

Large-Margin Classification with Multiple Decision Rules

Patrick K. Kimes¹ D. Neil Hayes² J. S. Marron^{1,3}
Yufeng Liu^{1,3,4,†} for the Alzheimer’s Disease Neuroimaging Initiative*

Abstract: Binary classification is a common statistical learning problem in which a model is estimated on a set of covariates for some outcome indicating the membership of one of two classes. In the literature, there exists a distinction between hard and soft classification. In soft classification, the conditional class probability is modeled as a function of the covariates. In contrast, hard classification methods only target the optimal prediction boundary. While hard and soft classification methods have been studied extensively, not much work has been done to compare the actual tasks of hard and soft classification. In this paper we propose a spectrum of statistical learning problems which span the hard and soft classification tasks based on fitting multiple decision rules to the data. By doing so, we reveal a novel collection of learning tasks of increasing complexity. We study the problems using the framework of large-margin classifiers and a class of piecewise linear convex surrogates, for which we derive statistical properties and a corresponding sub-gradient descent algorithm. We conclude by applying our approach to simulation settings and a magnetic resonance imaging (MRI) dataset from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) study.

Keywords: Conditional Probability Estimation, Excess Risk Bounds, Statistical Machine Learning, Supervised Learning

¹Department of Statistics and Operations Research, ²Lineberger Comprehensive Cancer Center, ³Department of Biostatistics, ⁴Carolina Center for Genome Sciences, University of North Carolina, Chapel Hill, NC 27599. [†]Correspondence to: Yufeng Liu (yfliu@email.unc.edu). *Data used in preparation of this article were obtained from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators can be found at: http://adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf.

1 Introduction

Classification is one of the most widely applied and well studied problems in supervised learning. Given a training set of observed covariates and outcomes, similar to the usual regression problem, in classification, the outcome is modeled as a function of the set of covariates. However, in contrast to standard regression with a continuous response variable, classification describes the setting where the outcome is a discrete class label. While generalizations to more than two classes exist, in this paper we focus on the standard binary problem where the label takes one of two possible values, typically denoted by $+1$ and -1 .

Given such a dataset, commonly, the goal is to build a model, either to predict the class of a new observation from the covariate space, or to estimate the probability of each class as a function of the covariates. The tasks correspond respectively to hard and soft classification. Briefly, we refer to methods which only target the optimal prediction rule as hard classifiers, and those which produce estimates of class probability as soft classifiers. Examples of hard classifiers include the support vector machine (SVM) [1, 2] and ψ -learning [3, 4], and examples of soft classifiers include logistic regression and other likelihood-based approaches. Often, soft classifiers are also used to obtain hard classification rules by predicting the class with greater estimated probability. These rules are commonly referred to as plug-in classifiers. While hard classification rules do not directly provide conditional class probability estimates, several approaches have been proposed for estimating class probabilities based on hard classifiers, including those of [5] and [6]. As such, methods which may be traditionally viewed as soft and hard classifiers are often used for either task. Naturally, a question of interest is: how are hard and soft classifiers related, and how do they differ in practice?

Recently, [7] introduced the Large-margin Unified Machines (LUM) family of margin-based classifiers, shedding some light on the relationship between hard and soft classifiers. The LUM family connects several popular margin-based classification methods, including SVM, distance-weighted discrimination (DWD) [8], and a new hybrid logistic loss. Their approach was further extended to the multi-category case in [9]. Margin-based approaches to classification are popular in practice for their accuracy and computational efficiency in both low and high-dimensional settings. While a flexible family of margin-based classifiers, the LUM approach examines only a specific parameterized collection of classifiers along the gradient of soft to hard classification. In this paper, we similarly focus on connecting hard and soft margin-based methods. However, we consider a more natural approach based on connecting the tasks of hard and soft classification rather than specific hard and soft classifiers. Specifically, we propose a novel framework of binary learning problems which

may be formulated as partial or full estimation of the conditional class probability based on fitting an arbitrary number of boundaries to the data. As an example, suppose we are interested in separating patients into four disease risk groups based on clinical measurements. One possible approach is to group patients according to whether their conditional probability of being positive for the disease is less than 25%, between 25% to 50%, between 50% to 75%, or greater than 75%. In this setting, the emphasis is not on the accuracy of class probability estimates, but instead, on the correct stratification of individuals into risk groups. Therefore, only partial estimation of the conditional class probability is required; in particular, at the three boundaries, 25%, 50%, and 75%. While stratification of the patient classes is possible using a soft classifier, an approach directly targeting the three boundaries may provide improved stratification by requiring weaker assumptions on the entire form of the underlying conditional class probability.

In addition to hard and soft classification, the proposed framework also encompasses rejection-option classification [10, 11, 12, 13] and weighted classification [14, 15], two other well-studied binary learning problems. Briefly, the rejection-option problem expands on standard binary classification by introducing a third option to reject, where neither label is predicted. Notably, it can be shown that the decision to reject directly corresponds to a prediction that the probability of belonging to either class does not exceed a specified threshold. Since the task requires estimation of more than a single classification boundary, but less than the full class conditional probability, it may be viewed as an intermediate problem to hard and soft classification, as in the example given above. Applications of rejection option classification include certain medical settings where predictions should only be made when a level of certainty is obtained. Additionally, weighted classification extends the standard classification problem by accounting for differences or biases in class populations. We define these problems more formally, along with hard and soft classification, in Section 2.

The remainder of this paper is organized as follows. In the first part of Section 2 we provide a review of margin-based learning. Then, in the remainder of Section 2, we define our family of binary learning problems and introduce a corresponding *theoretical loss*, which generalizes the standard misclassification error to connect class prediction with probability estimation. In Section 3 we provide necessary and sufficient conditions for consistency of a surrogate loss function, and propose a class of consistent piecewise linear surrogates akin to the SVM hinge loss for binary classification. In Section 4, we present theoretical bounds on the empirical performance of classification rules obtained using surrogate loss functions. In Section 5, we provide a sub-gradient descent (SGD) algorithm for solving the corresponding

optimization problem using the proposed piecewise linear surrogates. We then illustrate the behavior of our generalized family of classifiers using simulation in Section 6, and a real data example from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) database in Section 7. We conclude in Section 8 with a discussion of the proposed framework.

2 Methodology

In this section, we first briefly introduce margin-based classifiers, and formally define the notion of classification consistency for loss functions. We then state the general form of our unified framework of problems and introduce a corresponding family of theoretical loss functions which encompasses the standard misclassification error as a special case.

2.1 Margin-Based Classifiers

Let $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$ denote a training set of n covariate–label pairs drawn from $\mathcal{X} \times \mathcal{Y}$ according to some unknown distribution $\mathcal{P}(\mathbf{X}, Y)$. For binary problems, $\mathcal{Y} = \{-1, +1\}$ is used to denote the label space, and often $\mathcal{X} \subset \mathbb{R}^p$, with $p \geq 1$. Given a training set, margin-based classifiers minimize a penalized loss over a class, \mathcal{F} , of margin functions, $f : \mathcal{X} \rightarrow \mathbb{R}$. Typically, the corresponding optimization problem is written as:

$$\min_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n \underbrace{L(y_i f(\mathbf{x}_i))}_{\text{loss}} + \underbrace{\lambda J(f)}_{\text{penalty}}, \quad (1)$$

where $L : \mathbb{R} \rightarrow \mathbb{R}$ is a loss function defined with respect to the functional margin, $yf(\mathbf{x})$, and $J : \mathcal{F} \rightarrow \mathbb{R}$ is some roughness measure on \mathcal{F} with corresponding tuning parameter $\lambda \geq 0$. Both hard and soft classification may be formulated as margin-based problems. In the case of hard classification, with a little abuse of notation, we use $\hat{Y} \in \mathcal{Y}$ to denote a predicted class label, and $\hat{Y} : \mathbb{R} \rightarrow \mathcal{Y}$ to denote a prediction rule on \mathbb{R} . In margin-based classification, $\hat{Y}(\cdot)$ is combined with a margin function, $f \in \mathcal{F}$, to obtain predictions in \mathcal{Y} . Most commonly, in hard classification the sign rule, $\hat{Y}(f(\mathbf{X})) = \text{sign}(f(\mathbf{x}))$, is used, assuming $f(\mathbf{x}) \neq 0$ almost surely (*a.s.*). Thus, given a new (\mathbf{x}^*, y^*) pair with $f(\mathbf{x}^*) \neq 0$, correct classification is obtained if and only if $y^* f(\mathbf{x}^*) > 0$. Since the functional margin, $yf(x)$, serves as an approximate measure for classification correctness, the loss function, L , in (1) is often chosen to be a non-increasing function over $yf(\mathbf{x})$. A natural choice of L in

hard classification is the misclassification error, or 0–1 loss, given by:

$$\ell_{0-1}(Y, \hat{Y}) = \mathbf{I}\{\hat{Y} \neq Y\}, \quad (2)$$

where $\mathbf{I}\{\cdot\}$ is used to denote the indicator function. Using the sign rule, the loss may be equivalently written over the class of margin functions as: $L_{0-1}(Yf(\mathbf{X})) = \mathbf{I}\{Yf(\mathbf{X}) < 0\}$. However, direct optimization of the non-convex and discontinuous loss, L_{0-1} , is NP-hard and often infeasible in practice. Thus, continuous convex losses, called surrogates, are commonly used instead. Choices of the surrogate loss function corresponding to existing margin-based classifiers include the SVM hinge loss, $L(z) = \max\{0, 1 - z\}$, logistic loss, $L(z) = \log(1 + e^{-z})$, and the DWD loss, $L(z) = \frac{1}{4z} \cdot \mathbf{I}\{z \geq \frac{1}{2}\} + (1 - z) \cdot \mathbf{I}\{z < \frac{1}{2}\}$. Finally, the penalty term, $J(\cdot)$ is used to prevent over-fitting and improve generalizability of the resulting classifier. The amount of penalization is commonly determined by cross-validation over a grid of λ values. Here, we note that while in the literature there exists a natural theoretical loss for hard classification, i.e. the 0–1 loss, there is no equivalent theoretical loss targeting consistent probability estimation for soft classification. In addition to providing a spectrum of theoretical loss functions covering soft and hard classifications at the two extremes, our proposed framework also naturally defines precisely such a theoretical loss for the soft classification problem (Figure 2C).

In Section 1, we briefly discussed the learning tasks of rejection-option and weighted classification. As with hard and soft classification, these tasks may also be formulated as margin-based problems. We next describe how rejection-option classification may be formulated as a problem of the form (1). Borrowing the notation of [12], we use 0 to denote the rejection option such that a prediction, \hat{Y}_{rej} , takes values in $\mathcal{Y}_{rej} = \{+1, 0, -1\}$. Then, for some pre-specified *rejection cost* $\pi \in (0, \frac{1}{2})$, they propose the following theoretical loss for rejection-option classification:

$$\ell_{rej,\pi}(Y, \hat{Y}_{rej}) = \begin{cases} 1 & \text{if } \hat{Y}_{rej} \neq Y, \hat{Y}_{rej} \neq 0 \\ \pi & \text{if } \hat{Y}_{rej} = 0 \\ 0 & \text{otherwise} \end{cases}. \quad (3)$$

To express the loss as a function over $Yf(\mathbf{X})$, [12] propose the prediction rule $\hat{Y}_{rej}(f(\mathbf{X}); \delta) = \mathbf{I}\{|Yf(\mathbf{X})| > \delta\} \cdot \text{sign}(Yf(\mathbf{X}))$ for some appropriately chosen $\delta > 0$. Then, $\ell_{rej,\pi}$ may be

written as the following generalized 0–1 loss on $Yf(\mathbf{X})$:

$$L_{rej,\pi}(Yf(\mathbf{X}); \delta) = (1 - \pi)\mathbf{I}\{Yf(\mathbf{X}) \leq -\delta\} + \pi\mathbf{I}\{Yf(\mathbf{X}) < \delta\}.$$

We finally consider the task of weighted classification. In contrast to the problems mentioned thus far, to fit the form of (1), weighted classification requires specifying separate theoretical loss functions for observations from the +1 and –1 classes, denoted by $\ell_{w,\pi}^+$ and $\ell_{w,\pi}^-$. For simplicity, we use $\ell_{w,\pi}^Y$ to denote the loss for both classes. Similar to hard classification, the task is to predict class labels in $\mathcal{Y} = \{+1, -1\}$. The loss function depends on a weight parameter, π , which accounts for imbalances between the two classes. Commonly, π is constrained to the interval $(0, 1)$ without loss of generality. Then, for fixed weight π , the weighted loss is given by:

$$\begin{aligned} \ell_{w,\pi}^Y(Y, \hat{Y}) &= \mathbf{I}\{Y = +1\} \cdot \ell_{w,\pi}^+(\hat{Y}) + \mathbf{I}\{Y = -1\} \cdot \ell_{w,\pi}^-(\hat{Y}), \\ \ell_{w,\pi}^+(\hat{Y}) &= (1 - \pi) \cdot \mathbf{I}\{\hat{Y} \neq +1\}, \\ \ell_{w,\pi}^-(\hat{Y}) &= \pi \cdot \mathbf{I}\{\hat{Y} \neq -1\}. \end{aligned} \tag{4}$$

Note that the standard 0–1 loss corresponds to the special case of the weighted loss (4) when equal weight is assigned to the two classes with $\pi = \frac{1}{2}$. Using the same prediction rule as for hard classification, $\hat{Y}(f(\mathbf{x})) = \text{sign}(f(\mathbf{x}))$, the loss over the functional margin may be written:

$$\begin{aligned} L_{w,\pi}^Y(Yf(\mathbf{X})) &= \mathbf{I}\{Y = +1\} \cdot L_{w,\pi}^+(Yf(\mathbf{X})) + \mathbf{I}\{Y = -1\} \cdot L_{w,\pi}^-(Yf(\mathbf{X})), \\ L_{w,\pi}^+(Yf(\mathbf{X})) &= (1 - \pi) \cdot \mathbf{I}\{Yf(\mathbf{X}) < 0\}, \\ L_{w,\pi}^-(Yf(\mathbf{X})) &= \pi \cdot \mathbf{I}\{Yf(\mathbf{X}) < 0\}. \end{aligned}$$

As with the usual 0–1 loss, optimization of $L_{rej,\pi}$ and $L_{w,\pi}^Y$ is NP-hard, and in practice should be approximated using a convex surrogate loss. In the next section, we introduce the notion of consistency, an important statistical property of surrogate loss functions.

2.2 Classification Consistency

Much work has been done to study the statistical properties of classifiers of the *loss+penalty* form given in (1) [16, 17, 18, 19]. Of these, consistency of loss functions is one of the most fundamental. In general, a loss function is called consistent for a margin-based learning problem if it recovers in expectation the optimal rule, often called the Bayes rule, to the theoretical loss

function, e.g. l_{0-1} , $l_{rej,\pi}$ or $l_{w,\pi}^Y$. More formally, for a theoretical loss function, ℓ , and a surrogate loss, ϕ , let $Y_\ell^*(\mathbf{X}) = \operatorname{argmin}_{Y^*} \mathbb{E}_{Y|\mathbf{X}}\{\ell(Y, Y^*)\}$ and $f_\phi^*(\mathbf{X}) = \operatorname{argmin}_f \mathbb{E}_{Y|\mathbf{X}}\{\phi(Yf(\mathbf{X}))\}$ denote the Bayes rule and ϕ -optimal margin function, respectively. Then, we call ϕ consistent for ℓ if $\widehat{Y}_\ell(f_\phi^*(\mathbf{X})) = Y_\ell^*(\mathbf{X})$, where \widehat{Y}_ℓ is the appropriate prediction rule, e.g. the sign function. Equivalently, using the margin-based formulation of the theoretical loss, L , and letting $f_L^*(\mathbf{X}) = \operatorname{argmin}_f \mathbb{E}_{Y|\mathbf{X}}\{L(Yf(\mathbf{X}))\}$ denote the L -optimal margin function, consistency may be expressed as $\widehat{Y}_\ell(f_\phi^*(\mathbf{X})) = \widehat{Y}_\ell(f_L^*(\mathbf{X}))$. For rejection-option classification, the Bayes optimal rule is given by:

$$Y_{rej,\pi}^*(\mathbf{X}) = \begin{cases} +1 & \text{if } p(\mathbf{X}) \geq 1 - \pi \\ 0 & \text{if } p(\mathbf{X}) \in (\pi, 1 - \pi) . \\ -1 & \text{if } p(\mathbf{X}) \leq \pi \end{cases} \quad (5)$$

The Bayes optimal rule for weighted classification is given by:

$$Y_{w,\pi}^*(\mathbf{X}) = \begin{cases} +1 & \text{if } p(\mathbf{X}) > \pi \\ -1 & \text{if } p(\mathbf{X}) \leq \pi \end{cases} . \quad (6)$$

For hard classification, the Bayes optimal rule corresponds to $Y_{w,0.5}$, and consistency is often referred to as Fisher consistency or classification calibrated [18]. While no theoretical loss has been proposed for soft classification, using $p(\mathbf{X}) = \mathbb{P}(Y = +1|\mathbf{X})$ to denote the conditional class probability at $\mathbf{X} \in \mathcal{X}$, commonly, ϕ is called consistent for soft classification if there exists some monotone mapping, $C : \mathbb{R} \rightarrow [0, 1]$ such that $C(f_\phi^*(\mathbf{X})) = p(\mathbf{X})$. Naturally, $C(\cdot)$ may be viewed as an extension of the prediction rules $\widehat{Y}(\cdot)$ and $\widehat{Y}_{rej}(\cdot; \delta)$ given for hard and rejection-option classification. Necessary and sufficient conditions for Fisher, rejection-option, and probability estimation consistency have been described in [12, 20, 21].

In this paper, we propose a novel framework for unifying hard, soft, rejection-option, and weighted classification through a generalized formulation of their corresponding theoretical losses, corresponding Bayes optimal rules, and necessary and sufficient conditions for consistency. Our generalized formulation not only provides a platform for comparing existing binary classification tasks, but also introduces an entire family of new tasks which fills the gap between these problems. We next formally introduce our unified framework of binary learning problems.

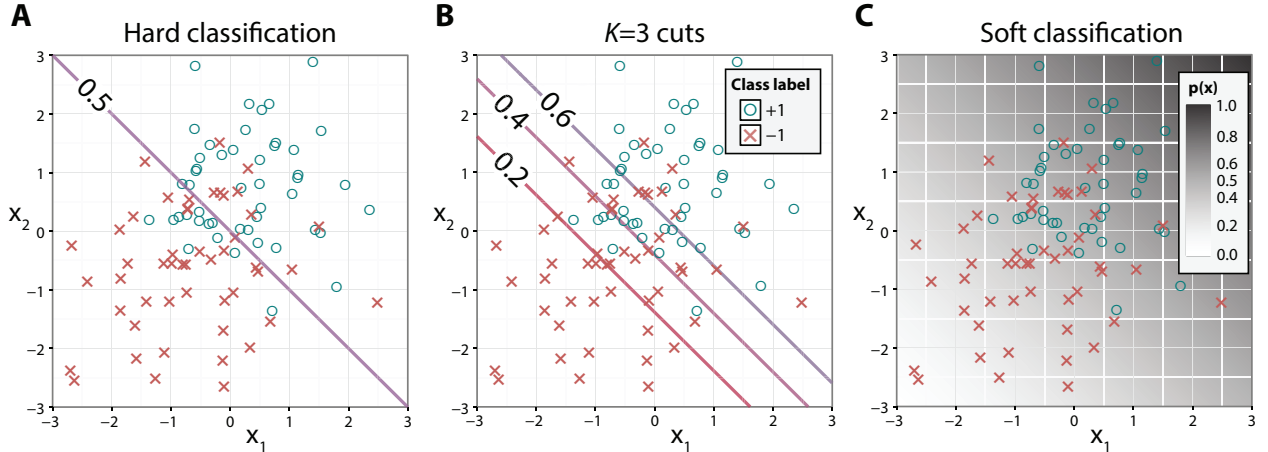


Figure 1: Boundaries are shown separating the input space, \mathbb{R}^2 into the $K + 1$ prediction sets, Ω_π . A sample of 100 observations drawn from the underlying class populations are overlaid to show the distribution over the space. (A) The $K = 1$ boundary for $\pi = \{0.5\}$ corresponding to hard classification is shown by the separating hyperplane corresponding to the set $\{\mathbf{x} \in \mathbb{R}^2 : p(\mathbf{x}) = 0.5\}$. (B) The set of $K = 3$ boundaries are shown for $\pi = \{0.2, 0.4, 0.6\}$ separating the 4 prediction sets. (C) The soft classification results are shown spanning the entire set of $\pi_k \in (0, 1)$. As $K \rightarrow \infty$, moving from hard to soft classification, the set of learning problems becomes increasingly complex.

2.3 Unified Framework

First, we note that all of the classification tasks described in Section 2.1 may be formulated as learning problems which target partial or complete estimation of the conditional class probability, $p(\mathbf{x})$. We propose our framework of unified margin-based learning problems based on this insight. Let Ω_π denote the ordered $(K + 1)$ partition of the interval $[0, 1]$ obtained by splitting at $\pi = \{\pi_1, \dots, \pi_K\}$, where $0 < \pi_1 < \dots < \pi_K < 1$. Assume $p(\mathbf{x}) \neq \pi_k$ *a.s.* for all k , such that observations belong to only a single region of interval. Letting $\pi_0 = 0$ and $\pi_{K+1} = 1$ for ease of notation, we write:

$$\Omega_\pi = \{\omega_0, \dots, \omega_K\},$$

where $\omega_0 = [\pi_0, \pi_1]$, and $\omega_k = (\pi_k, \pi_{k+1}]$, for $k \geq 1$. As our framework, we propose the class of problems which target a partition of the covariate space, \mathcal{X} , into the $K + 1$ regions, $\{\mathbf{x} : p(\mathbf{x}) \in \omega_k\}$. In Figure 1, we show a sample of 100 observations drawn from the same underlying distribution, $\mathcal{P}(\mathbf{X}, Y)$ along with optimal solutions to three representative problems from our proposed framework. Note that the extreme cases of $K = 1$ with $\pi =$

$\{0.5\}$ (Figure 1A), and $K = \infty$ with $\boldsymbol{\pi}$ dense on $(0, 1)$ (Figure 1C) correspond to hard and soft classification, respectively. We discuss these connections in more detail later in this section. To illustrate the spectrum of problems in our framework, we also show a new intermediate problem in Figure 1B, with $K = 3$ and $\boldsymbol{\pi} = \{0.2, 0.4, 0.6\}$.

Formally, we define our framework as the collection of minimization tasks of a theoretical loss which generalizes the 0–1 loss, over the collection of rules $\mathcal{G}_{\boldsymbol{\pi}} = \{g : \mathcal{X} \rightarrow \Omega_{\boldsymbol{\pi}}\}$. Recall the weighted 0–1 loss, ℓ_w , for weighted classification described above. For positive and negative class weights $(1 - \pi)$ and π where $\pi \in (0, 1)$, the weighted 0–1 loss has corresponding Bayes boundary at $\{\boldsymbol{x} : p(\boldsymbol{x}) = \pi\}$. Problems under our framework may be viewed as the task of simultaneously estimating K such boundaries. Intuitively, we formulate our theoretical loss as the average of K weighted 0–1 loss functions with corresponding weights $\boldsymbol{\pi}$. Throughout, we use $\ell_{\boldsymbol{\pi}}^+(g(\boldsymbol{x}))$ and $\ell_{\boldsymbol{\pi}}^-(g(\boldsymbol{x}))$ to denote the loss for positive and negative class observations, respectively. As with the weighted loss, we use $\ell_{\boldsymbol{\pi}}^Y$ to denote the loss for both classes:

$$\ell_{\boldsymbol{\pi}}^Y(g(\mathbf{X})) = \frac{2}{K} \sum_{k=1}^K \ell_{\pi_k}^Y(g(\mathbf{X})), \quad (7)$$

where

$$\begin{aligned} \ell_{\pi_k}^+(g(\mathbf{X})) &= (1 - \pi_k) \cdot \mathbf{I}\{g(\mathbf{X}) \leq \pi_k\}, \\ \ell_{\pi_k}^-(g(\mathbf{X})) &= \pi_k \cdot \mathbf{I}\{g(\mathbf{X}) > \pi_k\}, \end{aligned}$$

and the notion of inequalities is extended to elements of $\Omega_{\boldsymbol{\pi}}$ such that $(\pi_j, \pi_{j+1}] \leq \pi_k$ if $\pi_{j+1} \leq \pi_k$ and $(\pi_j, \pi_{j+1}] > \pi_k$ if $\pi_j \geq \pi_k$. As we show in Supplementary Section S1, our theoretical loss encompasses the usual 0–1 loss, its weighted variant, and the rejection-option loss proposed by [12]. The multiplicative constant, 2, is included in (7) such that $\ell_{\boldsymbol{\pi}}^Y$ reduces precisely to the usual 0–1 loss when $\boldsymbol{\pi} = \{0.5\}$. Note that since $\ell_{\boldsymbol{\pi}}^Y$ is effectively the average of K indicator functions scaled by 2, the function takes values in the interval $[0, 2]$. In Figure 2, we show $\ell_{\boldsymbol{\pi}}^Y$ as a function of $g(\boldsymbol{x}) \mapsto \Omega_{\boldsymbol{\pi}}$, corresponding to the problems in Figure 1. Along the horizontal axis, the range $[0, 1]$ is split into corresponding $\omega_j = (\pi_j, \pi_{j+1}]$ intervals. Note that the loss function is constant within each interval, giving the appearance of a step function, except in the extreme case when $K = \infty$. As K increases, the theoretical loss becomes smoother, with the limit at $\boldsymbol{\pi} = (0, 1)$ corresponding to the proposed theoretical loss for consistent soft classification described in Section 2.1. Additionally, note that while the loss functions, $\ell_{\boldsymbol{\pi}}^+$ and $\ell_{\boldsymbol{\pi}}^-$, are symmetric in Panels A and C of Figure 2, the same is not

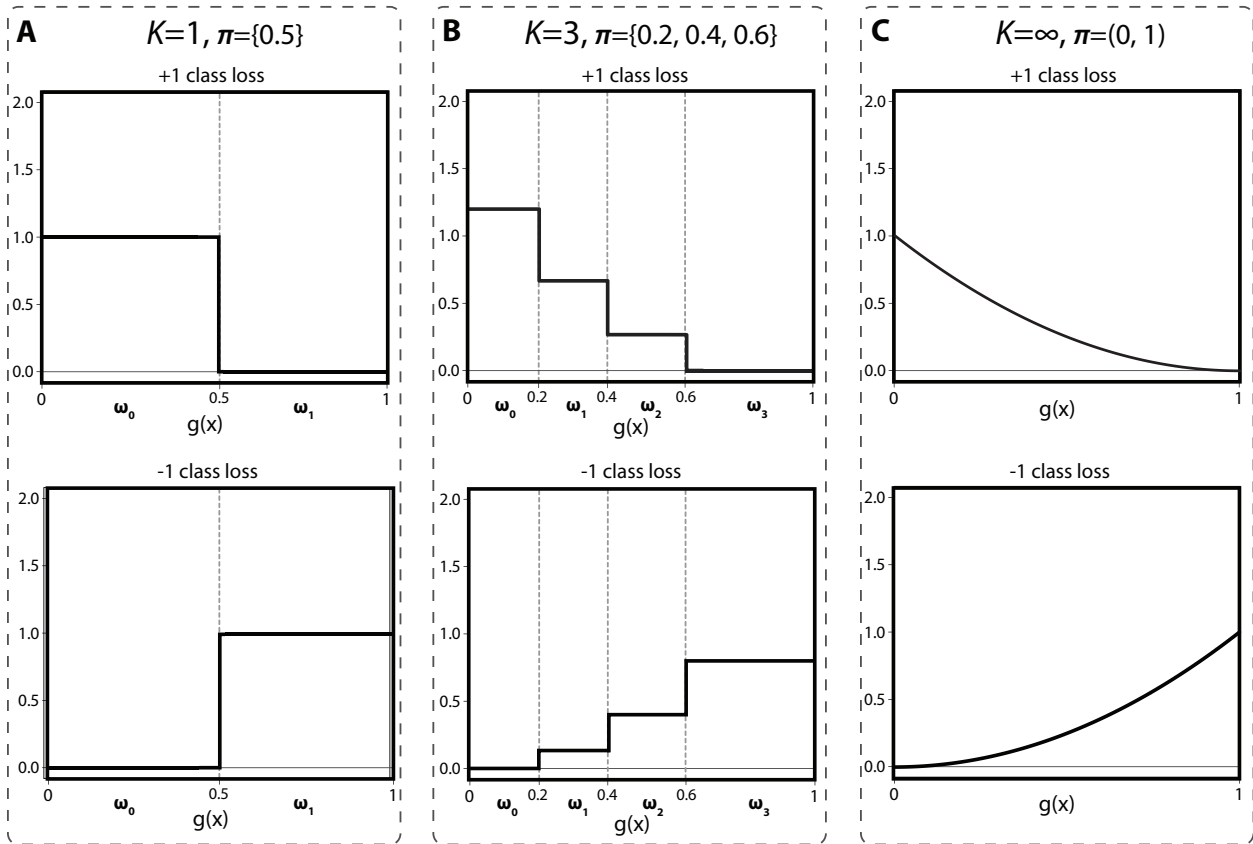


Figure 2: Examples of the theoretical loss functions, ℓ_{π}^Y , for observations from the positive and negative classes over $g(\mathbf{x})$ are shown for boundaries, π at (A) $\{0.5\}$, (B) $\{0.2, 0.4, 0.6\}$, and (C) $(0, 1)$, corresponding to the problems shown in Figure 1. The theoretical loss generalizes the standard 0–1 loss given in (A) by incorporating K steps. As K increases and the problem approaches soft classification, the theoretical loss becomes noticeably smoother.

true for the loss functions in Panel B. This is due to the fact that the boundaries of interest, π , are symmetric between the two classes, i.e. $\pi = 1 - \pi$, when $\pi = \{0.5\}$ and $\pi = (0, 1)$, but not when $\pi = \{0.2, 0.4, 0.6\}$.

The following result states that the class of problems defined by our theoretical loss indeed corresponds to the proposed framework of learning tasks. That is, the Bayes optimal rule given by $W_{\pi}^*(\mathbf{X}) = \operatorname{argmin}_g \mathbb{E}_{Y|\mathbf{X}} \{\ell_{\pi}^Y(g(\mathbf{X}))\}$, is precisely the partitioning task described above.

Theorem 1. *For fixed K and π defined as above, the Bayes optimal rule for the theoretical*

loss (7) is given by:

$$\begin{aligned} W_{\pi}^*(\mathbf{X}) &= \operatorname{argmin}_{g \in \mathcal{G}_{\pi}} \mathbb{E}_{Y|\mathbf{X}} \{ \ell_{\pi}^Y(g(\mathbf{X})) \} \\ &= \sum_{k=0}^K \omega_k \cdot \mathbf{I}\{p(\mathbf{X}) \in \omega_k\}. \end{aligned}$$

In addition to the results of Theorem 1, the theoretical loss functions for hard (2), rejection-option (3), and weighted (4) classification can be derived as special cases of (7). This is shown by first noting the equivalence of Ω_{π} to \mathcal{Y} and \mathcal{Y}_{rej} based on the Bayes optimal rules, (5) and (6). From this equivalence, (3) and (4) can be obtained directly from (7). For soft classification, we derive a new theoretical loss from the limiting form of (7):

$$\begin{aligned} \ell_{\pi}^Y(g(\mathbf{X})) &= \lim_{K \rightarrow \infty} \frac{2}{K} \sum_{k=1}^K \ell_{\pi_k}^Y(g(\mathbf{X})), \\ &= (\mathbf{I}\{Y = +1\} - g(\mathbf{X}))^2. \end{aligned}$$

The resulting theoretical loss is shown in Figure 2C. Since $\Omega_{\pi} = (0, 1)$, the Bayes rule is simply the conditional class probability, $g(\mathbf{X}) = p(\mathbf{X})$, corresponding to soft classification. All proofs, and a more complete derivation of these results may be found in the Supplementary Materials.

As with the problems described in Section 2.1, optimization of ℓ_{π} with respect to $g \in \mathcal{G}_{\pi}$ is NP-hard. Thus, we first reformulate ℓ_{π} as a function on \mathbb{R} to express the optimization over a collection of margin functions, \mathcal{F} . We then propose in Section 3 to solve the approximate problem using convex surrogate loss functions. Generalizing the approach of [12] for rejection-option classification, we frame the optimization task over the class of margin functions, \mathcal{F} , using a prediction rule $C : \mathbb{R} \times \mathbb{R}^K \rightarrow \Omega_{\pi}$ of the form:

$$C(f(\mathbf{x}); \boldsymbol{\delta}) = \sum_{k=0}^K \omega_k \cdot \mathbf{I}\{f(\mathbf{x}) \in (\delta_{k-1}, \delta_k]\}, \quad (8)$$

for monotone increasing $\boldsymbol{\delta} = \{\delta_1, \dots, \delta_K\}$, and $\delta_0 = -\infty$, $\delta_{K+1} = \infty$. Intuitively, each δ_k corresponds to the π_k -boundary along the range of the margin function, $f(\mathbf{X})$. As is common in margin-based learning, we write the theoretical loss as the following function over $Yf(\mathbf{X})$:

$$L_{\pi}^Y(Yf(\mathbf{X}); \boldsymbol{\delta}) = \ell_{\pi}^Y(C(f(\mathbf{X}); \boldsymbol{\delta}))$$

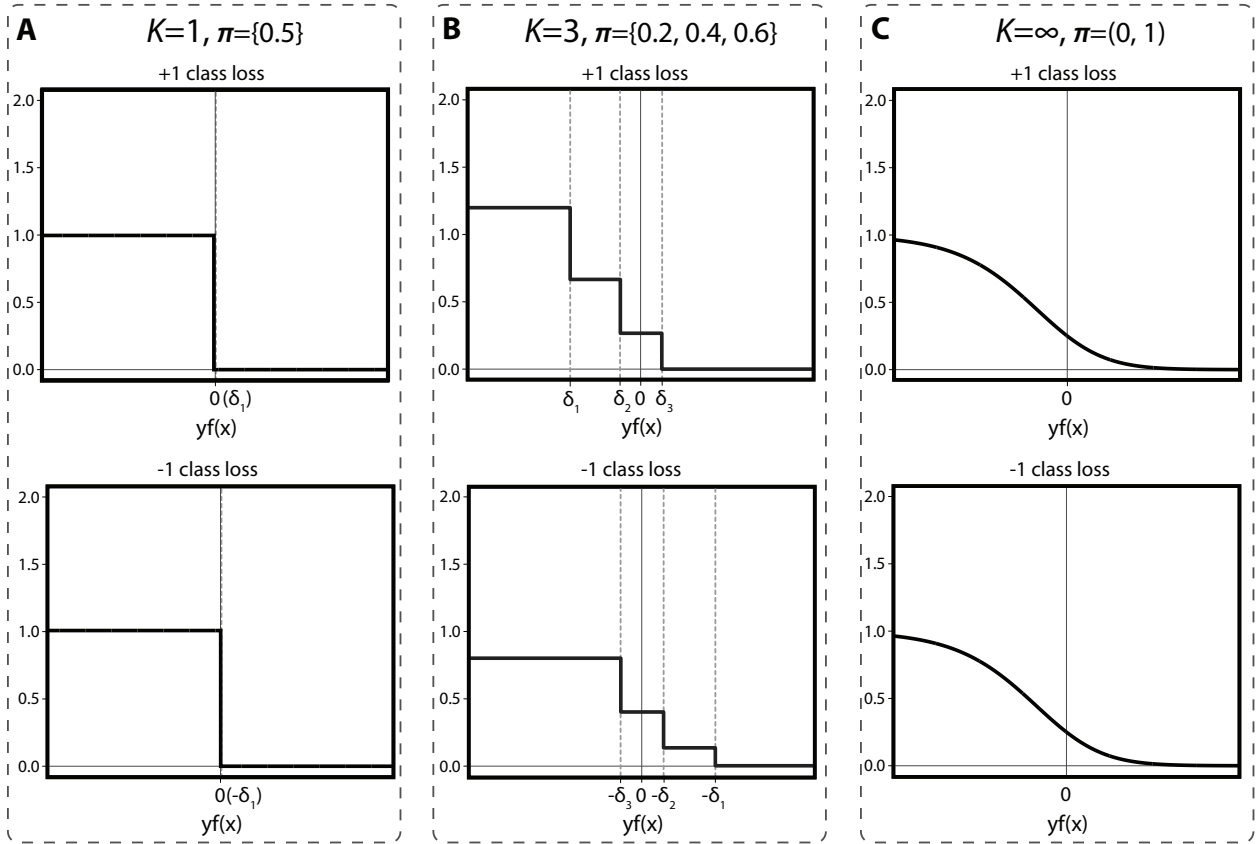


Figure 3: Examples of the margin-based formulation of the theoretical loss function, $L_{\pi}^Y(\cdot, \boldsymbol{\delta})$, for observations from the positive and negative classes over $yf(\mathbf{x})$ are shown for boundaries, $\boldsymbol{\pi}$, at (A) $\{0.5\}$, (B) $\{0.2, 0.4, 0.6\}$, and (C) $(0, 1)$, using well-chosen $\boldsymbol{\delta}$.

$$= \begin{cases} \frac{2}{K} \sum_{k=1}^K (1 - \pi_k) \cdot \mathbf{I}\{Y f(\mathbf{X}) \leq \delta_k\} & \text{if } Y = +1 \\ \frac{2}{K} \sum_{k=1}^K \pi_k \cdot \mathbf{I}\{Y f(\mathbf{X}) < -\delta_k\} & \text{if } Y = -1 \end{cases}. \quad (9)$$

In Figure 3, we plot the corresponding margin-based formulations of the theoretical loss functions shown in Figure 2, with well chosen $\boldsymbol{\delta}$. Intuitively, both $L_{\pi}^+(\cdot; \boldsymbol{\delta})$ and $L_{\pi}^-(\cdot; \boldsymbol{\delta})$ are non-increasing on $yf(\mathbf{x})$. We also note that ℓ_{π}^- and $L_{\pi}^-(\cdot; \boldsymbol{\delta})$ differ by a reflection along the vertical axis since $L_{\pi}^-(\cdot; \boldsymbol{\delta})$ is defined with respect to $yf(\mathbf{x}) = -f(\mathbf{x})$. Given the margin-based formulation (9), we propose to solve our class of problems using convex surrogate loss functions. In the following section, we first present necessary and sufficient conditions for a surrogate loss to be consistent to (7). We then introduce a class of consistent piecewise linear surrogates, which includes the SVM hinge loss as a special case.

3 Convex Surrogate Loss Functions

Since the proposed theoretical loss function (7) and its margin-based reformulation (9) are discontinuous and non-convex for any finite choice of K and $\boldsymbol{\pi}$, empirical minimization can quickly become intractable. Therefore, we propose to instead minimize a convex surrogate loss over the class of margin functions, as in hard and soft classification. In this section, we first provide necessary and sufficient conditions for a surrogate loss to be consistent for (7) with fixed K and $\boldsymbol{\pi}$. Then, we introduce a class of convex piecewise linear surrogates which includes the SVM hinge loss as a special case. Intuitively, the piecewise linear surrogates each consist of K non-zero segments, corresponding to the K boundaries, $\boldsymbol{\pi}$. In the limit, as $\boldsymbol{\pi}$ becomes dense on $(0, 1)$, the piecewise linear surrogate tends towards a smooth loss, as in Panel C of Figures 2 and 3.

3.1 Consistency

Throughout this section, we assume K and $\boldsymbol{\pi}$ to be fixed. First, let ϕ^+ and ϕ^- denote a pair of convex surrogate loss functions for $\ell_{\boldsymbol{\pi}}^+$ and $\ell_{\boldsymbol{\pi}}^-$. Further, let $f_{\phi}^* = \operatorname{argmin}_f \mathbb{E}_{Y|\mathbf{X}}\{\phi^Y(Yf(\mathbf{X}))\}$ denote the ϕ^Y -optimal rule over the class of all measurable functions. We call ϕ^Y consistent if there exists $\boldsymbol{\delta} \in \mathbb{R}^K$ such that the prediction rule (8) satisfies $C(f_{\phi}^*(\mathbf{x}); \boldsymbol{\delta}) = W_{\boldsymbol{\pi}}^*(\mathbf{x})$, i.e. if there exists a known monotone mapping from the ϕ^Y -optimal rule to the $K+1$ partition of \mathcal{X} to $\Omega_{\boldsymbol{\pi}}$. The following result provides necessary and sufficient conditions for the consistency of the surrogate loss ϕ^Y to $\ell_{\boldsymbol{\pi}}^Y$.

Theorem 2. *A pair of convex surrogate loss functions, ϕ^Y , are consistent for $\ell_{\boldsymbol{\pi}}^Y$ if and only if there exists $\boldsymbol{\delta} \in \mathbb{R}^K$ such that for each $k = 1, \dots, K$: $\phi^{+'}(\delta_k)$ and $\phi^{-'}(-\delta_k)$ exist, $\phi^{+'}(\delta_k)$ and $\phi^{-'}(-\delta_k) < 0$, and*

$$\frac{\phi^{-'}(-\delta_k)}{\phi^{-'}(-\delta_k) + \phi^{+'}(\delta_k)} = \pi_k. \tag{10}$$

Naturally, any surrogate loss satisfying the conditions of Theorem 2 for some $\boldsymbol{\pi}$, must also satisfy the set of conditions for any subset of the boundaries, $\boldsymbol{\pi}' \subseteq \boldsymbol{\pi}$. Thus, for surrogate loss functions consistent for soft classification, i.e. when $\boldsymbol{\pi} = (0, 1)$, there exists an appropriate $\boldsymbol{\delta}$ for any possible K and $\boldsymbol{\pi}$. Similar intuition is used to justify the use of soft classification based plug-in classifiers described in Section 1. Examples of surrogate losses consistent for soft classification include the logistic, squared hinge, exponential, and DWD losses. Values

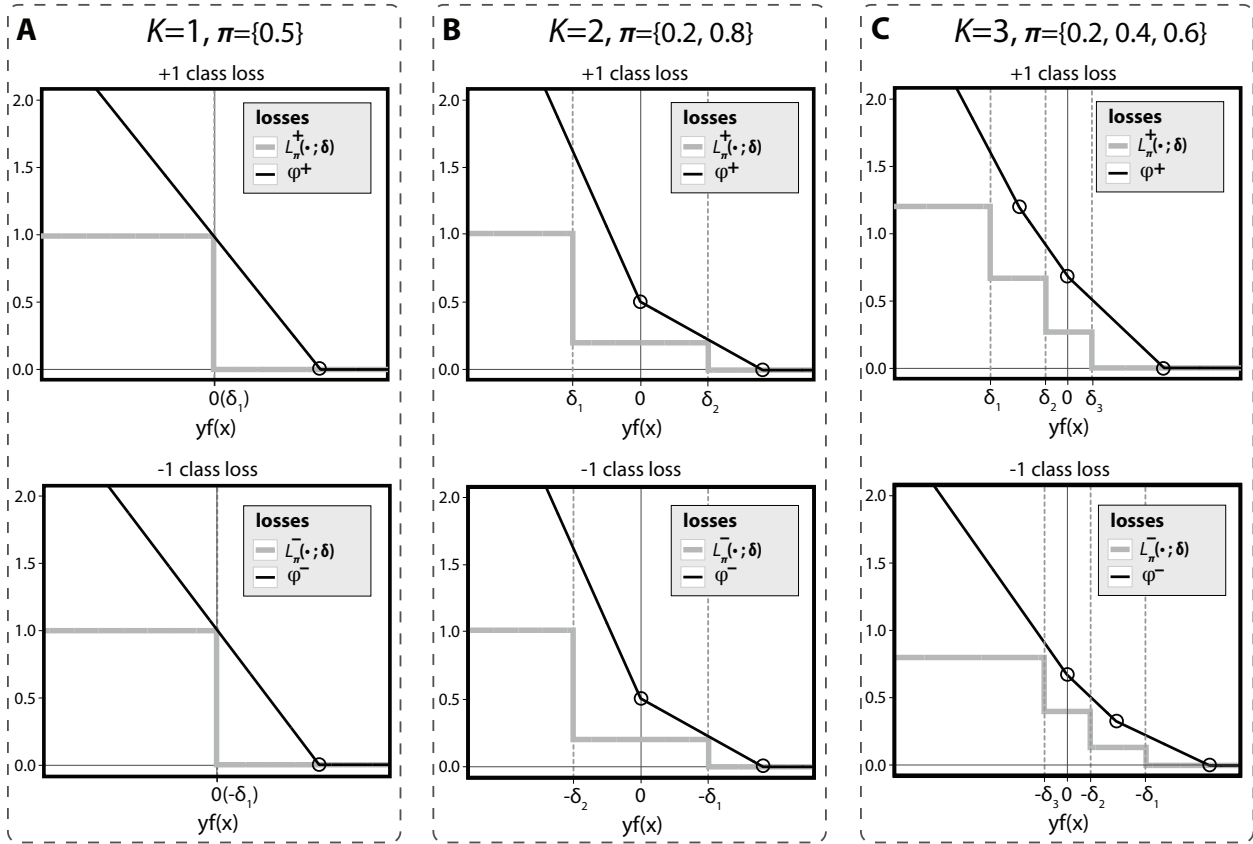


Figure 4: Examples of piecewise linear surrogates are shown along with the corresponding theoretical loss, $L_{\pi, \delta}$ for (A) $\pi = \{0.5\}$ (hard classification), (B) $\pi = \{0.2, 0.8\}$ (rejection-option classification), and (C) $\pi = \{0.2, 0.4, 0.6\}$.

of δ_k such that the conditions of Theorem 2 are met for these loss functions are provided in Corollaries 3-8 of [12]. In the next section, we introduce a class of piecewise linear surrogates which, similar to the SVM loss for hard classification, satisfy consistency for the π of interest, but not for any $\pi' \supset \pi$. We refer to such a piecewise linear surrogate as being minimally consistent for a corresponding set of boundaries, π . In contrast to soft classification losses which satisfy consistency for all $\pi \subseteq (0, 1)$, minimally consistent surrogates are well-tuned for a given ℓ_{π}^Y , and may provide improved stratification of \mathcal{X} to the sets, Ω_{π} .

3.2 Piecewise Linear Surrogates

Throughout, we use φ^+ and φ^- to denote piecewise linear surrogates. To build intuition, in the columns of Figure 4, we show examples of φ^Y for $K = 1, 2, 3$, corresponding to hard clas-

sification, rejection-option classification, and the new problem shown in Figure 1B. Circles are used to highlight the hinges, i.e. non-differentiable points, along the piecewise linear loss functions. The corresponding margin-based theoretical loss, $L_{\boldsymbol{\pi}}^Y(\cdot; \boldsymbol{\delta})$, is also shown in each panel using appropriately chosen $\boldsymbol{\delta}$. First, note that the losses in Panels A and B of Figure 4 correspond to the standard SVM hinge loss and generalized hinge loss of [11], respectively. Consider the new surrogate losses in Figure 4C for boundaries at $\boldsymbol{\pi} = \{0.2, 0.4, 0.6\}$. Note that φ^+ and φ^- each consist of K non-zero linear segments. Furthermore, each linear segment only spans a single δ_k or $-\delta_k$ for φ^+ and φ^- , respectively. We will refer to these pairs of linear segments as the π_k -consistent segments. This construction allows for the consistency of the surrogate loss for each $\pi_k \in \boldsymbol{\pi}$ to be controlled separately by the K pairs of π_k -consistent segments along the piecewise linear loss.

We formulate our collection of piecewise linear surrogate losses as the maximum of the K linear segments and 0. Consider first the surrogate loss for positive observations, φ^+ . Using $A^+(\pi)$, $B^+(\pi)$ to denote the intercept and slope of the π_k -consistent segment, we express the piecewise linear loss as:

$$\varphi^+(z) = \max\{0, A^+(\pi_1) + B^+(\pi_1) \cdot z, \dots, A^+(\pi_K) + B^+(\pi_K) \cdot z\}. \quad (11)$$

We similarly use $A^-(\pi)$ and $B^-(\pi)$ to denote the intercept and slope of the π_k -consistent segment for the negative class loss such that:

$$\varphi^-(z) = \max\{0, A^-(\pi_1) + B^-(\pi_1) \cdot z, \dots, A^-(\pi_K) + B^-(\pi_K) \cdot z\}. \quad (12)$$

By construction, the resulting piecewise linear losses are non-negative, convex and continuous. While (11) and (12) define a general class of piecewise linear losses, we focus on a subset of minimally consistent piecewise linear surrogates. In the following theorem, we provide a set of sufficient conditions for a piecewise linear loss to be minimally consistent for a specified $\boldsymbol{\pi}$.

Theorem 3. *Let $H^Y(\pi, \pi') = (A^Y(\pi) - A^Y(\pi')) / (B^Y(\pi') - B^Y(\pi))$ denote the location of the hinges along the respective loss functions between consecutive boundaries, $\pi < \pi'$. Then, φ^Y is a minimally consistent piecewise linear surrogate for $\boldsymbol{\pi}$ if the intercept and slope parameters, $A^Y(\pi)$ and $B^Y(\pi)$, satisfy the following conditions:*

(C1) $B^+(\pi)$ is non-decreasing, and $B^-(\pi)$ is non-increasing in π .

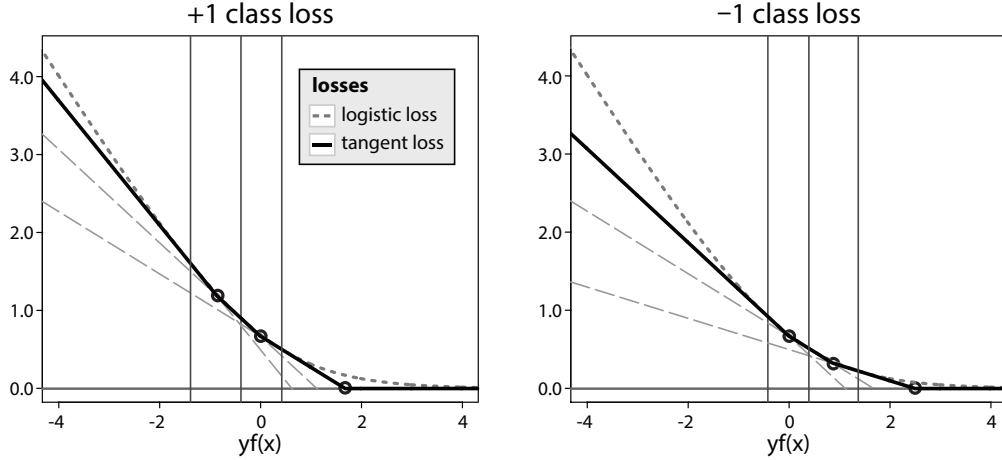


Figure 5: A pair of piecewise linear loss functions, φ^Y , obtained from the logistic loss for $\boldsymbol{\pi} = \{0.2, 0.4, 0.6\}$ is shown along with the logistic loss (dotted lines), and the set of tangent lines used to derive $A^Y(\boldsymbol{\pi})$ and $B^Y(\boldsymbol{\pi})$ (dashed lines).

(C2) The hinge points are such that:

$$\begin{aligned}
 -H^-(\pi_{k-1}, \pi_k) &= H^+(\pi_{k-1}, \pi_k) && \text{for } k = 2, \dots, K, \\
 H^+(\pi_{k-1}, \pi_k) &< H^+(\pi_k, \pi_{k+1}) && \text{for } k = 2, \dots, K-1, \\
 A^-(\pi_1)/B^-(\pi_1) &< H^+(\pi_1, \pi_2), \\
 A^+(\pi_K)/B^+(\pi_K) &> H^-(\pi_{K-1}, \pi_K).
 \end{aligned}$$

(C3) $B^+(\boldsymbol{\pi}), B^-(\boldsymbol{\pi})$ satisfy:

$$\frac{B^-(\pi_k)}{B^-(\pi_k) + B^+(\pi_k)} = \pi_k \quad \text{for } 1 \leq k \leq K.$$

Conditions (C1) and (C2) guarantee that the linear segments are well-ordered and non-degenerate along $Yf(\mathbf{X})$ with appropriately aligned hinge points. Condition (C3) guarantees the consistency of φ^Y to the corresponding $\ell_{\boldsymbol{\pi}}$. Most importantly, by aligning the hinge points, $-H^-(\pi_{k-1}, \pi_k)$ and $H^+(\pi_{k-1}, \pi_k)$, we ensure that there does not exist a $\delta \in \mathbb{R}$ such that (10) is satisfied for any $\boldsymbol{\pi} \notin \boldsymbol{\pi}$. Next, we present an approach to obtaining $A^Y(\boldsymbol{\pi})$ and $B^Y(\boldsymbol{\pi})$ which satisfy the conditions of Theorem 3 using the logistic loss as an example.

3.3 Logistic Derived Surrogates

In this section, we propose to construct piecewise linear losses by choosing $A^Y(\pi_k) + B^Y(\pi_k) \cdot z$ to be the tangent lines to the logistic loss at $Y \cdot \log(\frac{\pi_k}{1-\pi_k})$. A similar approach was used by [22] to construct a piecewise linear loss for the rejection-option problem. The following Proposition states that piecewise linear loss functions constructed using this approach satisfy the conditions of Theorem 3 for any choice of K and $\boldsymbol{\pi}$.

Proposition 1. *For fixed K and $\boldsymbol{\pi}$, let φ^Y be the piecewise linear loss constructed from the tangent lines to the logistic loss such that $A^Y(\pi)$ and $B^Y(\pi)$ are defined as:*

$$\begin{aligned} A^+(\pi) &= A^-(1 - \pi) = -\pi \log(\pi) - (1 - \pi) \log(1 - \pi), \\ B^+(\pi) &= B^-(1 - \pi) = -(1 - \pi). \end{aligned}$$

Then, φ^Y is a minimally consistent piecewise linear surrogate for $\boldsymbol{\pi}$ satisfying the conditions of Theorem 3.

In Figure 5, we illustrate the logistic-derived piecewise linear loss for $\boldsymbol{\pi} = \{0.2, 0.4, 0.6\}$. The logistic loss is shown by dotted lines, with the piecewise linear surrogate functions for the positive and negative classes shown in solid black. Thin vertical lines are used to denote the tangent points where the losses are equal, and thin dashed lines give the tangent lines to the logistic loss corresponding to $A^Y(\pi_k) + B^Y(\pi_k) \cdot yf(\mathbf{x})$ for $\pi_k \in \boldsymbol{\pi}$. Additionally, the non-differentiable hinge points are highlighted by circles. While the loss functions appear roughly equivalent within the region of the tangent points, the difference is non-negligible above and below these bounds. Notably, the piecewise linear losses diverge slower as $yf(\mathbf{x})$ tends to $-\infty$, suggesting the losses may be more robust to outliers [7]. Additionally, the logistic derived loss functions provide a natural spectrum for comparing the impact of targeting different partitions, $\Omega_{\boldsymbol{\pi}}$, on the same dataset. We explore these issues using simulation in Section 6.

4 Statistical Properties

We next derive statistical properties for surrogate loss functions to the theoretical loss, $\ell_{\boldsymbol{\pi}}^Y$. In Subsection 4.1, we first show that the excess risk with respect to $\ell_{\boldsymbol{\pi}}^Y$ may be bounded by the excess risk of a consistent surrogate loss. Then, in Subsection 4.2, we use these risk bounds to derive convergence rates for the empirical minimizer of a surrogate loss to the

Bayes optimal rule. Our results generalize and extend those derived for the particular case of rejection-option classification in [10, 11, 12], to an arbitrary number of boundaries.

4.1 Excess Risk Bounds

For a rule $g \in \mathcal{G}_\pi$, we define the ℓ_π^Y -risk of g to be the expected loss of the rule, denoted by $R(g) = \mathbb{E}_{Y, \mathbf{X}}\{\ell_\pi^Y(g(\mathbf{X}))\}$. In statistical machine learning, a natural measure of the performance of a rule is its excess risk: $\Delta R(g) = R(g) - R(W_\pi^*)$, where $R(W_\pi^*) = \min_{g \in \mathcal{G}_\pi} R(g)$ such that $\Delta R(g) \geq 0$. In this section, we derive convergence rates on $\Delta R(g)$ for rules obtained using consistent surrogate loss functions. For a surrogate loss ϕ^Y , we similarly define the ϕ -risk and excess ϕ -risk over the class of margin functions, \mathcal{F} , to be $Q(f) = \mathbb{E}_{Y, \mathbf{X}}\{\phi^Y(Yf(\mathbf{X}))\}$ and $\Delta Q(f) = Q(f) - Q(f_\phi^*)$. To obtain convergence rates on $\Delta R(g)$, we first show that under certain conditions, the excess ϕ -risk of a margin function f can be used to bound the corresponding excess ℓ_π^Y -risk of $g = C(f; \boldsymbol{\delta})$. Using this bound, we then derive rates of convergence on $\Delta R(g)$ through rates of convergence on $\Delta Q(g)$. The following additional notation is used to denote excess conditional ℓ_π^Y -risk and excess conditional ϕ -risk:

$$\begin{aligned} R_p(g) &:= \mathbb{E}_{Y|\mathbf{X}}\{\ell_\pi^Y(g(\mathbf{X}))\}, & Q_p(f) &:= \mathbb{E}_{Y|\mathbf{X}}\{\phi^Y(Yf(\mathbf{X}))\}, \\ \Delta R_p(g) &:= R_p(g) - R_p(W_\pi^*), & \Delta Q_p(f) &:= Q_p(f) - Q_p(f_\phi^*). \end{aligned}$$

In the following results, we provide conditions under which there exists some function, $\rho: \mathbb{R} \rightarrow \mathbb{R}$, such that $\rho(\Delta Q(f))$ can be used to bound the corresponding $\Delta R(C(f; \boldsymbol{\delta}))$.

Theorem 4. *Let ϕ^Y be a consistent surrogate loss for ℓ_π^Y satisfying the conditions for Theorem 2 at $\boldsymbol{\delta}$. Furthermore, suppose there exist constants $C > 0$ and $s \geq 1$ such that for all k ,*

$$|p(\mathbf{X}) - \pi_k|^s \leq C^s \Delta Q_p(\delta_k). \quad (13)$$

Then,

$$\Delta R(C(f; \boldsymbol{\delta})) \leq C[2 \cdot \Delta Q(f)]^{1/s}.$$

The above bound may be tightened as in [12] by the additional assumption:

$$\mathbb{P}\{|p(\mathbf{X}) - \pi_k| \leq t\} \leq At^\alpha, \quad k = 1, \dots, K, \quad (14)$$

for some $\alpha \geq 0$, $A \geq 1$. The bound (14) generalizes the margin condition introduced by [23] and used in [10].

Theorem 5. *In addition to the assumptions of Theorem 4, assume that there exists $\alpha \geq 0$ and $A \geq 1$, such that (14) holds for $t \in [0, \min_k \{\pi_k - \pi_{k-1}, \pi_{k+1} - \pi_k\})$. Then, for some D depending on A, α ,*

$$\Delta R(C(f; \boldsymbol{\delta})) \leq D \cdot \Delta Q(f)^{1/(s+\beta-\beta s)}$$

where $\beta = \alpha/(1 + \alpha)$.

Note that when $\alpha = 0$, Theorem 5 provides the same bound as Theorem 4. However, as $\alpha \rightarrow \infty$, the bound becomes tighter, with $1/(s + \beta + \beta s)$ limiting to 1. While neither result depends explicitly on $\boldsymbol{\pi}$, Theorem 5 suggests that tighter bounds may be achieved by only targeting $\boldsymbol{\pi}$ such that the margin condition is satisfied with large α . This reiterates the motivating intuition for our proposed framework, in which we formalize a class of learning problems for settings where more information than hard classification is desired, but soft classification may not be appropriate.

Corresponding values of C and s for the exponential, logistic, squared hinge and DWD losses, are provided in Corollaries 13–16 of [12]. In the following result, we derive values of C and s for our class of minimally consistent piecewise linear surrogates.

Corollary 1. *For minimally consistent piecewise linear loss, φ^Y , defined as in (11) and (12) and satisfying the conditions of Theorem 3 for boundaries $\boldsymbol{\pi}$, the inequality (13) is satisfied by $s = 1$ and*

$$C = \max \left\{ -\frac{\pi_k}{B^-(\pi_k) \cdot |\delta_k - H_j|} : k = 1, \dots, K; j = 0, \dots, K \right\},$$

where H_0 is used to denote $A^-(\pi_1)/B^-(\pi_1)$, H_j to denote $H^+(\pi_j, \pi_{j+1})$ for $j = 2, \dots, K - 1$, and H_K to denote $A^+(\pi_K)/B^+(\pi_K)$.

Consider now a sequence of margin functions, $\{f_n\}_{n \geq 1}$. By Theorems 4 and 5, to show that the excess $\ell_{\boldsymbol{\pi}}^Y$ -risk, $\Delta R(C(f_n; \boldsymbol{\delta}))$, converges to 0 as $n \rightarrow \infty$, it suffices to show that $\Delta Q(f_n) \rightarrow 0$ as $n \rightarrow \infty$. In the following results, we derive convergence rates for $\Delta R(C(\cdot; \boldsymbol{\delta}))$ for the sequence of functions, $\{\hat{f}_n\}_{n \geq 1}$, where \hat{f}_n is used to denote the empirical minimizer of the surrogate loss over a training set of size n .

4.2 Rates of Convergence

In this section, we derive convergence results for two classes of surrogate loss functions separately. We first consider Lipschitz continuous and differentiable surrogate loss functions

which satisfy a modulus of convexity condition specified below. Examples of such loss functions include the exponential, logistic, squared hinge and DWD losses. We then separately consider the class of piecewise linear surrogates described in Section 3.

Let ϕ^Y denote a Lipschitz continuous and differentiable surrogate loss function. Assume that the corresponding ϕ -risk, $Q(\cdot)$, has modulus of convexity,

$$\delta(\epsilon) = \inf \left\{ \frac{Q(f) + Q(g)}{2} - Q\left(\frac{f+g}{2}\right) : \mathbb{E}[(f-g)^2(\mathbf{X})] \geq \epsilon^2 \right\} \quad (15)$$

satisfying $\delta(\epsilon) > c\epsilon^2$ for some $c > 0$. Furthermore, let $L < \infty$ denote the Lipschitz constant, such that $|\phi^y(\mathbf{x}) - \phi^y(\mathbf{x}')| \leq L|\mathbf{x} - \mathbf{x}'|$ for all $\mathbf{x}, \mathbf{x}' \in \mathbb{R}$ and $y = +1, -1$. Letting \mathcal{F}_B denote the class of uniformly bounded functions such that $|f| \leq B$ for all $f \in \mathcal{F}_B$, we use $N_n = N(\frac{1}{n}, L_\infty, \mathcal{F}_B)$ to denote the cardinality of the set of closed balls with radius $\frac{1}{n}$ in L_∞ needed to cover \mathcal{F}_B . Finally, as stated above, let $\hat{f}_n = \operatorname{argmin}_{f \in \mathcal{F}_B} \sum_{i=1}^n \phi^{y_i}(y_i f(\mathbf{x}_i))$ denote the empirical minimizer of ϕ^Y over the training set $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$. For the following corollary, we make use of Theorem 18 from [12] which provides a bound on the expected estimation error, $Q(\hat{f}_n) - \inf_{f \in \mathcal{F}_B} Q(f)$, for consistent loss functions satisfying the modulus of convexity condition stated above. Combining Theorem 18 of [12] with the excess risk bounds of Theorems 4 and 5, we obtain the following result.

Corollary 2. *If ϕ^Y satisfies the assumptions of Theorems 2 and 4, and has modulus of convexity (15) satisfying $\delta(\epsilon) > c\epsilon^2$ for some $c > 0$, then with probability at least $1 - \gamma$,*

$$\Delta R(C(\hat{f}_n; \boldsymbol{\delta})) \leq C \cdot 2^{1/s} \left\{ \inf_{f \in \mathcal{F}_B} \Delta Q(f) + \frac{3L}{n} + 8 \left(\frac{L^2}{2c} + \frac{LB}{3} \right) \frac{\log(N_n/\gamma)}{n} \right\}^{1/s}.$$

Furthermore, if the generalized margin condition of Theorem 5 holds, then with probability at least $1 - \gamma$,

$$\Delta R(C(\hat{f}_n; \boldsymbol{\delta})) \leq D \left\{ \inf_{f \in \mathcal{F}_B} \Delta Q(f) + \frac{3L}{n} + 8 \left(\frac{L^2}{2c} + \frac{LB}{3} \right) \frac{\log(N_n/\gamma)}{n} \right\}^{1/(s+\beta-\beta s)}, \quad (16)$$

for constants $C, D > 0$ defined as in Theorems 4 and 5.

From the bound on excess risk obtained in Corollary 2, corresponding rates of convergence can be derived based on the cardinality, N_n , of the class of functions, \mathcal{F}_B .

Due to the non-differentiability of the loss at hinge points, our class of piecewise linear surrogates do not satisfy the modulus of convexity condition (15). The following theorem provides separate convergence results for our class of minimally consistent piecewise linear

surrogates. Again, we use \mathcal{F}_B to denote a class of uniformly bounded functions, and let $\hat{f}_n = \operatorname{argmin}_{f \in \mathcal{F}_B} \sum_{i=1}^n \varphi^{y_i}(y_i f(\mathbf{x}_i))$ denote the empirical minimizer of φ^Y .

Theorem 6. *If φ^Y is a minimally consistent piecewise linear loss satisfying the conditions of Theorem 3, satisfying the generalized margin condition of Theorem 5, then with probability at least $1 - \gamma$,*

$$\Delta Q(\hat{f}_n) \leq \frac{3L}{n} + \frac{4LB}{3} \cdot G(\gamma) + \left(\left(\frac{4LB}{3} \cdot G(\gamma) \right)^2 + 8 \cdot B' \cdot G(\gamma) \right)^{1/2},$$

where $G(\gamma) = \log(N_n/\gamma)/n$, and $B' > 0$ is some constant depending on B , φ^Y , and margin constants A, α .

Combining Theorems 4, 5, and 6, we obtain the following corollary.

Corollary 3. *If φ^Y is a minimally consistent piecewise linear loss satisfying the assumptions of Theorems 2, 4, and 5, then with probability at least $1 - \gamma$,*

$$\Delta R(C(\hat{f}_n; \delta)) \leq D \left\{ \frac{3L}{n} + \frac{4LB}{3} \cdot G(\gamma) + \left(\left(\frac{4LB}{3} \cdot G(\gamma) \right)^2 + 8 \cdot B' \cdot G(\gamma) \right)^{1/2} \right\}^{1/(s+\beta-\beta s)}, \quad (17)$$

for constants $C, D > 0$ defined as in Theorems 4 and 5.

As in Theorem 5, while the convergence rate of Theorem 6 does not depend on $\boldsymbol{\pi}$ explicitly, it does depend on the parameters of the margin condition (14). Therefore, Theorem 6 further suggests the advantage of targeting $\boldsymbol{\pi}$ for which the data show strong separation with large α . Furthermore, in contrast to Theorem 18 of [12] which provides a bound on the expected estimation error, Theorem 6 bounds the total φ^Y -risk, including both the expected estimation error, and expected approximation error of the class of functions \mathcal{F}_B . As a result, while the bounds in Corollary 2 include the separate approximation error term, $\inf_{f \in \mathcal{F}_B} \Delta Q(f)$, the piecewise linear bound in Corollary 3, does not.

Based on the bounds in (16) and (17), rates of convergence can be obtained as in [12]. As an example, we consider the case when \mathcal{F}_B is the class of linear combinations of decision stumps, f_λ ,

$$f_\lambda(x) = \sum_{j=1}^M \lambda_j f_j(x)$$

where $\sum_j |\lambda_j| \leq B$, and $|f_j| < 1$. By (16) and (17), the same rate, $(M \log n/n)^{1/(s+\beta-\beta_s)}$, can be obtained as in [12] for both classes of surrogate losses considered above.

5 Computational Algorithm

For a piecewise linear surrogate, φ^Y , and convex penalty, $J(f)$, the objective (1) is a non-differentiable convex problem. Several approaches have been proposed for solving the similar non-differentiable and convex SVM objective, most commonly by reformulating (1) as a quadratic program (QP) with $2n$ constraints. The penalized objective (1) with φ^Y may also be formulated as a QP with $(K+1)n$ constraints. However, as with the SVM problem, the complexity of the problem grows almost cubically with the number of constraints, making the problem computationally intensive for moderately large K and n [24]. We therefore propose a projected sub-gradient descent algorithm similar to the PEGASOS algorithm [25].

We first rewrite (1) with piecewise linear surrogate, φ^Y defined as in (11) and (12) as:

$$\min_{h,b} \frac{1}{n} \sum_{i=1}^n \left(\max_{k=1,\dots,K} \{A^{y_i}(\pi_k) + B^{y_i}(\pi_k) \cdot y_i(h(\mathbf{x}_i) + b)\} \right)_+ + \frac{\lambda}{2} \|h\|_{\mathcal{H}}^2, \quad (18)$$

where $(z)_+ = \max\{0, z\}$, and \mathcal{H} is some Reproducing Kernel Hilbert Space (RKHS) with norm $\|\cdot\|_{\mathcal{H}}$ and corresponding kernel function $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$. Commonly, the margin function is formulated with a non-penalized intercept parameter, b . A more complete review of RKHS may be found in [26, 27]. In margin-based learning, kernel methods are commonly used to estimate non-linear classification boundaries. In the case of linear learning, i.e. $h(\mathbf{x}) = \langle \mathbf{w}, \mathbf{x} \rangle$ for $\mathbf{w} \in \mathbb{R}^p$, the penalty $\|h\|_{\mathcal{H}}^2$ reduces to $\|\mathbf{w}\|^2$ and (18) may be written as:

$$\min_{\mathbf{w},b} \frac{1}{n} \sum_{i=1}^n \left(\max_{k=1,\dots,K} \{A^{y_i}(\pi_k) + B^{y_i}(\pi_k) \cdot y_i(\langle \mathbf{w}, \mathbf{x}_i \rangle + b)\} \right)_+ + \frac{\lambda}{2} \|\mathbf{w}\|^2.$$

We next describe our iterative algorithm for the linear learning setting. Let $\mathbf{w}^{(m)}$ and $b^{(m)}$ denote the estimated parameters at the m -th iteration. Furthermore, at each iteration, let B_i^* denote the sub-gradient of φ^{y_i} at $\langle \mathbf{w}^{(m)}, \mathbf{x}_i \rangle + b^{(m)}$ for $i = 1, \dots, n$. Using a decreasing step-size parameter, $\eta_m = (\lambda m)^{-1}$, we iterate the following updates until $\mathbf{w}^{(m)}$ and $b^{(m)}$ converge:

1. $\mathbf{w}^{(m)} = \mathbf{w}^{(m-1)} + \eta_m (\frac{1}{n} \sum_i B_i^* y_i \mathbf{x}_i - \lambda \mathbf{w}^{(m-1)})$,
2. $b^{(m)} = b^{(m-1)} + \eta_m (\frac{1}{n} \sum_i B_i^* y_i)$,

$$3. [\mathbf{w}^{(m)}, b^{(m)}] = \min\left\{1, \frac{\lambda^{-1/2}}{\|[\mathbf{w}^{(m)}, b^{(m)}]\|}\right\}[\mathbf{w}^{(m)}, b^{(m)}],$$

where B_i^* is used to denote the sub-gradient of φ^{y_i} at $y_i(\langle \mathbf{x}_i, \mathbf{w} \rangle + b)$. The final projection step is included to ensure $\|[\mathbf{w}^{(m)}, b^{(m)}]\|^2 \leq \lambda^{-1}$ at each iteration [25, 28]. In the following section, we apply our projected sub-gradient descent algorithm to simulated datasets to illustrate the utility of our class of problems.

6 Simulations

In this section, we use simulations to illustrate the performance achieved by targeting different binary learning problems. Namely, we compare the performance of several minimal consistent piecewise linear losses against the standard logistic classifier, when the underlying conditional class probability, $p(\mathbf{X})$, is piecewise constant. Piecewise linear loss functions are derived from the logistic loss as described in Section 3.3, and the sets of boundaries, $\boldsymbol{\delta}$, are chosen by the tangent points to the logistic loss. In each simulation, we consider piecewise linear losses with $\boldsymbol{\pi}_1 = \{1/2\}$, $\boldsymbol{\pi}_2 = \{1/3, 2/3\}$, and $\boldsymbol{\pi}_3 = \{1/4, 2/4, 3/4\}$. All methods are tuned over a grid of penalty parameters $\lambda \in \{2^{-15}, 2^{-14}, \dots, 2^{10}\}$, using training and tuning sets of 100 observations each. Piecewise linear classifiers and the logistic classifier are tuned with respect to the correspond theoretical loss (7) and likelihood function, respectively. The performance of each estimated model is evaluated using a test set of 10,000 observations. Each simulation was replicated 100 times.

6.1 Simulation 1

In this setting, data are simulated uniformly from $[-8, 8] \times [-1, 1]^{p-1}$ for $p = 2, 10, 50$, subject to a random rotation in the p -dimensional space. We consider three variations of this setting, in which the data were simulated with underlying conditional class probability, defined with respect to the sampling space prior to rotation:

$$1.1 \quad p(\mathbf{X}) = \frac{1}{4}\mathbf{I}\{x_1 \in [-8, 0)\} + \frac{3}{4}\mathbf{I}\{x_1 \in [0, 8]\},$$

$$1.2 \quad p(\mathbf{X}) = \frac{1}{6}\mathbf{I}\{x_1 \in [-8, -\frac{8}{3})\} + \frac{3}{6}\mathbf{I}\{x_1 \in [-\frac{8}{3}, \frac{8}{3})\} + \frac{5}{6}\mathbf{I}\{x_1 \in [\frac{8}{3}, 8]\},$$

$$1.3 \quad p(\mathbf{X}) = \frac{1}{8}\mathbf{I}\{x_1 \in [-8, -4)\} + \frac{3}{8}\mathbf{I}\{x_1 \in [-4, 0)\} + \frac{5}{8}\mathbf{I}\{x_1 \in [0, 4)\} + \frac{7}{8}\mathbf{I}\{x_1 \in [4, 8]\}.$$

Settings 1.1, 1.2, and 1.3 have one, two and three natural boundaries due to the piecewise constant form of $p(\mathbf{X})$. In Figure 6A, we show 1000 observations drawn from simulation

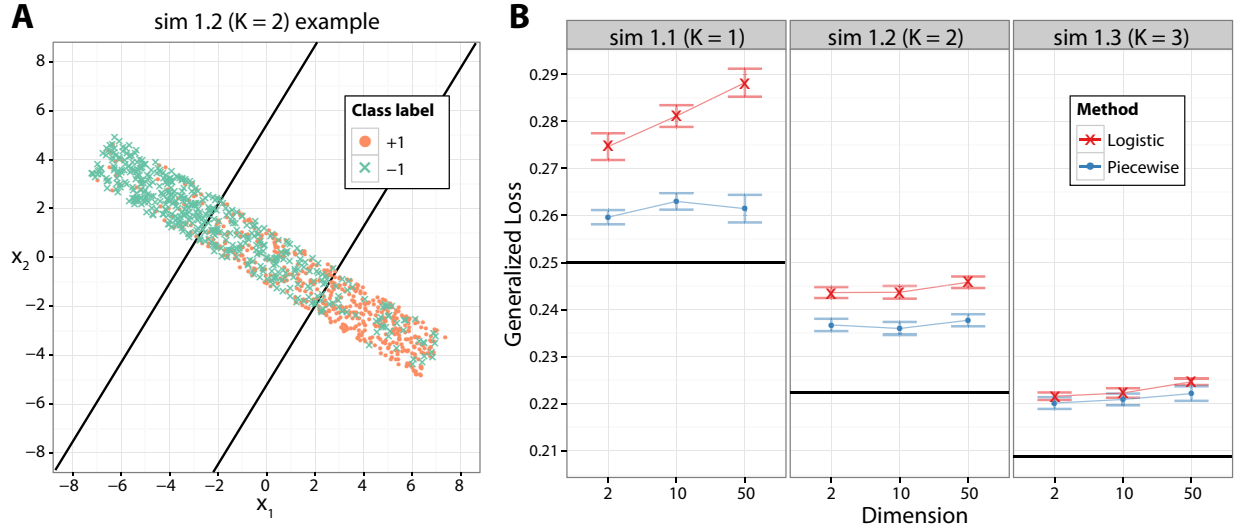


Figure 6: (A) Sample dataset of 1000 observations drawn from the generating distribution for Simulation 1.2. The two Bayes optimal boundaries separating the three regions of constant $p(\mathbf{X})$ are shown with black lines. (B) Comparison of the performance of the piecewise linear and logistic classifiers for the three settings of Simulation 1 across varying dimension. In each panel, the median loss and standard error over 100 replications is shown along with the Bayes minimal loss in black.

setting 1.2, with observations from the positive and negative class shown in orange and green. The Bayes optimal boundaries are also shown in black. For settings 1.1, 1.2, and 1.3, we use the piecewise linear losses with boundaries at π_1 , π_2 , and π_3 , respectively. In each setting, the performance of the piecewise linear and logistic classifiers is evaluated using the theoretical loss for boundaries at π_1 , π_2 , and π_3 . In these simulations, we aim to illustrate the advantage of minimizing and tuning with respect to an appropriate theoretical loss, which matches the underlying form of the data.

The results are shown in Figure 6B, along with the Bayes minimal loss, which provides a lower bound on the theoretical loss in each setting. In all settings, the piecewise linear classifier outperforms the logistic classifier, with the improvement decreasing as the number of boundaries, K increases. This makes intuitive sense, as the piecewise linear loss converges to the logistic loss as $K \rightarrow \infty$. The most significant improvement is seen in setting 1.1, in which the piecewise linear classifier and theoretical loss correspond to the standard SVM and misclassification error. These results confirm previous results highlighting the advantage of hard classifiers over soft classifiers when the underlying $p(\mathbf{X})$ is piecewise constant [7].

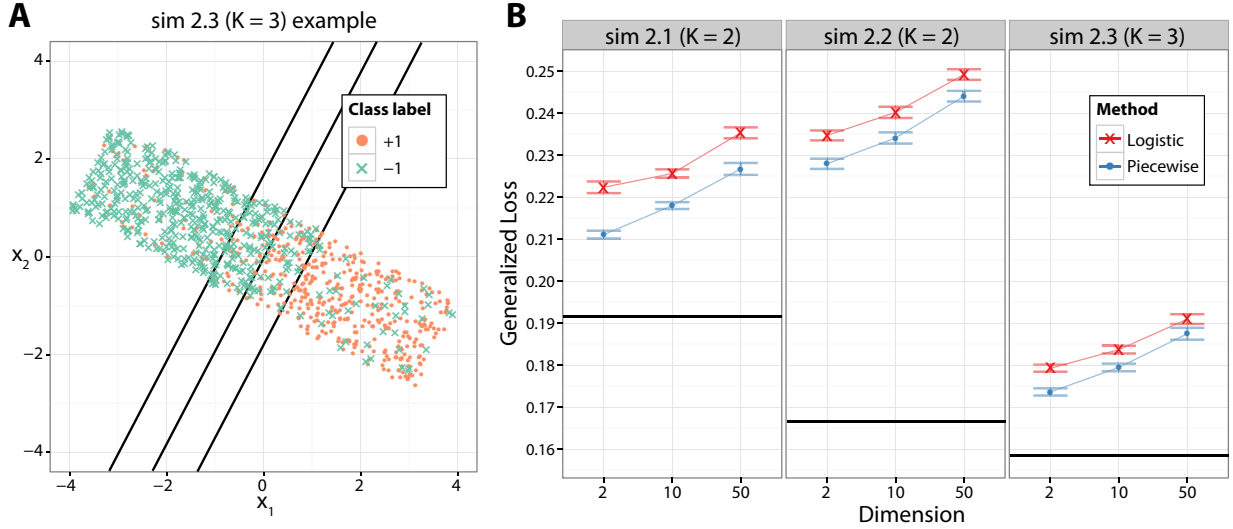


Figure 7: (A) Sample dataset of 1000 observations drawn from the generating distribution for Simulation 2.3. The three Bayes optimal boundaries separating the four regions of constant $p(\mathbf{X})$ are shown with black lines. (B) Comparison of the performance of the piecewise linear and logistic classifiers for the three settings of Simulation 2 across varying dimension. In each panel, the median loss and standard error over 100 replications is shown along with the Bayes minimal loss in black.

Furthermore, the complete set of results illustrates the transition of this behavior as the number of boundaries increases.

6.2 Simulation 2

In Simulation 1, the piecewise constant regions of $p(\mathbf{X})$ were of equal size. In our second set of simulations, we consider unequally spaced conditional class probabilities. Observations were uniformly sampled over $[-4, 4] \times [-1, 1]^{p-1}$, for $p = 2, 10, 50$, again subject to a random rotation. The following conditional class probabilities were considered, again, with respect to the sampling space prior to rotation:

$$2.1 \quad p(\mathbf{X}) = \frac{1}{6}\mathbf{I}\{x_1 \in [-4, -0.6]\} + \frac{3}{6}\mathbf{I}\{x_1 \in [-0.6, 0.6]\} + \frac{5}{6}\mathbf{I}\{x_1 \in [0.6, 4]\},$$

$$2.2 \quad p(\mathbf{X}) = \frac{1}{6}\mathbf{I}\{x_1 \in [-4, -2]\} + \frac{3}{6}\mathbf{I}\{x_1 \in [-2, 0]\} + \frac{5}{6}\mathbf{I}\{x_1 \in [0, 4]\},$$

$$2.3 \quad p(\mathbf{X}) = \frac{1}{8}\mathbf{I}\{x_1 \in [-4, -0.8]\} + \frac{3}{8}\mathbf{I}\{x_1 \in [-0.8, 0]\} + \frac{5}{8}\mathbf{I}\{x_1 \in [0, 0.8]\} + \frac{7}{8}\mathbf{I}\{x_1 \in [0.8, 4]\}.$$

In settings 2.1 and 2.3, we consider $p(\mathbf{X})$ with heavy tails, and in setting 2.2, we consider the case with asymmetric $p(\mathbf{X})$. A sample of 1000 observations drawn from setting 2.3

is shown in Figure 7A, with the Bayes optimal boundaries in black. For settings 2.1, 2.2, and 2.3, we use the piecewise linear losses with boundaries at π_2 , π_2 , and π_3 , respectively. The performance of the piecewise linear and logistic classifiers is again evaluated using the corresponding theoretical loss function. Simulation results are shown in Figure 7B. As in Simulation 1, the piecewise linear classifier outperforms the logistic classifier in all cases. Again, the improvement is greater in settings 2.1 and 2.2 than in setting 2.3, as the piecewise linear loss converges to the logistic loss with increasing K .

7 ADNI Data Analysis

In this section, we apply the proposed interval estimation procedure to a MRI dataset of healthy normal control (NC) and early Alzheimer’s disease (AD) subjects. Data were obtained from the ADNI database (adni.loni.usc.edu). The ADNI was launched in 2003 by the National Institute on Aging (NIA), the National Institute of Biomedical Imaging and Bioengineering (NIBIB), the Food and Drug Administration (FDA), private pharmaceutical companies and non-profit organizations as a \$60 million, 5-year public-private partnership. The Principal Investigator of this initiative is Michael W. Weiner, MD, VA Medical Center and University of California - San Francisco. ADNI is the result of efforts of many co-investigators from a broad range of academic institutions and private corporations, and subjects have been recruited from over 50 sites across the U.S. and Canada. For up-to-date information, see www.adni-info.org.

The dataset we use consists of 93 MRI features measured for 225 NC and 186 AD subjects, and was processed as described in [29]. As in Section 6, the logistic-derived piecewise linear loss is used to target the conditional class probability of AD at $\pi = \{1/4, 2/4, 3/4\}$. Two-fold cross validation is used to determine the optimal λ over $\{2^{-15}, 2^{-14}, \dots, 2^5\}$. The first two principal components (PCs) of the 411 NC and AD subjects are shown in Figure 8A, along with the estimated interval for each subject. Interestingly, the four distinct probability groups appear to separate along the first PC direction.

In addition to NC and AD subjects, the dataset also includes subjects with mild cognitive impairment (MCI), further classified as either progressive MCI (pMCI, 167 subjects) or stable MCI (sMCI, 226 subjects), depending on whether or not the subject progressed to develop AD during the study. The sMCI and pMCI may be considered as intermediary states between the NC and AD subjects. As such, in Figure 8B, we show the distribution of margin values, $\hat{f}(\mathbf{x})$, for NC, sMCI, pMCI, and AD subjects to investigate the transition between

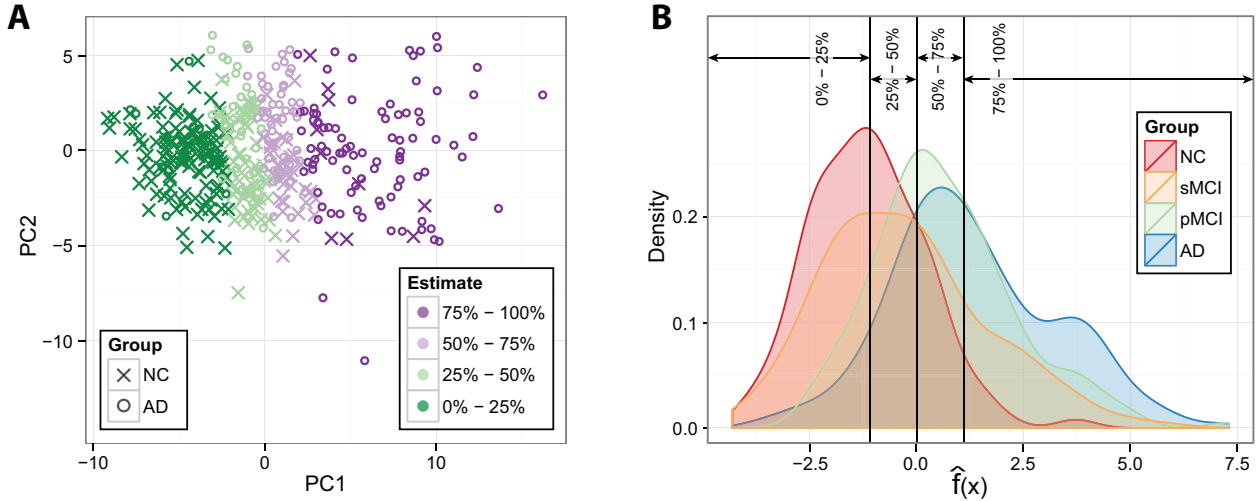


Figure 8: Analysis of ADNI MRI dataset with $\pi = \{0.25, 0.5, 0.75\}$. (A) Scatterplot of first two PCs for AD and NC subjects colored by estimated interval. (B) Density plots of predicted $\hat{f}(x)$ for AD, NC, and two intermediary subject groups, sMCI and pMCI. Corresponding interval cutoffs are shown with vertical lines.

the four distinct groups. The corresponding interval boundaries are shown by vertical lines. Interestingly, while not well-differentiated, the four groups appear to peak within each of the four intervals, with the densities shifting in the expected order. Overall, our method appears to appropriately divide the subject according to the severity of the disease.

8 Discussion

Supervised learning tasks with a discrete class label are commonly encountered in practice. Several problems have been formally defined and studied within this context, including hard, soft, and rejection-option classification. In this paper, we introduce a unified framework of binary learning tasks targeting partial or complete estimation of the conditional class probability, $p(\mathbf{X})$, which encompassing these problems. In contrast to previous frameworks connecting hard and soft classification, our approach spans a space of learning problems, rather than specific loss functions or classification methods. Our approach thus provides a unique perspective to study the transition between hard and soft classification.

We formalize our family of binary learning problems through a unified theoretical loss (7), a corresponding margin based relaxation (9), and a proposed class of minimally consistent piecewise linear surrogates. Simulation studies using the class of piecewise linear

loss functions reinforce previous results on hard and soft classification, and illustrate the transitional behavior between the class of problems. Finally, an application of our interval estimation approach to a MRI dataset from the ADNI study further illustrates the utility of our proposed class of problems.

Acknowledgments:

The authors are supported in part by National Institutes of Health (NIH) Grants U24 CA143848 (Hayes) and U24 CA143848-02S1 (Kimes). Data collection and sharing for this project was funded by the Alzheimer’s Disease Neuroimaging Initiative (ADNI) (NIH Grant U01 AG024904). ADNI is funded by the National Institute on Aging, the National Institute of Biomedical Imaging and Bioengineering, and through generous contributions from the following: Abbott; Alzheimer’s Association; Alzheimer’s Drug Discovery Foundation; Amorfix Life Sciences Ltd.; AstraZeneca; Bayer HealthCare; BioClinica, Inc.; Biogen Idec Inc.; Bristol-Myers Squibb Company; Eisai Inc.; Elan Pharmaceuticals Inc.; Eli Lilly and Company; F. Hoffmann-La Roche Ltd and its affiliated company Genentech, Inc.; GE Healthcare; Innogenetics, N.V.; IXICO Ltd.; Janssen Alzheimer Immunotherapy Research & Development, LLC.; Johnson & Johnson Pharmaceutical Research & Development LLC.; Medpace, Inc.; Merck & Co., Inc.; Meso Scale Diagnostics, LLC.; Novartis Pharmaceuticals Corporation; Pfizer Inc.; Servier; Synarc Inc.; and Takeda Pharmaceutical Company. The Canadian Institutes of Health Research is providing funds to support ADNI clinical sites in Canada. Private sector contributions are facilitated by the Foundation for the National Institutes of Health (www.fnih.org). The grantee organization is the Northern California Institute for Research and Education, and the study is coordinated by the Alzheimer’s Disease Cooperative Study at the University of California, San Diego. ADNI data are disseminated by the Laboratory for Neuro Imaging at the University of Southern California. This research was also supported by NIH grants P30 AG010129 and K01 AG030514.

References

- [1] V. Vapnik, *The nature of statistical learning theory*. Springer, 1995.
- [2] V. Vapnik, *Statistical learning theory*. Wiley, 1998.

- [3] X. Shen, G. C. Tseng, X. Zhang, and W. H. Wong, “On ψ -Learning,” *Journal of the American Statistical Association*, vol. 98, no. 463, pp. 724–734, 2003.
- [4] Y. Liu and X. Shen, “Multicategory ψ -Learning,” *Journal of the American Statistical Association*, vol. 101, no. 474, pp. 500–509, 2006.
- [5] J. C. Platt, “Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods,” *Advances in large margin classifiers*, vol. 10, no. 3, pp. 61–74, 1999.
- [6] J. Wang, X. Shen, and Y. Liu, “Probability estimation for large-margin classifiers,” *Biometrika*, vol. 95, no. 1, pp. 149–167, 2008.
- [7] Y. Liu, H. H. Zhang, and Y. Wu, “Hard or Soft Classification? Large-margin Unified Machines,” *Journal of the American Statistical Association*, vol. 106, no. 493, pp. 166–177, 2011.
- [8] J. S. Marron, M. J. Todd, and J. Ahn, “Distance-Weighted Discrimination,” *Journal of the American Statistical Association*, vol. 102, no. 480, pp. 1267–1271, 2007.
- [9] C. Zhang and Y. Liu, “Multicategory Large-Margin Unified Machines,” *Journal of Machine Learning Research*, vol. 14, pp. 1349–1386, 2013.
- [10] R. Herbei and M. H. Wegkamp, “Classification with reject option,” *Canadian Journal of Statistics*, vol. 34, no. 4, pp. 709–721, 2006.
- [11] P. L. Bartlett and M. H. Wegkamp, “Classification with a reject option using a hinge loss,” *Journal of Machine Learning Research*, vol. 9, pp. 1823–1840, 2008.
- [12] M. Yuan and M. H. Wegkamp, “Classification methods with reject option based on convex risk minimization,” *Journal of Machine Learning Research*, vol. 11, pp. 111–130, 2010.
- [13] M. H. Wegkamp and M. Yuan, “Support vector machines with a reject option,” *Bernoulli*, vol. 17, no. 4, pp. 1368–1385, 2011.
- [14] Y. Lin, Y. Lee, and G. Wahba, “Support Vector Machines for Classification in Nonstandard Situations,” *Machine Learning*, vol. 46, pp. 191–202, 2002.
- [15] X. Qiao and Y. Liu, “Adaptive weighted learning for unbalanced multicategory classification,” *Biometrics*, vol. 65, no. 1, pp. 159–168, 2009.
- [16] I. Steinwart and C. Scovel, “Fast rates for support vector machines using Gaussian kernels,” *The Annals of Statistics*, vol. 35, no. 2, pp. 575–607, 2007.
- [17] G. Blanchard, O. Bousquet, and P. Massart, “Statistical performance of support vector machines,” *The Annals of Statistics*, vol. 36, no. 2, pp. 489–531, 2008.
- [18] P. L. Bartlett, M. I. Jordan, and J. D. McAuliffe, “Convexity, Classification, and Risk

- Bounds,” *Journal of the American Statistical Association*, vol. 101, no. 473, pp. 138–156, 2006.
- [19] N. Cristianini and J. Shawe-Taylor, *An Introduction to Support Vector Machines and Other Kernel-based Learning Methods*. Cambridge University Press, 1 ed., 2000.
- [20] Y. Lin, “Support Vector Machines and the Bayes Rule in Classification,” *Data Mining and Knowledge Discovery*, vol. 6, pp. 259–275, 2002.
- [21] C. Zhang, Y. Liu, and Z. Wu, “On the effect and remedies of shrinkage on classification probability estimation,” *The American Statistician*, vol. 67, no. 3, pp. 134–142, 2013.
- [22] Y. Grandvalet, A. Rakotomamonjy, J. Keshet, and S. Canu, “Support vector machines with a reject option,” in *Advances in Neural Information Processing Systems 21*, pp. 537–544, Curran Associates, Inc., 2009.
- [23] E. Mammen and A. B. Tsybakov, “Smooth Discrimination Analysis,” *The Annals of Statistics*, vol. 27, no. 6, pp. 1808–1829, 1999.
- [24] L. Bottou and C.-J. Lin, “Support Vector Machine Solvers,” in *Large Scale Kernel Machines*, pp. 301–320, Cambridge, MA.: MIT Press, 2007.
- [25] S. Shalev-Shwartz, Y. Singer, N. Srebro, and A. Cotter, “Pegasos: primal estimated sub-gradient solver for SVM,” *Mathematical Programming*, vol. 127, no. 1, pp. 3–30, 2010.
- [26] N. Aronszajn, “Theory of Reproducing Kernels,” *Transactions of the American Mathematical Society*, vol. 68, no. 3, pp. 337–404, 1950.
- [27] G. Wahba, “Support vector machines, reproducing kernel Hilbert spaces and the randomized GACV,” *Advances in Kernel Methods-Support Vector Learning*, vol. 6, pp. 69–87, 1999.
- [28] P. H. Calamai and J. J. Moré, “Projected gradient methods for linearly constrained problems,” *Mathematical Programming*, vol. 39, no. 1, pp. 93–116, 1987.
- [29] G. Yu, Y. Liu, K.-H. Thung, and D. Shen, “Multi-Task Linear Programming Discriminant Analysis for the Identification of Progressive MCI Individuals,” *PLoS ONE*, vol. 9, no. 5, p. e96458, 2014.

SUPPLEMENTARY MATERIALS

S1 Common learning problems as special cases

In this section, we show that our class of problems encompasses hard, weighted, rejection-option, and soft classification. For hard, weighted and rejection-option classification, the equivalence is derived by showing that for specific choices of $\boldsymbol{\pi}$, the theoretical loss (7) reduces to the standard losses given in (2), (3), and (4). For soft classification, the equivalence is shown by deriving the limiting form of the theoretical loss (7), and showing that the limiting loss is optimized by $p(\mathbf{X})$.

S1.1 Hard and Weighted Classification

For hard classification, let $\boldsymbol{\pi} = \{0.5\}$, such that $\Omega_{\boldsymbol{\pi}} = \{\omega_0, \omega_1\} = \{[0, 0.5], (0.5, 1]\}$. Then, the theoretical loss (7) may be simplified as:

$$\begin{aligned}
 \ell_{\boldsymbol{\pi}}^Y(g(\mathbf{X})) &= 2 \cdot \ell_{0.5}^Y(g(\mathbf{X})) \\
 &= \begin{cases} 2 \cdot (1 - 0.5) \cdot \mathbf{I}\{g(\mathbf{X}) \leq 0.5\} & \text{if } Y = +1 \\ 2 \cdot 0.5 \cdot \mathbf{I}\{g(\mathbf{X}) > 0.5\} & \text{if } Y = -1 \end{cases} \\
 &= \mathbf{I}\{(g(\mathbf{X}) \leq 0.5, Y = +1) \text{ or } (g(\mathbf{X}) > 0.5, Y = -1)\} \\
 &= \mathbf{I}\{(g(\mathbf{X}) = \omega_0, Y = +1) \text{ or } (g(\mathbf{X}) = \omega_1, Y = -1)\}. \tag{S1}
 \end{aligned}$$

The equivalence of (S1) to the 0–1 loss (2) follows by noting that for the Bayes optimal rule for hard classification (6), predictions of $\widehat{Y} = +1, -1$ correspond to $p(\mathbf{x}) \in \omega_0$ and $p(\mathbf{x}) \in \omega_1$, respectively. More generally, the weighted 0–1 loss (4) may be similarly recovered up to a multiplicative constant by letting $\boldsymbol{\pi} = \{\pi\}$ for any $\pi \in (0, 1)$ such that $\Omega_{\boldsymbol{\pi}} = \{\omega_0, \omega_1\} = \{[0, \pi], (\pi, 1]\}$, and

$$\begin{aligned}
 \ell_{\boldsymbol{\pi}}^Y(g(\mathbf{X})) &= 2 \cdot \ell_{\pi}^Y(g(\mathbf{X})) \\
 &= \begin{cases} 2 \cdot (1 - \pi) \cdot \mathbf{I}\{g(\mathbf{X}) \leq \pi\} & \text{if } Y = +1 \\ 2 \cdot \pi \cdot \mathbf{I}\{g(\mathbf{X}) > \pi\} & \text{if } Y = -1 \end{cases} \\
 &\propto (1 - \pi) \cdot \mathbf{I}\{g(\mathbf{X}) = \omega_0, Y = +1\} \\
 &\quad + \pi \cdot \mathbf{I}\{g(\mathbf{X}) = \omega_1, Y = -1\}. \tag{S2}
 \end{aligned}$$

Again, the equivalence of (S2) to the weighted 0–1 loss (4) follows from the form of the Bayes optimal rule for weighted classification (6).

S1.2 Rejection-Option Classification

For rejection-option classification as formulated in [12], let $\boldsymbol{\pi} = \{\pi, 1 - \pi\}$ for some $\pi \in (0, 0.5)$, such that $\Omega_{\boldsymbol{\pi}} = \{\omega_0, \omega_1, \omega_2\} = \{[0, \pi], (\pi, 1 - \pi], (1 - \pi, 1]\}$. We first rewrite (7) as:

$$\begin{aligned}
\ell_{\boldsymbol{\pi}}^Y(g(\mathbf{X})) &= \frac{2}{2} \sum_{k=1}^2 \ell_{\pi_k}^Y(g(\mathbf{X})) \\
&= \begin{cases} \sum_{k=1}^2 (1 - \pi_k) \cdot \mathbf{I}\{g(\mathbf{X}) \leq \pi_k\} & \text{if } Y = +1 \\ \sum_{k=1}^2 \pi_k \cdot \mathbf{I}\{g(\mathbf{X}) > \pi_k\} & \text{if } Y = -1 \end{cases} \\
&= \begin{cases} (1 - \pi) \cdot \mathbf{I}\{g(\mathbf{X}) \leq \pi\} + \pi \cdot \mathbf{I}\{g(\mathbf{X}) \leq 1 - \pi\} & \text{if } Y = +1 \\ \pi \cdot \mathbf{I}\{g(\mathbf{X}) > \pi\} + (1 - \pi) \cdot \mathbf{I}\{g(\mathbf{X}) > 1 - \pi\} & \text{if } Y = -1 \end{cases} \\
&= \begin{cases} (1 - \pi) \cdot \mathbf{I}\{g(\mathbf{X}) = \omega_0\} + \pi \cdot \mathbf{I}\{g(\mathbf{X}) \neq \omega_2\} & \text{if } Y = +1 \\ \pi \cdot \mathbf{I}\{g(\mathbf{X}) \neq \omega_0\} + (1 - \pi) \cdot \mathbf{I}\{g(\mathbf{X}) = \omega_2\} & \text{if } Y = -1 \end{cases} \\
&= \begin{cases} 1 & \text{if } (g(\mathbf{X}) = \omega_0, Y = +1) \text{ or } (g(\mathbf{X}) = \omega_2, Y = -1) \\ \pi & \text{if } g(\mathbf{X}) = \omega_1 \\ 0 & \text{otherwise} \end{cases} .
\end{aligned}$$

The equivalence of (7) to the rejection option loss (3) is established by noting the correspondence between predictions of $\{\omega_0, \omega_1, \omega_2\}$ and $\widehat{Y}_{rej} = \{-1, 0, +1\}$ for the Bayes optimal rejection-option rule (5).

S1.3 Soft Classification

Although not traditionally formulated as the minimization of a theoretical loss, the soft classification problem may be derived as the special case of (7) when $K = \infty$ and $\boldsymbol{\pi}$ becomes dense on $(0, 1)$, such that $\Omega_{\boldsymbol{\pi}} = (0, 1)$. The limiting form of (7), which we define as the average of K functions, may be expressed as the following integral:

$$\ell_{\boldsymbol{\pi}}^Y(g(\mathbf{X})) = \lim_{K \rightarrow \infty} \frac{2}{K} \sum_{k=1}^K \ell_{\pi_k}^Y(g(\mathbf{X}))$$

$$\begin{aligned}
&= \begin{cases} 2 \int_0^1 (1 - \pi) \cdot \mathbf{I}\{g(\mathbf{X}) \leq \pi\} d\pi & \text{if } Y = +1 \\ 2 \int_0^1 \pi \cdot \mathbf{I}\{g(\mathbf{X}) > \pi\} d\pi & \text{if } Y = -1 \end{cases} \\
&= \begin{cases} (1 - g(\mathbf{X}))^2 & \text{if } Y = +1 \\ g(\mathbf{X})^2 & \text{if } Y = -1 \end{cases} \\
&= (\mathbf{I}\{Y = +1\} - g(\mathbf{X}))^2.
\end{aligned}$$

Thus, the limiting loss is minimized by the prediction $g(\mathbf{X}) = \mathbb{E}_{Y|\mathbf{X}}\{\mathbf{I}\{Y = +1\}\} = p(\mathbf{X})$, corresponding to the conditional class probability estimation task of soft classification.

S2 Proof of Theorem 1

Let $\boldsymbol{\pi} = \{\pi_1, \dots, \pi_K\}$ for some $K \geq 1$ such that $0 < \pi_1 < \dots < \pi_K < 1$. Furthermore, let $h \in \{0, \dots, K\}$ denote the index for some predicted $\omega_h \in \Omega_{\boldsymbol{\pi}}$. Then,

$$\begin{aligned}
\mathbb{E}_{Y|\mathbf{X}}\{\ell_{\boldsymbol{\pi}}^Y(\omega_h)\} &= p(\mathbf{X}) \cdot \ell_{\boldsymbol{\pi}}^+(\omega_h) + (1 - p(\mathbf{X})) \cdot \ell_{\boldsymbol{\pi}}^-(\omega_h). \\
&\propto p(\mathbf{X}) \sum_{k=h+1}^{K+1} (1 - \pi_k) + (1 - p(\mathbf{X})) \sum_{k=1}^h \pi_k,
\end{aligned}$$

Letting $\pi_0 = 0$, $\pi_{K+1} = 1$, we can express the above as:

$$\mathbb{E}_{Y|\mathbf{X}}\{\ell_{\boldsymbol{\pi}}^Y(\omega_h)\} = \sum_{k=0}^{K+1} \left\{ p(\mathbf{X})(1 - \pi_k) \cdot \mathbf{I}_{\{k > h\}} + \pi_k(1 - p(\mathbf{X})) \cdot \mathbf{I}_{\{k \leq h\}} \right\}.$$

The sum is minimized by choosing h such that $p(\mathbf{X})(1 - \pi_k) \geq \pi_k(1 - p(\mathbf{X}))$ for all $k \leq h$ and $p(\mathbf{X})(1 - \pi_k) \leq \pi_k(1 - p(\mathbf{X}))$ for all $k > h$. Thus, the optimal solution is given by $h^* = \operatorname{argmax}_k \{\pi_k < p(\mathbf{X})\}$. The equivalence between ω_{h^*} and $\sum_{k=0}^K \omega_k \cdot \mathbf{I}\{p(\mathbf{X}) \in \omega_k\}$ is immediate from the fact that $p(\mathbf{X}) \in (\pi_{h^*}, \pi_{h^*+1}] = \omega_{h^*}$, and the additional assumption that $p(\mathbf{X}) \neq \pi_k$ a.s. for all k .

S3 Proof of Theorem 2

Let $\boldsymbol{\pi}$ and $\boldsymbol{\delta}$ be appropriately defined boundaries in $(0, 1)$ and \mathbb{R} . Note that surrogate losses, ϕ^+ , ϕ^- are consistent for boundaries at $\boldsymbol{\pi}$ with $\boldsymbol{\delta}$, i.e. $\boldsymbol{\pi}, \boldsymbol{\delta}$ -consistent, if and only if they are π_k, δ_k -consistent for each k separately. Thus, conditions for $\boldsymbol{\pi}, \boldsymbol{\delta}$ -consistency are simply the

union of the conditions for π_k, δ_k -consistency. Necessary and sufficient conditions for ϕ^Y to be π_k, δ_k -consistent were provided by Theorem 1 of [12].

S4 Proof of Theorem 3

Let $\boldsymbol{\pi}$ be an appropriately defined set of boundaries in $(0, 1)$. Assume φ^Y to be defined as in (11) and (12) such that (C1)–(C3) are satisfied. We wish to show that for all $\pi_k \in \boldsymbol{\pi}$, there exists some δ_k such that (10) is satisfied, and furthermore, that there does not exist any δ such that (10) is satisfied for $\pi \in (0, 1) \setminus \boldsymbol{\pi}$. Equivalently, we wish to show that $\varphi^{-'}(x)/(\varphi^{-'}(x) + \varphi^{+'}(x))$ only takes values in $\boldsymbol{\pi}$ over the set of x such that $\varphi^{-'}(x) < 0$ and $\varphi^{+'}(x) < 0$ are defined. Note that $\varphi^{+'}$ and φ^{-}' are only undefined at the hinge points, $H^Y(\pi_k, \pi_{k+1})$, $A^-(\pi_1)/B^-(\pi_1)$, and $A^+(\pi_K)/B^+(\pi_K)$. By (C2), the set of possible $\varphi^{+'}, \varphi^{-}'$ pairs are given by:

$$\begin{array}{cccccc} \hline \varphi^{+'}: & B^+(\pi_1) & B^+(\pi_1) & \cdots & B^+(\pi_K) & 0 \\ \hline \varphi^{-}': & 0 & B^-(\pi_1) & \cdots & B^-(\pi_K) & B^-(\pi_K) \\ \hline \end{array}$$

Excluding the cases when $\varphi^{+'}(x) = 0$ or $\varphi^{-}'(x) = 0$, the set of possible consistent boundaries satisfying (10) are given by:

$$\frac{\varphi^{-'}(x)}{\varphi^{-'}(x) + \varphi^{+'}(x)} = \frac{B^{-'}(\pi_k)}{B^{-'}(\pi_k) + B^{+'}(\pi_k)} = \pi_k \quad \text{for } k = 1, \dots, K,$$

where the final equality is given by (C3).

S5 Proof of Proposition 1

Let $\boldsymbol{\pi}$ be an appropriately defined set of boundaries in $(0, 1)$. We wish to show that (C1)–(C3) of Theorem 3 are satisfied for $A^+(\pi) = A^-(1 - \pi) = -\pi \log(\pi) - (1 - \pi) \log(1 - \pi)$, and $B^+(\pi) = B^-(1 - \pi) = -(1 - \pi)$.

Trivially, (C1) is satisfied, as $B^+(\pi) = \pi - 1$ and $B^-(\pi) = -\pi$ are non-decreasing and non-increasing, respectively, in π . To show that (C2) is satisfied, we derive the hinge points for the positive and negative class losses:

$$\begin{aligned} H^+(\pi, \pi') &= \frac{A^+(\pi) - A^+(\pi')}{B^+(\pi') - B^+(\pi)} = \frac{A^+(\pi) - A^+(\pi')}{\pi' - \pi} \\ H^-(\pi, \pi') &= \frac{A^-(\pi) - A^-(\pi')}{B^-(\pi') - B^-(\pi)} = -\frac{A^+(\pi) - A^+(\pi')}{\pi' - \pi}, \end{aligned}$$

where the final equality is obtained by noting $A^+(\pi) = A^+(1 - \pi)$. The first equality of (C2) is clearly satisfied by the above derivations. We next show that the remaining three inequalities of (C2) are also satisfied. Let $k \in \{2, \dots, K - 1\}$. By the concavity of $A^+(\pi)$:

$$\begin{aligned} H^+(\pi_{k-1}, \pi_k) &= \frac{A^+(\pi_{k-1}) - A^+(\pi_k)}{\pi_k - \pi_{k-1}} \\ &= -\frac{A^+(\pi_k) - A^+(\pi_{k-1})}{\pi_k - \pi_{k-1}} \\ &< -(A^+)'(\pi_k) \\ &< -\frac{A^+(\pi_{k+1}) - A^+(\pi_k)}{\pi_{k+1} - \pi_k} = H^+(\pi_k, \pi_{k+1}), \end{aligned}$$

Similarly, by the convexity of $A^-(\pi)$ and the fact that $\lim_{\pi \rightarrow 0} A^+(\pi) = \lim_{\pi \rightarrow 1} A^+(\pi) = 0$, we have:

$$\begin{aligned} \frac{A^-(\pi_1)}{B^-(\pi_1)} &= -\frac{A^-(\pi_1) - \lim_{\pi \rightarrow 0} A^-(\pi)}{\pi_1 - 0} \\ &< -(A^-)'(\pi_1) \\ &< -\frac{A^-(\pi_1) - A^-(\pi_2)}{\pi_1 - \pi_2} = H^+(\pi_1, \pi_2) \\ \frac{A^+(\pi_K)}{B^+(\pi_K)} &= -\frac{A^+(\pi_K) - \lim_{\pi \rightarrow 1} A^+(\pi)}{\pi_K - 1} \\ &> -(A^+)'(\pi_K) \\ &> -\frac{A^+(\pi_{K-1}) - A^+(\pi_K)}{\pi_{K-1} - \pi_K} = H^-(\pi_{K-1}, \pi_K). \end{aligned}$$

Thus (C2) is satisfied. Finally, (C3) holds, since for any $k = 1, \dots, K$:

$$\frac{B^-(\pi_k)}{B^-(\pi_k) + B^+(\pi_k)} = \frac{-\pi_k}{-\pi_k - (1 - \pi_k)} = \pi_k.$$

S6 Proof of Theorem 4

Let ϕ^Y be a consistent surrogate loss for appropriately defined boundaries $\boldsymbol{\pi}$ in $(0, 1)$ at $\boldsymbol{\delta}$. First, note that the excess condition ϕ -risk for a rule $g \in \mathcal{G}$ may be written as:

$$R_p(g(\mathbf{x})) = \frac{2}{K} \left[(1 - p(\mathbf{x})) \sum_k \pi_k \mathbf{I}\{g(\mathbf{x}) > \pi_k\} + p(\mathbf{x}) \sum_k (1 - \pi_k) \mathbf{I}\{g(\mathbf{x}) \leq \pi_k\} \right].$$

Consider a candidate rule $g \in \mathcal{G}$, and recall the Bayes optimal rule over \mathcal{G} , $W_{\pi}^*(\mathbf{x})$, defined in Theorem 1. Suppose that $\mathbf{x} \in \mathcal{X}$ is such that $g(\mathbf{x}) > W_{\pi}^*(\mathbf{x})$. Then, letting

$$\mathcal{K} = \begin{cases} \{k : g(\mathbf{X}) \leq \pi_k < W_{\pi}^*(\mathbf{X})\} & \text{if } W_{\pi}^*(\mathbf{X}) > g(\mathbf{X}) \\ \{k : W_{\pi}^*(\mathbf{X}) \leq \pi_k < g(\mathbf{X})\} & \text{if } W_{\pi}^*(\mathbf{X}) < g(\mathbf{X}), \\ \emptyset & \text{otherwise} \end{cases}$$

the excess condition ϕ -risk may be expressed as:

$$\begin{aligned} \Delta R_p(g) &= \frac{2}{K} \left[(1 - p(\mathbf{x})) \sum_k \pi_k \mathbf{I}\{\pi_k : W_{\pi}^*(\mathbf{x}) \leq \pi_k < g(\mathbf{x})\} \right. \\ &\quad \left. - p(\mathbf{x}) \sum_k (1 - \pi_k) \mathbf{I}\{\pi_k : W_{\pi}^*(\mathbf{x}) \leq \pi_k < g(\mathbf{x})\} \right] \\ &= \frac{2}{K} \sum_{\mathcal{K}} [(1 - p(\mathbf{x}))\pi_k - p(\mathbf{x})(1 - \pi_k)] \\ &= \frac{2}{K} \sum_{\mathcal{K}} [\pi_k - p(\mathbf{x})]. \end{aligned}$$

Similarly, for $g(\mathbf{x}) < W_{\pi}^*(\mathbf{x})$, $\Delta R_p(g) = \frac{2}{K} \sum_{\mathcal{K}} [p(\mathbf{x}) - \pi_k]$. If $g(\mathbf{x}) = W_{\pi}^*(\mathbf{x})$, we have that $\Delta R_p(g) = 0$, such that:

$$\Delta R_p(g) = \frac{2}{K} \sum_{\mathcal{K}} |p(\mathbf{x}) - \pi_k|,$$

for all $\mathbf{x} \in \mathcal{X}$.

By the stated assumptions, for $g(\mathbf{x}) = C(f(\mathbf{x}); \boldsymbol{\delta}) \in \mathcal{G}$, we immediately have the following result:

$$\begin{aligned} (\Delta R_p(g))^s &= \left(\frac{2}{K} \sum_{\mathcal{K}} |p(\mathbf{X}) - \pi_k| \right)^s \\ &\leq \frac{2}{K} \sum_{\mathcal{K}} |p(\mathbf{X}) - \pi_k|^s \\ &\leq \frac{2}{K} C^s \sum_{\mathcal{K}} \Delta Q_p(\delta_k) \\ \Delta R_p(g) &\leq C \left(\frac{2}{K} \sum_{\mathcal{K}} \Delta Q_p(\delta_k) \right)^{1/s}. \end{aligned}$$

Since $\Delta Q_p \geq 0$, it suffices to show that $\sum_{\mathcal{K}} \Delta Q_p(\delta_k) \leq K \cdot \Delta Q_p(f)$. Since $|\mathcal{K}| \leq K$, we

complete the proof by showing $\Delta Q_p(f) \geq \Delta Q_p(\delta_k)$ for all $k \in \mathcal{K}$. Without loss of generality, suppose \mathbf{x} is such that $g(\mathbf{x}) < W_{\pi}^*(\mathbf{x})$ and let $k \in \mathcal{K}$. Note that $\pi_k < g(\mathbf{x})$ is equivalent to $\delta_k < f(\mathbf{x})$. By this fact and the convexity and consistency of ϕ^Y , the following inequalities hold:

$$\frac{\phi^+(f(\mathbf{x})) - \phi^+(\delta_k)}{f(\mathbf{x}) - \delta_k} \geq \phi^{+'}(\delta_k) \quad \frac{\phi^-(-f(\mathbf{x})) - \phi^-(-\delta_k)}{-f(\mathbf{x}) + \delta_k} \leq \phi^{-'}(-\delta_k).$$

Thus,

$$\begin{aligned} Q_p(f) - Q_p(\delta_k) &= p(\mathbf{x})(\phi^+(f(\mathbf{x})) - \phi^+(\delta_k)) + (1 - p(\mathbf{x}))(\phi^-(-f(\mathbf{x})) - \phi^-(-\delta_k)) \\ &\geq p(\mathbf{x})(f(\mathbf{x}) - \delta_k)\phi^{+'}(\delta_k) - (1 - p(\mathbf{x}))(f(\mathbf{x}) - \delta_k)\phi^{-'}(-\delta_k) \\ &\geq (f(\mathbf{x}) - \delta_k)\{p(\mathbf{x})(\phi^{+'}(\delta_k) + \phi^{-'}(-\delta_k)) - \phi^{-'}(-\delta_k)\} \\ &\geq (f(\mathbf{x}) - \delta_k)\{p(\mathbf{x})\frac{\phi^{-'}(-\delta_k)}{\pi_k} - \phi^{-'}(-\delta_k)\} \\ &\geq (f(\mathbf{x}) - \delta_k)\phi^{-'}(-\delta_k)\left(\frac{p(\mathbf{x})}{\pi_k} - 1\right). \end{aligned}$$

Since $f(\mathbf{x}) - \delta_k > 0$, $\phi^{-'}(-\delta_k) < 0$, and $p(\mathbf{x}) < \pi_k$, $Q_p(f) - Q_p(\delta_k) \geq 0$. The case when $g(\mathbf{x}) < W_{\pi}^*(\mathbf{x})$ follows similarly, and the proof is complete.

S7 Proof of Theorem 5

Let ϕ^Y be a consistent surrogate loss for appropriately defined boundaries $\boldsymbol{\pi}$ in $(0, 1)$ at $\boldsymbol{\delta}$. Throughout, we use $g = C(f; \boldsymbol{\delta})$ to denote the corresponding rule in \mathcal{G} for some margin function $f \in \mathcal{F}$. From the proof of Theorem 4, we have that:

$$\begin{aligned} \Delta R(g) &= \frac{2}{K} \cdot \mathbb{E}\left\{\sum_{\mathcal{K}} |p(\mathbf{X}) - \pi_k|\right\} \\ &= \frac{2}{K} \cdot \mathbb{E}\left\{\sum_{k=1}^K |p(\mathbf{X}) - \pi_k| \cdot \mathbf{I}\{k \in \mathcal{K}\}\right\} \\ &= \frac{2}{K} \cdot \sum_{k=1}^K \mathbb{E}\left\{|p(\mathbf{X}) - \pi_k| \cdot \mathbf{I}\{k \in \mathcal{K}\}\right\}, \end{aligned}$$

where \mathcal{K} is defined as in the proof of Theorem 4 (Section S6). Additionally, note that for fixed $k \in \{1, \dots, K\}$:

$$\mathbb{E}\left\{|p(\mathbf{X}) - \pi_k| \cdot \mathbf{I}\{k \in \mathcal{K}\}\right\} \geq t \cdot \mathbb{P}\{(k \in \mathcal{K}) \cap |p(\mathbf{X}) - \pi_k| > t\}$$

$$\begin{aligned}
&= t \cdot \mathbb{P}\{|p(\mathbf{X}) - \pi_k| > t\} - t \cdot \mathbb{P}\{(k \notin \mathcal{K}) \cap |p(\mathbf{X}) - \pi_k| > t\} \\
&\geq t \cdot (1 - At^\alpha) - t \cdot \mathbb{P}\{k \notin \mathcal{K}\} \\
&= t \cdot (\mathbb{P}\{k \in \mathcal{K}\} - At^\alpha).
\end{aligned}$$

Combining the above inequalities, we have:

$$\begin{aligned}
\Delta R(g) &\geq \frac{2t}{K} \cdot \left(\sum_{k=1}^K \mathbb{P}\{k \in \mathcal{K}\} - KAt^\alpha \right) \\
&\geq \frac{2t}{K} \cdot \left(\mathbb{P}\{f \neq f^*\} - KAt^\alpha \right).
\end{aligned}$$

Letting $t = \left(\frac{\mathbb{P}\{f \neq f^*\}}{2KA}\right)^{1/\alpha}$ and using β to denote $\alpha/(1 + \alpha)$,

$$\begin{aligned}
\Delta R(g) &\geq \frac{2}{K} \cdot \left(\frac{\mathbb{P}\{f \neq f^*\}}{2KA} \right)^{1/\alpha} \cdot \left(\frac{\mathbb{P}\{f \neq f^*\}}{2} \right) \\
&= \frac{\mathbb{P}\{f \neq f^*\}^{(1+\alpha)/\alpha}}{(2A)^{1/\alpha} K^{(1+\alpha)/\alpha}} \\
\frac{\mathbb{P}\{f \neq f^*\}}{K} &\leq ((2A)^{1/\alpha} \Delta R(g))^\beta.
\end{aligned}$$

Now consider,

$$\begin{aligned}
\Delta R(g) &= \frac{2}{K} \cdot \sum_{k=1}^K \mathbb{E}(|p(\mathbf{X}) - \pi_k| \cdot \mathbf{I}\{k \in \mathcal{K}\}) \\
&= \frac{2}{K} \cdot \sum_{k=1}^K \mathbb{E}(|p(\mathbf{X}) - \pi_k| \cdot \mathbf{I}\{k \in \mathcal{K}\} \cdot \mathbf{I}\{|p(\mathbf{X}) - \pi_k| > \epsilon\}) \\
&\quad + \frac{2}{K} \cdot \sum_{k=1}^K \mathbb{E}(|p(\mathbf{X}) - \pi_k| \cdot \mathbf{I}\{k \in \mathcal{K}\} \cdot \mathbf{I}\{|p(\mathbf{X}) - \pi_k| \leq \epsilon\}).
\end{aligned}$$

Using the inequality: $|x| \cdot \mathbf{I}\{|x| \geq \epsilon\} \leq |x|^s \cdot \epsilon^{1-s}$ for $s \geq 1$, we have:

$$\begin{aligned}
\Delta R(g) &\leq \frac{2}{K} \cdot \sum_{k=1}^K \mathbb{E}(|p(\mathbf{X}) - \pi_k|^s \cdot \epsilon^{1-s} \cdot \mathbf{I}\{k \in \mathcal{K}\}) \\
&\quad + \frac{2\epsilon}{K} \cdot \sum_{k=1}^K \mathbb{P}\{k \in \mathcal{K}\}.
\end{aligned}$$

From the proof of Theorem 4 (Section S6), $|p(\mathbf{X}) - \pi_k|^s \leq C^s \Delta Q_p(\delta_k) \leq C^s \Delta Q_p(f)$ for $k \in \mathcal{K}$. Therefore,

$$\Delta R(g) \leq 2\epsilon^{1-s} C^s \Delta Q(f) + \frac{2\epsilon}{K} \cdot \mathbb{P}\{f \neq f^*\}.$$

Combining with the previous bound on $\mathbb{P}\{f \neq f^*\}$,

$$\Delta R(g) \leq 2\epsilon^{1-s} C^s \Delta Q(f) + 2\epsilon \cdot ((2A)^{1/\alpha} \Delta R(g))^\beta$$

Further choosing $\epsilon = \Delta R(g)^{1-\beta}$,

$$\begin{aligned} \Delta R(g) &\leq 2\Delta R(g)^{(1-\beta)(1-s)} C^s \Delta Q(f) + 2(2A)^{1/\alpha} \Delta R(g) \\ (1 - 2(2A)^{1/\alpha}) \Delta R(g)^{s+\beta-s\beta} &\leq 2C^s \Delta Q(f) \\ \Delta R(g) &\leq \left(\frac{2C^s}{1 - 2(2A)^{1/\alpha}} \right)^{1/(s+\beta-s\beta)} \cdot \Delta Q(f)^{1/(s+\beta-s\beta)} \end{aligned}$$

Letting D denote the exponentiated fraction on the right of the inequality,

$$\Delta R(g) \leq D \cdot \Delta Q(f)^{1/(s+\beta-s\beta)}.$$

S8 Proof of Corollary 1

Let φ^Y be a minimally consistent piecewise linear surrogate loss for appropriately defined boundaries $\boldsymbol{\pi}$ in $(0, 1)$ at $\boldsymbol{\delta}$. The φ^Y -optimal margin function, denoted by f_φ^* , is given by:

$$\begin{aligned} f_\varphi^*(\mathbf{X}) &= \operatorname{argmin}_f \mathbb{E}_{Y|\mathbf{X}} \{ \varphi^Y(Yf(\mathbf{X})) \} \\ &= \operatorname{argmin}_f \{ p(\mathbf{X})\varphi^+(f(\mathbf{X})) + (1 - p(\mathbf{X}))\varphi^-(-f(\mathbf{X})) \} \\ &= \begin{cases} A^-(\pi_1)/B^-(\pi_1) & \text{if } p(\mathbf{X}) \in [0, \pi_1) \\ H^+(\pi_1, \pi_2) & \text{if } p(\mathbf{X}) \in [\pi_1, \pi_2) \\ \dots & \dots \\ H^+(\pi_{K-1}, \pi_K) & \text{if } p(\mathbf{X}) \in [\pi_{K-1}, \pi_K) \\ A^+(\pi_K)/B^+(\pi_K) & \text{if } p(\mathbf{X}) \in [\pi_K, 1] \end{cases} \end{aligned}$$

For any $k \in \{1, \dots, K\}$,

$$\begin{aligned}
\Delta Q_p(\delta_k) &= Q_p(\delta_k) - Q_p(f_\varphi^*(\mathbf{X})) \\
&= p(\mathbf{X})(\varphi^+(\delta_k) - \varphi^+(f_\varphi^*(\mathbf{X}))) + (1 - p(\mathbf{X}))(\varphi^-(-\delta_k) - \varphi^-(-f_\varphi^*(\mathbf{X}))) \\
&= p(\mathbf{X})B^+(\pi_k)(\delta_k - f_\varphi^*(\mathbf{X})) - (1 - p(\mathbf{X}))B^-(\pi_k)(\delta_k - f_\varphi^*(\mathbf{X})) \\
&= p(\mathbf{X})(B^+(\pi_k) + B^-(\pi_k))(\delta_k - f_\varphi^*(\mathbf{X})) - B^-(\pi_k)(\delta_k - f_\varphi^*(\mathbf{X})) \\
&= p(\mathbf{X})(B^-(\pi_k) \cdot \pi_k^{-1})(\delta_k - f_\varphi^*(\mathbf{X})) - B^-(\pi_k)(\delta_k - f_\varphi^*(\mathbf{X})) \\
&= B^-(\pi_k) \cdot \pi_k^{-1} \cdot (p(\mathbf{X}) - \pi_k)(\delta_k - f_\varphi^*(\mathbf{X})).
\end{aligned}$$

Since $f_\varphi^*(\mathbf{X}) > \delta_k$ when $p(\mathbf{X}) > \pi_k$, and similarly $f_\varphi^*(\mathbf{X}) < \delta_k$ when $p(\mathbf{X}) < \pi_k$, $(p(\mathbf{X}) - \pi_k)(\delta_k - f_\varphi^*(\mathbf{X})) \leq 0$ must always hold. Therefore,

$$\begin{aligned}
\Delta Q_p(\delta_k) &= -\frac{B^-(\pi_k) \cdot |\delta_k - f_\varphi^*(\mathbf{X})|}{\pi_k} \cdot |p(\mathbf{X}) - \pi_k| \\
&\geq C^{-1} \cdot |p(\mathbf{X}) - \pi_k|,
\end{aligned}$$

where $C = \max \left\{ -\frac{\pi_k}{B^-(\pi_k) \cdot |\delta_k - H_j|} : k = 1, \dots, K; j = 0, \dots, K \right\} > 0$. Letting $s = 1$, the desired bound is achieved.

S9 Proof of Theorem 6

Let φ^Y be a minimally consistent piecewise linear surrogate loss for appropriately defined boundaries $\boldsymbol{\pi}$ in $(0, 1)$ at $\boldsymbol{\delta}$. We first show that $\mathcal{H} = \{h_f(\mathbf{x}, y) = \varphi^y(yf(\mathbf{x})) - \varphi^y(yf_\varphi^*(\mathbf{x})) : f \in \mathcal{F}\}$ is a Bernstein class of functions, i.e. that there exists some $B > 1$, $\beta \in (0, 1]$ such that:

$$\mathbb{E}\{h_f(\mathbf{X}, Y)^2\} \leq B \cdot \mathbb{E}\{h_f(\mathbf{X}, Y)\}^\beta.$$

Then, given that h_f is a Bernstein class, we complete the proof by obtaining a tail bound on $\mathbb{E}h_f(\mathbf{X}, Y) - 2\frac{1}{n}\sum_i h_f(\mathbf{x}_i, y_i)$. Following the approach of [11], to derive the Bernstein property of h_f , we first show that $\Delta Q_p(f)$ can be bounded below by a pseudo-norm between f and f_φ^* , denoted $\rho_{\mathbf{X}}(f, f_\varphi^*)$. Then, we show that $\mathbb{E}\{h_f(\mathbf{X}, Y)^2\}$ can be bounded above by $\mathbb{E}\{\rho_{\mathbf{X}}(f, f_\varphi^*)\}$, and combine the two results to show the Bernstein property of h_f . Let

$\rho_{\mathbf{X}}(f, f_\varphi^*)$ be defined as:

$$\rho_{\mathbf{X}}(f, f_\varphi^*) = \begin{cases} p(\mathbf{X})|f - f_\varphi^*| & \text{if } p(\mathbf{X}) < \pi_1, f < H_0 \\ (1 - p(\mathbf{X}))|f - f_\varphi^*| & \text{if } p(\mathbf{X}) > \pi_K, f > H_K \\ |f - f_\varphi^*| & \text{otherwise} \end{cases}$$

Lemma 1. For $p(\mathbf{X}) \in [0, 1]$,

$$\Delta Q_p(f) \geq D^* \cdot \min\{|p(\mathbf{X}) - \pi_1|, |p(\mathbf{X}) - \pi_K|, (1 - \pi_1), \pi_K\} \cdot \rho_{\mathbf{X}}(f, f_\varphi^*).$$

Proof. Since $Q_p(f)$ is convex, $Q_p(f) \geq Q_p(f_\varphi^*) + r \cdot (f - f_\varphi^*)$ for any subgradient, r , of $Q_p(\cdot)$ at f_φ^* . Since φ^Y is piecewise linear, and f_φ^* is as defined above, the set of subgradients are given by:

$$r = \begin{cases} \begin{cases} p(\mathbf{X})B^+(\pi_1) \\ \text{and } p(\mathbf{X})B^+(\pi_1) + (1 - p(\mathbf{X}))B^-(\pi_1) \end{cases} & \text{for } f_\varphi^* = H_0 \\ \begin{cases} p(\mathbf{X})B^+(\pi_1) + (1 - p(\mathbf{X}))B^-(\pi_1) \\ \text{and } p(\mathbf{X})B^+(\pi_1) + (1 - p(\mathbf{X}))B^-(\pi_1) \end{cases} & \text{for } f_\varphi^* = H_1, \dots, H_{K-1} \\ \begin{cases} (1 - p(\mathbf{X}))B^-(\pi_K) \\ \text{and } p(\mathbf{X})B^+(\pi_K) + (1 - p(\mathbf{X}))B^-(\pi_K) \end{cases} & \text{for } f_\varphi^* = H_K \end{cases}$$

Therefore,

$$\begin{aligned} Q_p(f) &\geq Q_p(f_\varphi^*) + r \cdot (f - f_\varphi^*) \\ \Delta Q_p(f) &\geq r \cdot (f - f_\varphi^*) \\ &\geq \begin{cases} (p(\mathbf{X})B^+(\pi_1)) \cdot (f - f_\varphi^*) & \text{if } p(\mathbf{X}) < \pi_1, f < H_0 \\ (p(\mathbf{X})B^+(\pi_1) - (1 - p(\mathbf{X}))B^-(\pi_1)) \cdot (f - f_\varphi^*) & \text{if } p(\mathbf{X}) < \pi_1, f > H_0 \\ (p(\mathbf{X})B^+(\pi_k) - (1 - p(\mathbf{X}))B^-(\pi_k)) \cdot (f - f_\varphi^*) & \text{if } p(\mathbf{X}) \in [\pi_k, \pi_{k+1}), f < H_k \\ (p(\mathbf{X})B^+(\pi_{k+1}) - (1 - p(\mathbf{X}))B^-(\pi_{k+1})) \cdot (f - f_\varphi^*) & \text{if } p(\mathbf{X}) \in [\pi_k, \pi_{k+1}), f > H_k \\ (p(\mathbf{X})B^+(\pi_K) - (1 - p(\mathbf{X}))B^-(\pi_K)) \cdot (f - f_\varphi^*) & \text{if } p(\mathbf{X}) > \pi_K, f < H_K \\ (1 - p(\mathbf{X}))B^-(\pi_K) \cdot (f - f_\varphi^*) & \text{if } p(\mathbf{X}) > \pi_K, f > H_K \end{cases} \end{aligned}$$

Since by definition, $B^+(\pi_1) \leq B^+(\pi_2) \leq \dots \leq B^+(\pi_K)$ and $B^-(\pi_1) \geq B^-(\pi_2) \geq \dots \geq$

$B^-(\pi_K)$, we have:

$$\begin{aligned} -p(\mathbf{X})B^+(\pi_k) + (1 - p(\mathbf{X}))B^-(\pi_k) &\geq -p(\mathbf{X})B^+(\pi_K) + (1 - p(\mathbf{X}))B^-(\pi_K) \\ p(\mathbf{X})B^+(\pi_{k+1}) - (1 - p(\mathbf{X}))B^-(\pi_{k+1}) &\leq p(\mathbf{X})B^+(\pi_1) - (1 - p(\mathbf{X}))B^-(\pi_1). \end{aligned}$$

Therefore, the bound on $\Delta Q_p(f)$ may be rewritten as:

$$\Delta Q_p(f) \geq \begin{cases} |p(\mathbf{X})B^+(\pi_1)| \cdot |f - f_\varphi^*| & \text{if } p(\mathbf{X}) < \pi_1, f < H_0 \\ |p(\mathbf{X})B^+(\pi_1) - (1 - p(\mathbf{X}))B^-(\pi_1)| \cdot |f - f_\varphi^*| & \text{if } p(\mathbf{X}) < \pi_1, f > H_0 \\ |p(\mathbf{X})B^+(\pi_K) - (1 - p(\mathbf{X}))B^-(\pi_K)| \cdot |f - f_\varphi^*| & \text{if } p(\mathbf{X}) > \pi_K, f < H_K \\ |(1 - p(\mathbf{X}))B^-(\pi_K)| \cdot |f - f_\varphi^*| & \text{if } p(\mathbf{X}) > \pi_K, f > H_K \\ \min \{ |p(\mathbf{X})B^+(\pi_1) - (1 - p(\mathbf{X}))B^-(\pi_1)|, \\ |p(\mathbf{X})B^+(\pi_K) - (1 - p(\mathbf{X}))B^-(\pi_K)| \} \cdot |f - f_\varphi^*| & \text{otherwise} \end{cases}.$$

By the consistency of φ^Y , $B^-(\pi_k)/(B^+(\pi_k) + B^-(\pi_k)) = \pi_k$ for all k . Thus, letting $D^* = \min_{k=1, \dots, K} \{|B^+(\pi_k) + B^-(\pi_k)|\} > 0$, $p(\mathbf{X})B^+(\pi_k) - (1 - p(\mathbf{X}))B^-(\pi_k) = (p(\mathbf{X}) - \pi_k)(B^+(\pi_k) + B^-(\pi_k)) \geq D^* \cdot |p(\mathbf{X}) - \pi_k|$. Therefore,

$$\Delta Q_p(f) \geq \begin{cases} |B^+(\pi_1)| \cdot \rho_{\mathbf{X}}(f, f_\varphi^*) & \text{if } p(\mathbf{X}) < \pi_1, f < H_0 \\ D^* \cdot |p(\mathbf{X}) - \pi_1| \cdot \rho_{\mathbf{X}}(f, f_\varphi^*) & \text{if } p(\mathbf{X}) < \pi_1, f > H_0 \\ D^* \cdot |p(\mathbf{X}) - \pi_K| \cdot \rho_{\mathbf{X}}(f, f_\varphi^*) & \text{if } p(\mathbf{X}) > \pi_K, f < H_K \\ |B^-(\pi_K)| \cdot \rho_{\mathbf{X}}(f, f_\varphi^*) & \text{if } p(\mathbf{X}) > \pi_K, f > H_K \\ D^* \cdot \min \{ |p(\mathbf{X}) - \pi_1|, |p(\mathbf{X}) - \pi_K| \} \cdot \rho_{\mathbf{X}}(f, f_\varphi^*) & \text{otherwise} \end{cases}.$$

Since $|B^+(\pi_1)| \geq D^* \cdot (1 - \pi_1)$, $|B^-(\pi_K)| \geq D^* \cdot \pi_K$, we have for $p(\mathbf{X}) \in [0, 1]$:

$$\Delta Q_p(f) \geq D^* \cdot \min \{ |p(\mathbf{X}) - \pi_1|, |p(\mathbf{X}) - \pi_K|, (1 - \pi_1), \pi_K \} \cdot \rho_{\mathbf{X}}(f, f_\varphi^*).$$

□

Lemma 2. *If $|f| < B$ for all $f \in \mathcal{F}$,*

$$\mathbb{E}_{Y|\mathbf{X}} \{h_f(\mathbf{X}, Y)^2\} \leq L^2(B + M) \cdot \rho_{\mathbf{X}}(f, f_\varphi^*)$$

for $L, M \geq 0$.

Proof. We first decompose the conditional expectation as:

$$\begin{aligned}\mathbb{E}_{Y|\mathbf{X}}\{h_f(\mathbf{X}, Y)^2\} &= \mathbb{E}_{Y|\mathbf{X}}\left\{\left(\varphi^Y(Yf(\mathbf{X})) - \varphi^Y(Yf^*(\mathbf{X}))\right)^2\right\} \\ &= p(\mathbf{X})(\varphi^+(f(\mathbf{X})) - \varphi^+(f_\varphi^*(\mathbf{X})))^2 \\ &\quad + (1 - p(\mathbf{X}))(\varphi^-(f(\mathbf{X})) - \varphi^-(f_\varphi^*(\mathbf{X})))^2.\end{aligned}$$

Note that if $f(\mathbf{X}) \leq H_0$ and $p(\mathbf{X}) \leq \pi_1$, then $\varphi^-(f(\mathbf{X})) = 0$ and $\varphi^-(f_\varphi^*(\mathbf{X})) = 0$. Similarly, if $f(\mathbf{X}) \geq H_K$ and $p(\mathbf{X}) \geq \pi_K$, then $\varphi^+(f(\mathbf{X})) = 0$ and $\varphi^+(f_\varphi^*(\mathbf{X})) = 0$. Therefore,

$$E_{Y|\mathbf{X}}\{h_f(\mathbf{X}, Y)^2\} = \begin{cases} p(\mathbf{X})(\varphi^+(f(\mathbf{X})) - \varphi^+(f_\varphi^*(\mathbf{X})))^2 & \text{if } f(\mathbf{X}) \leq H_0, p(\mathbf{X}) \leq \pi_1 \\ (1 - p(\mathbf{X}))(\varphi^-(f(\mathbf{X})) - \varphi^-(f_\varphi^*(\mathbf{X})))^2 & \text{if } f(\mathbf{X}) \geq H_K, p(\mathbf{X}) \geq \pi_K \\ p(\mathbf{X})(\varphi^+(f(\mathbf{X})) - \varphi^+(f_\varphi^*(\mathbf{X})))^2 \\ \quad + (1 - p(\mathbf{X}))(\varphi^-(f(\mathbf{X})) - \varphi^-(f_\varphi^*(\mathbf{X})))^2 & \text{otherwise} \end{cases}.$$

Let $L = \max\{B^+(\pi_1), B^-(\pi_K)\}$ denote the Lipschitz constant for φ^Y , and let $M = \max\{|H_0|, |H_K|\}$ denote the bound on f_φ^* , such that $|f_\varphi^*(\mathbf{X})| \leq M$ for all \mathbf{X} . Then,

$$E_{Y|\mathbf{X}}\{h_f(\mathbf{X}, Y)^2\} \leq L^2(B + M) \cdot \rho_{\mathbf{X}}(f, f_\varphi^*),$$

where $\rho_{\mathbf{X}}$ is as defined above. □

Lemma 3. *If $p(\mathbf{X})$ satisfies the margin condition (14) at $\boldsymbol{\pi} = \{\pi_1, \dots, \pi_K\}$ with parameters A, α , then for any class \mathcal{F} of measurable uniformly bounded functions, the class $\mathcal{H} = \{h_f(\mathbf{X}, Y) : f \in \mathcal{F}\}$ is a Bernstein class with exponent $\beta = \alpha/(1 + \alpha)$.*

Proof. Let E_1 denote the event that $|p(\mathbf{X}) - \pi|$ is the minimizer over the set $\{|p(\mathbf{X}) - \pi_1|, |p(\mathbf{X}) - \pi_K|, (1 - \pi_1), \pi_K\}$, and let E_2, E_3, E_4 similarly denote the corresponding events for $|p(\mathbf{X}) - \pi_K|, (1 - \pi_1)$ and π_K . Using \mathbf{I}_E to denote the indicator for event E , by Lemma 2 we have:

$$\begin{aligned}\mathbb{E}\{h_f(\mathbf{X}, Y)\} &\geq D^* \cdot \mathbb{E}\left\{\min\{|p(\mathbf{X}) - \pi_1|, |p(\mathbf{X}) - \pi_K|, (1 - \pi_1), \pi_K\} \cdot \rho_{\mathbf{X}}(f, f_\varphi^*)\right\} \\ &= D^* \cdot \mathbb{E}\left\{\rho_{\mathbf{X}}(f, f_\varphi^*) \cdot \{\mathbf{I}_{E_1} \cdot |p(\mathbf{X}) - \pi_1| + \mathbf{I}_{E_2} \cdot |p(\mathbf{X}) - \pi_K| \right. \\ &\quad \left. + \mathbf{I}_{E_3} \cdot (1 - \pi_1) + \mathbf{I}_{E_4} \cdot \pi_K\}\right\}.\end{aligned}$$

Let $t_{\max} = \min_{k=1, \dots, K+1}\{\pi_k - \pi_{k-1}\}$, where $\pi_0 = 0, \pi_{K+1} = 1$. Given the margin condition,

for all k , there exists some $A \geq 0, \alpha \geq 0$ such that for all $t \in [0, t_{\max})$,

$$\mathbb{P}\{|p(\mathbf{X}) - \pi_k| \leq t\} \leq At^\alpha,$$

for $k = 1, \dots, K$. Therefore, letting B and M denote the bounds on f and f_φ^* given in the proof of Lemma 2,

$$\begin{aligned} \mathbb{E}\{\rho_{\mathbf{X}}(f, f_\varphi^*) \cdot |p(\mathbf{X}) - \pi_1| \cdot \mathbf{I}_{E_1}\} &\geq t \cdot \mathbb{E}\{\rho_{\mathbf{X}}(f, f_\varphi^*) \cdot \mathbf{I}\{|p(\mathbf{X}) - \pi_1| > t\} \cdot \mathbf{I}_{E_1}\} \\ &\geq t \cdot [\mathbb{E}\{\rho_{\mathbf{X}}(f, f_\varphi^*) \cdot \mathbf{I}_{E_1}\} - (B + M) \cdot At^\alpha], \end{aligned}$$

and similarly,

$$\begin{aligned} \mathbb{E}\{\rho_{\mathbf{X}}(f, f_\varphi^*) \cdot |p(\mathbf{X}) - \pi_K| \cdot \mathbf{I}_{E_2}\} &\geq t \cdot [\mathbb{E}\{\rho_{\mathbf{X}}(f, f_\varphi^*) \cdot \mathbf{I}_{E_2}\} - (B + M) \cdot At^\alpha] \\ \mathbb{E}\{\rho_{\mathbf{X}}(f, f_\varphi^*) \cdot (1 - \pi_1) \cdot \mathbf{I}_{E_3}\} &\geq t \cdot [\mathbb{E}\{\rho_{\mathbf{X}}(f, f_\varphi^*) \cdot \mathbf{I}_{E_3}\} - (B + M) \cdot \mathbf{I}\{(1 - \pi_1) < t, (1 - \pi_1) \leq \pi_K\}] \\ \mathbb{E}\{\rho_{\mathbf{X}}(f, f_\varphi^*) \cdot \pi_K \cdot \mathbf{I}_{E_4}\} &\geq t \cdot [\mathbb{E}\{\rho_{\mathbf{X}}(f, f_\varphi^*) \cdot \mathbf{I}_{E_4}\} - (B + M) \cdot \mathbf{I}\{\pi_K < t, \pi_K < (1 - \pi_1)\}]. \end{aligned}$$

Assume without loss of generality that $\pi_K < (1 - \pi_1)$. Let

$$t = \left(\frac{\mathbb{E}\{\rho_{\mathbf{X}}(f, f_\varphi^*)\}}{C \cdot 2A(B + M)} \right)^{1/\alpha},$$

where $C \geq \max\{2, (2A\pi_K^\alpha)^{-1}\}$. Then, since $\mathbb{E}\{\rho_{\mathbf{X}}(f, f_\varphi^*)\} \leq (B + M)$, we have $t < \pi_K$. Combining the above inequalities, we have:

$$\begin{aligned} \mathbb{E}\{h_f(\mathbf{X}, Y)\} &\geq D^* \cdot t \cdot [\mathbb{E}\{\rho_{\mathbf{X}}(f, f_\varphi^*)\} - (B + M)(2At^\alpha)] \\ &\geq D^* \cdot \left(\frac{\mathbb{E}\{\rho_{\mathbf{X}}(f, f_\varphi^*)\}}{C \cdot 2A(B + M)} \right)^{1/\alpha} [\mathbb{E}\{\rho_{\mathbf{X}}(f, f_\varphi^*)\} - C^{-1}\mathbb{E}\{\rho_{\mathbf{X}}(f, f_\varphi^*)\}] \\ &\geq D^* \cdot \left(\frac{C - 1}{C} \right) \cdot \left(\frac{1}{C \cdot 2A(B + M)} \right)^{1/\alpha} \cdot \mathbb{E}\{\rho_{\mathbf{X}}(f, f_\varphi^*)\}^{(1+\alpha)/\alpha} \\ \mathbb{E}\{\rho_{\mathbf{X}}(f, f_\varphi^*)\} &\leq \left[\left(\frac{C}{C - 1} \right) \cdot D^* \cdot (C \cdot 2A(B + M))^{1/\alpha} \right]^{\alpha/(1+\alpha)} \cdot \mathbb{E}\{h_f(\mathbf{X}, Y)\}^{\alpha/(1+\alpha)}. \end{aligned}$$

Combining with the result of Lemma 2, and noting that $\mathbb{E}\{\rho_{\mathbf{X}}(f, f_\varphi^*)\} = \mathbb{E}_{\mathbf{X}}\{\rho_{\mathbf{X}}(f, f_\varphi^*)\}$, we have:

$$\mathbb{E}\{h_f(\mathbf{X}, Y)^2\} = \mathbb{E}_{\mathbf{X}}\{\mathbb{E}_{Y|\mathbf{X}}\{h_f(\mathbf{X}, Y)^2\}\}$$

$$\begin{aligned}
&\leq L^2(B + M) \cdot \mathbb{E}_{\mathbf{X}}\{\rho_{\mathbf{X}}(f, f_{\varphi}^*)\} \\
&\leq L^2(B + M) \cdot \left[\left(\frac{C}{C-1} \right) \cdot D^* \cdot (C \cdot 2A(B + M))^{1/\alpha} \right]^{\alpha/(1+\alpha)} \cdot \mathbb{E}\{h_f(\mathbf{X}, Y)\}^{\alpha/(1+\alpha)},
\end{aligned}$$

such that h_f is a Bernstein class. \square

Let B' and β be defined such that $\mathbb{E}\{h_f(\mathbf{X}, Y)^2\} \leq B' \cdot \mathbb{E}\{h_f(\mathbf{X}, Y)\}^{\beta}$. Let \hat{f}_n denote the empirical minimizer in \mathcal{F} of $\varphi^y(yf(\mathbf{x}))$ over a training sample of size n . We first bound the excess φ -risk by:

$$\begin{aligned}
\Delta Q(\hat{f}_n) &= \mathbb{E}\{h_{\hat{f}_n}(\mathbf{X}, Y)\} \\
&= 2 \left(\frac{1}{n} \sum_{i=1}^n h_{\hat{f}_n}(\mathbf{x}_i, y_i) \right) + \left(\mathbb{E}\{h_{\hat{f}_n}(\mathbf{X}, Y)\} - 2 \left(\frac{1}{n} \sum_{i=1}^n h_{\hat{f}_n}(\mathbf{x}_i, y_i) \right) \right) \\
&\leq \sup_{f \in \mathcal{F}_B} \left(\mathbb{E}\{h_f(\mathbf{X}, Y)\} - 2 \left(\frac{1}{n} \sum_{i=1}^n h_f(\mathbf{x}_i, y_i) \right) \right).
\end{aligned}$$

Note that,

$$\sup_{f \in \mathcal{F}_B} \left(\mathbb{E}\{h_f(\mathbf{X}, Y)\} - 2 \left(\frac{1}{n} \sum_{i=1}^n h_f(\mathbf{x}_i, y_i) \right) \right) \leq \frac{3L}{n} + \sup_{f \in \mathcal{F}_n} \left(\mathbb{E}\{h_f(\mathbf{X}, Y)\} - 2 \left(\frac{1}{n} \sum_{i=1}^n h_f(\mathbf{x}_i, y_i) \right) \right),$$

where \mathcal{F}_n is a minimal $1/n$ -net of \mathcal{F}_B . Now applying Bernstein's inequality,

$$\begin{aligned}
&\mathbb{P} \left\{ \sup_{f \in \mathcal{F}_n} \left(\mathbb{E}\{h_f(\mathbf{X}, Y)\} - 2 \left(\frac{1}{n} \sum_{i=1}^n h_f(\mathbf{x}_i, y_i) \right) \right) \geq t \right\} \\
&\leq N_n \cdot \exp \left\{ - \frac{n(t + \mathbb{E}\{h_f(\mathbf{X}, Y)\})^2/8}{\mathbb{E}\{h_f(\mathbf{X}, Y)^2\} + (2LB)(t + \mathbb{E}\{h_f(\mathbf{X}, Y)\})/6} \right\}.
\end{aligned}$$

Using the fact that h_f is a Bernstein class, and noting that for $\beta \in [0, 1)$, $z^{\beta} \leq 1 + z$ for all $z > 0$,

$$\begin{aligned}
\frac{\mathbb{E}\{h_f(\mathbf{X}, Y)^2\}}{t + \mathbb{E}\{h_f(\mathbf{X}, Y)\}} &\leq B' \cdot \frac{\mathbb{E}\{h_f(\mathbf{X}, Y)\}^{\beta}}{t + \mathbb{E}\{h_f(\mathbf{X}, Y)\}} \\
&\leq B' \cdot \frac{1 + \mathbb{E}\{h_f(\mathbf{X}, Y)\}}{t + \mathbb{E}\{h_f(\mathbf{X}, Y)\}} \\
&\leq B' \cdot t^{-1}.
\end{aligned}$$

Therefore,

$$\begin{aligned}
\mathbb{P}\left\{\sup_{f \in \bar{\mathcal{F}}_n} \left(\mathbb{E}\{h_f(\mathbf{X}, Y)\} - 2\left(\frac{1}{n} \sum_{i=1}^n h_f(\mathbf{x}_i, y_i)\right)\right) \geq t\right\} \\
\leq N_n \cdot \exp\left\{-\frac{n}{8} \cdot \frac{(t + \mathbb{E}\{h_f(\mathbf{X}, Y)\})}{B' + \frac{LB}{3}}\right\} \\
\leq N_n \cdot \exp\left\{-\frac{nt}{8} \cdot \left(\frac{B'}{t} + \frac{LB}{3}\right)^{-1}\right\}.
\end{aligned}$$

The proof is complete by noting that the necessary bound holds with probability γ for:

$$t = 4 \cdot \frac{LB}{3} \cdot \frac{\log(N_n/\gamma)}{n} + \left(\left(4 \cdot \frac{LB}{3} \cdot \frac{\log(N_n/\gamma)}{n}\right)^2 + 8 \cdot B' \cdot \frac{\log(N_n/\gamma)}{n}\right)^{1/2}.$$