δ_x that concentrates on a point x in the last inequality. Next, we will upper bound the first term. Note that the conditional risk and the best-in class conditional risk of $\ell_{\Phi,h}$ can be expressed as

$$C_{\ell_{\Phi,h}}(r,x) = \sum_{y \in \mathcal{Y}} p(x,y) \mathbb{1}_{\mathsf{h}(x) \neq y} \Phi(-r(x)) + c \Phi(r(x))$$

$$C_{\ell_{\Phi,h}}^*(\mathcal{R},x) = \inf_{r \in \mathcal{R}} \left(\sum_{y \in \mathcal{Y}} p(x,y) \mathbb{1}_{\mathsf{h}(x) \neq y} \Phi(-r(x)) + c \Phi(r(x)) \right).$$
(15)

Let $p_1 = \frac{\sum_{y \in \mathcal{Y}} p(x,y) \mathbb{1}_{h(x)\neq y}}{\sum_{y \in \mathcal{Y}} p(x,y) \mathbb{1}_{h(x)\neq y}+c}$ and $p_2 = \frac{c}{\sum_{y \in \mathcal{Y}} p(x,y) \mathbb{1}_{h(x)\neq y}+c}$. Then, the first term can be rewritten as 1

$$\begin{split} &\sum_{y \in \mathcal{Y}} p(x, y) \mathbb{1}_{\mathsf{h}(x) \neq y} \mathbb{1}_{r(x) > 0} + c \mathbb{1}_{r(x) \le 0} - \min \left\{ \sum_{y \in \mathcal{Y}} p(x, y) \mathbb{1}_{\mathsf{h}(x) \neq y}, c \right\} \\ &= \left(\sum_{y \in \mathcal{Y}} p(x, y) \mathbb{1}_{\mathsf{h}(x) \neq y} + c \right) \left[p_1 \mathbb{1}_{r(x) > 0} + p_2 \mathbb{1}_{r(x) \le 0} - \inf_{r \in \mathcal{R}} \left(p_1 \mathbb{1}_{r(x) > 0} + p_2 \mathbb{1}_{r(x) \le 0} \right) \right] \end{split}$$

By Lemma 21, we have

 \sim

$$p_{1}1_{r(x)>0} + p_{2}1_{r(x)\leq0} - \inf_{r\in\mathbb{R}} \left(p_{1}1_{r(x)>0} + p_{2}1_{r(x)\leq0} \right)$$

$$\leq \Gamma_{2} \left(p_{1}\Phi(-r(x)) + p_{2}\Phi(r(x)) - \inf_{r\in\mathbb{R}} (p_{1}\Phi(-r(x)) + p_{2}\Phi(r(x))) \right)$$

$$= \Gamma_{2} \left(\frac{\mathcal{C}_{\ell_{\Phi,h}}(r,x) - \mathcal{C}^{*}_{\ell_{\Phi,h}}(\mathcal{R},x)}{\sum_{y\in\mathcal{Y}} p(x,y)\mathbb{1}_{h(x)\neq y} + c} \right).$$

Therefore, the first term can be upper bounded as

$$\begin{split} &\sum_{y \in \mathcal{Y}} p(x,y) \mathbb{1}_{\mathsf{h}(x) \neq y} \mathbb{1}_{r(x) > 0} + c \mathbb{1}_{r(x) \leq 0} - \min\left\{\sum_{y \in \mathcal{Y}} p(x,y) \mathbb{1}_{\mathsf{h}(x) \neq y}, c\right\} \\ &= \left(\sum_{y \in \mathcal{Y}} p(x,y) \mathbb{1}_{\mathsf{h}(x) \neq y} + c\right) \left[p_1 \mathbb{1}_{r(x) > 0} + p_2 \mathbb{1}_{r(x) \leq 0} - \inf_{r \in \mathcal{R}} \left(p_1 \mathbb{1}_{r(x) > 0} + p_2 \mathbb{1}_{r(x) \leq 0} \right) \right] \\ &\leq \left(\sum_{y \in \mathcal{Y}} p(x,y) \mathbb{1}_{\mathsf{h}(x) \neq y} + c\right) \Gamma_2 \left(\frac{\mathcal{C}_{\ell_{\Phi,h}}(r,x) - \mathcal{C}^*_{\ell_{\Phi,h}}(\mathcal{R},x)}{\sum_{y \in \mathcal{Y}} p(x,y) \mathbb{1}_{\mathsf{h}(x) \neq y} + c} \right) \\ &\leq \left\{ \frac{\Gamma_2 \left(\mathcal{C}_{\ell_{\Phi,h}}(r,x) - \mathcal{C}^*_{\ell_{\Phi,h}}(\mathcal{R},x) \right)}{(1+c)\Gamma_2 \left(\frac{\mathcal{C}_{\ell_{\Phi,h}}(r,x) - \mathcal{C}^*_{\ell_{\Phi,h}}(\mathcal{R},x)}{c} \right)} \right. \text{ otherwise} \end{split}$$

where we use the fact that $c \leq \sum_{y \in \mathcal{Y}} p(x, y) \mathbb{1}_{h(x)\neq y} + c \leq 1 + c$ and Γ_2 is non-decreasing in the last inequality. After upper bounding the first term and the second term in (14) as above, taking the expectation on both sides, using the fact that Γ_1 and Γ_2 are concave, we obtain

$$\begin{split} & \mathbb{E}_{X} \Big[\mathbb{C}_{\mathsf{L}_{\mathrm{abs}}}(h, r, x) - \mathbb{C}_{\mathsf{L}_{\mathrm{abs}}}^{*}(\mathcal{H}, \mathcal{R}, x) \Big] \\ & \leq \begin{cases} \Gamma_{2} \Big(\mathbb{E}_{X} \Big[\mathbb{C}_{\ell_{\Phi,h}}(r, x) - \mathbb{C}_{\ell_{\Phi,h}}^{*}(\mathcal{R}, x) \Big] \Big) + \Gamma_{1} \Big(\mathbb{E}_{X} \Big[\mathbb{C}_{\ell}(h, x) - \mathbb{C}_{\ell}^{*}(\mathcal{H}, x) \Big] \Big) & \Gamma_{2} \text{ is linear} \\ (1+c)\Gamma_{2} \Big(\frac{1}{c} \mathbb{E}_{X} \Big[\mathbb{C}_{\ell_{\Phi,h}}(r, x) - \mathbb{C}_{\ell_{\Phi,h}}^{*}(\mathcal{R}, x) \Big] \Big) + \Gamma_{1} \Big(\mathbb{E}_{X} \Big[\mathbb{C}_{\ell}(h, x) - \mathbb{C}_{\ell}^{*}(\mathcal{H}, x) \Big] \Big) & \text{otherwise} \end{cases} \end{split}$$

Since the three expected terms can be expressed as

we have

$$\begin{split} \mathcal{E}_{\mathsf{L}_{\mathrm{abs}}}(h,r) - \mathcal{E}^{*}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) + \mathcal{M}_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) &\leq \Gamma_{1}(\mathcal{E}_{\ell}(h) - \mathcal{E}^{*}_{\ell}(\mathcal{H}) + \mathcal{M}_{\ell}(\mathcal{H})) \\ &+ (1+c)\Gamma_{2}\left(\frac{\mathcal{E}_{\ell_{\Phi,h}}(r) - \mathcal{E}^{*}_{\ell_{\Phi,h}}(\mathcal{R}) + \mathcal{M}_{\ell_{\Phi,h}}(\mathcal{R})}{c}\right), \end{split}$$

where the constant factors (1 + c) and $\frac{1}{c}$ can be removed when Γ_2 is linear.

G.5. Proof of realizable $(\mathcal{H}, \mathcal{R})$ -consistency bounds for single-stage surrogates (Theorem 9)

Theorem 9 Assume that \mathcal{H} and \mathcal{R} are closed under scaling. Let $\Psi(0) = 0$ and Φ satisfy Assumption 1. Then, for any ℓ that satisfies Assumption 2, the following $(\mathcal{H}, \mathcal{R})$ -consistency bound holds for any $(\mathcal{H}, \mathcal{R})$ -realizable distribution, $h \in \mathcal{H}$ and $r \in \mathcal{R}$:

$$\mathcal{E}_{\mathsf{L}_{\mathrm{abs}}}(h,r) - \mathcal{E}^*_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) \le \mathcal{E}_{\mathsf{L}}(h,r) - \mathcal{E}^*_{\mathsf{L}}(\mathcal{H},\mathcal{R}).$$

Proof It is straightforward to see that L serves as an upper bound for L_{abs} when ℓ serves as an upper bound for ℓ_{0-1} under Assumption 1. By definition, for any $(\mathcal{H}, \mathcal{R})$ -realizable distribution, there exists $h^* \in \mathcal{H}$ and $r^* \in \mathcal{R}$ such that $\mathcal{E}_{L_{abs}}(h^*, r^*) = \mathcal{E}^*_{L_{abs}}(\mathcal{H}, \mathcal{R}) = 0$. Then, by the assumption that \mathcal{H} and \mathcal{R} are closed under scaling, for any $\nu > 0$,

$$\begin{aligned} \mathcal{E}_{\mathsf{L}}^{*}(\mathcal{H},\mathcal{R}) &\leq \mathcal{E}_{\mathsf{L}}(\nu h^{*},\nu r^{*}) \\ &= \mathbb{E}[\mathsf{L}(\nu h^{*},\nu r^{*},x,y) \mid r^{*} < 0]\mathbb{P}(r^{*} < 0) + \mathbb{E}[\mathsf{L}(\nu h^{*},\nu r^{*},x,y) \mid r^{*} > 0]\mathbb{P}(r^{*} > 0) \end{aligned}$$

Next, we investigate the two terms. The first term is when $r^* < 0$, then we must have c = 0 since the data is realizable. By taking the limit, we obtain:

$$\lim_{\nu \to +\infty} \mathbb{E}[\mathsf{L}(\nu h^*, \nu r^*, x, y) \mid r^* < 0] \mathbb{P}(r^* < 0)$$

$$= \lim_{\nu \to +\infty} \mathbb{E}[\ell(\nu h^*, x, y) \Phi(-\alpha \nu r^*(x)) + \Psi(c) \Phi(\beta \nu r^*(x)) \mid r^* < 0] \mathbb{P}(r^* < 0)$$

$$= \lim_{\nu \to +\infty} \mathbb{E}[\ell(\nu h^*, x, y) \Phi(-\alpha \nu r^*(x)) \mid r^* < 0] \mathbb{P}(r^* < 0) \qquad (c = 0 \text{ and } \Psi(0) = 0)$$

$$= 0. \qquad \text{(by the Lebesgue dominated convergence theorem and } \lim_{t \to +\infty} \Phi(t) = 0)$$

The second term is when $r^* > 0$, then we must have $h^*(x, y) - \max_{y' \neq y} h^*(x, y') > 0$ since the data is realizable. Thus, using the fact that $\lim_{\nu \to +\infty} \ell(\nu h^*, x, y) = 0$ and taking the limit, we obtain

$$\lim_{\nu \to +\infty} \mathbb{E}[\mathsf{L}(\nu h^*, \nu r^*, x, y) \mid r^* < 0] \mathbb{P}(r^* < 0)$$

=
$$\lim_{\nu \to +\infty} \mathbb{E}[\ell(\nu h^*, x, y)\Phi(-\alpha\nu r^*(x)) + \Psi(c)\Phi(\beta\nu r^*(x)) \mid r^* < 0] \mathbb{P}(r^* < 0)$$

= 0

(by the Lebesgue dominated convergence theorem, $\lim_{t\to+\infty} \Phi(t) = 0$, $\lim_{\nu\to+\infty} \ell(\nu h^*, x, y) = 0$)

Therefore, by combining the above two analysis, we obtain

$$\mathcal{E}_{\mathsf{L}}^{*}(\mathcal{H},\mathcal{R}) \leq \lim_{\nu \to +\infty} \mathcal{E}_{\mathsf{L}}(\nu h^{*},\nu r^{*}) = 0.$$

By using the fact that L serves as an upper bound for L_{abs} and $\mathcal{E}^*_{L_{abs}}(\mathcal{H},\mathcal{R}) = 0$, we conclude that

$$\mathcal{E}_{\mathsf{L}_{\mathrm{abs}}}(h,r) - \mathcal{E}^*_{\mathsf{L}_{\mathrm{abs}}}(\mathcal{H},\mathcal{R}) \leq \mathcal{E}_{\mathsf{L}}(h,r) - \mathcal{E}^*_{\mathsf{L}}(\mathcal{H},\mathcal{R}).$$

G.6. Proof of realizable $(\mathcal{H}, \mathcal{R})$ -consistency for two-stage surrogates (Theorem 12)

Theorem 12 Assume that \mathcal{H} and \mathcal{R} are closed under scaling. Let ℓ be any multi-class surrogate loss that is realizable \mathcal{H} -consistent with respect to ℓ_{0-1} when \mathcal{H} is closed under scaling and Φ satisfies Assumption 1. Let \hat{h} be the minimizer of \mathcal{E}_{ℓ} and \hat{r} be the minimizer of $\mathcal{E}_{\ell_{\Phi,\hat{h}}}$. Then, for any $(\mathcal{H}, \mathcal{R})$ -realizable distribution, $\mathcal{E}_{\mathsf{L}_{abs}}(\hat{h}, \hat{r}) = 0$.

Proof It is straightforward to see that $\ell_{\Phi,h}$ upper bounds the abstention loss L_{abs} under Assumption 1. By definition, for any $(\mathcal{H}, \mathcal{R})$ -realizable distribution, there exists $h^* \in \mathcal{H}$ and $r^* \in \mathcal{R}$ such that $\mathcal{E}_{\mathsf{L}_{abs}}(h^*, r^*) = 0$. Let \hat{h} be the minimizer of \mathcal{E}_{ℓ} and \hat{r} be the minimizer of $\mathcal{E}_{\ell_{\Phi,\hat{h}}}$. Then, using the fact that $\ell_{\Phi,h}$ upper bounds the abstention loss $\mathcal{E}_{\mathsf{L}_{abs}}$, we have $\mathcal{E}_{\mathcal{E}_{\mathsf{L}_{abs}}}(\hat{h}, \hat{r}) \leq \mathcal{E}_{\ell_{\Phi,\hat{h}}}(\hat{r})$.

Next, we analyze two cases. If for a point x, abstention happens, that is $r^*(x) < 0$, then we must have c = 0 since the data is realizable. Therefore, there exists an optimal r^{**} abstaining all the points with zero cost: $r^{**}(x) < 0$ for all $x \in \mathcal{X}$. Then, by the assumption that \mathcal{R} is closed under scaling and the Lebesgue dominated convergence theorem, using the fact that $\lim_{t\to+\infty} \Phi(t) = 0$, we obtain

$$\mathcal{E}_{\mathsf{L}_{abs}}(\hat{h},\hat{r}) \leq \mathcal{E}_{\ell_{\Phi,\hat{h}}}(\hat{r})$$

$$\leq \lim_{\nu \to +\infty} \mathcal{E}_{\ell_{\Phi,\hat{h}}}(\nu r^{**}) \qquad (\hat{r} \text{ is the minimizer of } \mathcal{E}_{\ell_{\Phi,\hat{h}}})$$

$$= \lim_{\nu \to +\infty} \mathbb{E}\left[\mathbb{1}_{\hat{\ell}_{e_{\nu}}} - \Phi(-\nu r^{**}(x)) + c\Phi(\nu r^{**}(x))\right] \qquad (By (5))$$

$$= \lim_{\nu \to +\infty} \mathbb{E} \Big[\mathbb{1}_{\hat{\mathsf{h}}(x) \neq y} \Phi(-\nu r^{**}(x)) \Big]$$
(By (5))
$$= \lim_{\nu \to +\infty} \mathbb{E} \Big[\mathbb{1}_{\hat{\mathsf{h}}(x) \neq y} \Phi(-\nu r^{**}(x)) \Big]$$
(c = 0)

= 0. (by the Lebesgue dominated convergence theorem and
$$\lim_{t\to+\infty} \Phi(t) = 0$$
)

On the other hand, if no abstention occurs for any point, that is $r^*(x) > 0$ for any $x \in \mathcal{X}$, then we must have $\mathbb{1}_{h^*(x)\neq y} = 0$ for all $(x, y) \in \mathcal{X} \times \mathcal{Y}$ since the data is realizable. Using the fact that ℓ is realizable \mathcal{H} -consistent with respect to ℓ_{0-1} when \mathcal{H} is closed under scaling, we obtain $\mathbb{1}_{\hat{h}(x)\neq y} = 0$ for all $(x, y) \in \mathcal{X} \times \mathcal{Y}$. Then, by the assumption that \mathcal{R} is closed under scaling and the Lebesgue dominated convergence theorem, using the fact that $\lim_{t\to+\infty} \Phi(t) = 0$, we obtain

= 0. (by the Lebesgue dominated convergence theorem and $\lim_{t\to+\infty} \Phi(t) = 0$)

This completes the proof.