
Generalizing Consistent Multi-Class Classification with Rejection to be Compatible with Arbitrary Losses

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 *Classification with rejection* (CwR) refrains from making a prediction to avoid critical
2 misclassification when encountering test samples that are difficult to classify.
3 Though previous methods for CwR have been provided with theoretical guarantees,
4 they are only compatible with certain loss functions, making them not flexible
5 enough when the loss needs to be changed with the dataset in practice. In this paper,
6 we derive a novel formulation for CwR that can be equipped with arbitrary loss
7 functions while maintaining the theoretical guarantees. First, we show that K -class
8 CwR is equivalent to a $(K + 1)$ -class classification problem on the original data
9 distribution with an augmented class, and propose an empirical risk minimization
10 formulation to solve this problem with an estimation error bound. Then, we find a
11 necessary and sufficient condition for the learning *consistency* of the surrogates con-
12 structed on our proposed formulation equipped with any classification-calibrated
13 multi-class losses, where consistency means the surrogate risk minimization implies
14 the target risk minimization for CwR. Finally, experiments on benchmark
15 datasets validate the effectiveness of our proposed method.

16 1 Introduction

17 In risk-sensitive multi-class classification applications (e.g., medical diagnosis, healthcare, au-
18 tonomous driving, and product inspections [12, 21, 43]), misclassification can cause serious or
19 even fatal consequences. To alleviate this issue, many studies have been conducted on *classification*
20 *with rejection* (CwR) [10, 6, 61, 12, 13, 15, 21, 51, 47, 43, 8], which can abstain from making an
21 unsure prediction to prevent such critical misclassification.

22 Most of the previous studies follow the framework that provides the reject option with a pre-defined
23 cost c which is lower than the misclassification cost 1. Given cost c , the problem is further formulated
24 as a risk minimization problem that aims to minimize the expectation of the zero-one- c loss, i.e., the
25 zero-one- c risk. With the risk minimization process, the obtained classifier can balance the cost of
26 rejection and prediction by choosing to incur a rejection cost c if the misclassification risk is high.

27 Due to the discontinuous nature of the zero-one- c loss, recent works focused on finding its continuous
28 surrogates to make the optimization problem tractable. A basic requirement for surrogate losses
29 is the *consistency* [63, 7, 53, 46], i.e., the surrogate risk minimization implies the zero-one- c risk
30 minimization. Moreover, compared with the traditional K -class classification task where decisions
31 are normally made from the index of the maximum coordinate of a K -dimensional scoring function,
32 the design of decision criteria in the CwR task is more elusive due to the existence of a reject option.

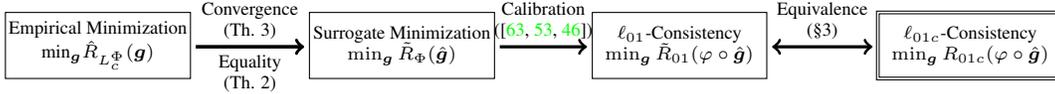


Figure 1: Overview of the construction of consistent surrogates for classification with rejection in this work.

33 By adopting different classification and rejection criteria, various surrogates of the zero-one- c loss
34 have been proposed with consistency analyses [6, 61, 12, 13, 47, 43, 8].

35 Classical studies focused on developing *confidence-based methods*[6, 61, 47, 43], which use the
36 outputs of classifiers as confidence values and set a real-valued threshold as the rejection rule.
37 Representative methods [61, 43] used surrogates that depend on *class-posterior possibility estimation*
38 (CPE) [49, 56], which is challenging when using deep models [24]. Though some of them [6, 47, 34,
39 23] could avoid CPE, most of them applied the modification of non-differentiable hinge/ramp-like
40 surrogates, and their performance was only validated with linear models.

41 To avoid the use of the confidence threshold, Cortes et al. [12] provided an upper bound of the zero-
42 one- c loss as the surrogate that allows the use of a separated rejector and can be trained simultaneously
43 with the classifier, which is regarded as *classifier-rejector methods*. Though these methods achieved
44 state-of-the-art performance in binary classification scenarios, they only provided a consistency
45 guarantee for hinge-like and exponential losses and cannot be directly generalized to the multi-class
46 scenario as shown in Ni et al. [43]. Charoenphakdee et al. [8] showed that K -class CwR can be
47 decomposed into K binary cost-sensitive classification problems [16, 50, 11] and proposed a family
48 of surrogates are the ensembles of arbitrary binary classification losses, which can avoid CPE and the
49 use of confidence threshold with properly chosen losses when the cost function is constant. Mozannar
50 and Sontag [40] provided a modified version of the cross entropy loss as the surrogate for the task of
51 learning to defer [40, 41] that can also be used in CwR, while its optimal solution still relies on CPE.
52 In summary, previous works only took limited types of losses into consideration, and there lacks a
53 theoretically grounded framework that can cover all the surrogates used in multi-class classification.

54 In this paper, we propose a novel framework for CwR that allows the use of arbitrary surrogate losses
55 used in traditional multi-class classification as long as they are classification-calibrated, including
56 but not limited to the well-known cross entropy loss, mean absolute error, focal loss [32, 9], and the
57 pairwise/one-versus-all generalizations of binary margin losses [63]. Thanks to the flexible choices of
58 losses, we be free of the restricted analyses on the consistency of certain surrogates. An overview of
59 our framework is shown in Figure 1. We summarize the main contributions of this work as follows:

- 60 • We disclose the equivalence between K -class CwR and a $(K + 1)$ -class classification problem
61 on the original data distribution with an augmented class, by showing the equality between their
62 classification risks.
- 63 • We propose a formulation of surrogates for ℓ_{01c} that can recover the surrogate risk of a $(K + 1)$ -
64 class classification task only with the K -class training distribution, and derive an estimation error
65 bound for its empirical risk minimization.
- 66 • We find a necessary and sufficient condition for the consistency of the proposed family of surrogates
67 *w.r.t.* the zero-one- c loss that allows the use of any calibrated multi-class surrogates.
- 68 • We for the first time provide an analysis on the calibration of the *generalized cross entropy* loss
69 [64] that benefits from both the cross entropy loss and mean absolute error, and experimentally
70 demonstrate that it is suitable for our proposed framework.

71 2 Preliminaries

72 In this section, we provide preliminary knowledge of CwR and calibrated surrogate losses, and
73 discuss the consistency in CwR.

74 **2.1 Classification with Rejection**

The problem setting of CwR is based on the cost-based framework [10]. Let us denote by \mathcal{X} the feature space, $\mathcal{Y} = \{1, 2, \dots, K\}$ the label space, and $\mathcal{Y}^\circledast = \{1, 2, \dots, K, \circledast\}$ the label space with a reject option. We are given instance-label pairs $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$ independently and identically drawn from an underlying distribution with probability density $p(\mathbf{x}, y)$. The goal of CwR is to train a classifier $f : \mathcal{X} \rightarrow \mathcal{Y}^\circledast$ that can abstain from making a decision, where \circledast denotes the reject option. The evaluation metric of this task is the zero-one- c loss ℓ_{01c} , which can be expressed as a variant of the traditional zero-one loss $\ell_{01}(f(\mathbf{x}), y) = \mathbb{I}[f(\mathbf{x}) \neq y]$:

$$\ell_{01c}(f(\mathbf{x}), y) = \begin{cases} c, & f(\mathbf{x}) = \circledast, \\ \mathbb{I}[f(\mathbf{x}) \neq y], & f(\mathbf{x}) \in \{1, 2, \dots, k\}, \end{cases}$$

75 where $\mathbb{I}[\cdot]$ is the Iverson bracket notation as suggested by Knuth [28] and the cost c can be further
76 extended to an instance-dependent function $c(\mathbf{x})$. Our goal is to train a classifier that can minimize
77 the expectation of ℓ_{01c} over the data distribution:

$$R_{01c}(f) = \mathbb{E}_{p(\mathbf{x}, y)}[\ell_{01c}(f(\mathbf{x}), y)]. \quad (1)$$

78 Let us denote by $f^* = \operatorname{argmin}_f R_{01c}(f)$ the Bayes optimal solution and $\boldsymbol{\eta}(\mathbf{x}) = \{p(y|\mathbf{x})\}_{y=1}^K$ the
79 posterior probabilities. When evaluated by ℓ_{01c} , a classifier receives a standard classification error in
80 $\{0, 1\}$ if it makes a prediction and a cost of c if it does not make a prediction (i.e., chooses the reject
81 option). Intuitively, an optimal solution f^* should balance the possibility of misclassification and the
82 rejection cost c . This explanation is theoretically justified by Chow’s rule [10]:

Definition 1. (Chow’s Rule) A classifier $f : \mathcal{X} \rightarrow \mathcal{Y}^\circledast$ is the optimal solution of (1) if and only if it meets the following condition almost surely:

$$f(\mathbf{x}) = \begin{cases} \circledast, & \max_y \eta_y(\mathbf{x}) \leq 1 - c, \\ \operatorname{argmax}_y \eta_y(\mathbf{x}), & \text{else.} \end{cases}$$

83 Chow’s rule shows that the optimal solution should refrain from making a decision if the most
84 competent prediction of an example is still not confident enough given a rejection cost c .

85 **2.2 Calibrated Surrogate Losses**

86 Most classification problems can be formalized as the minimization of the target risk, which is the
87 expectation of a target loss. Then *empirical risk minimization* (ERM) is conducted to obtain models
88 with performance guarantees. However, most of the target losses are discontinuous, e.g., the zero-one
89 loss in multi-class classification and the Hamming/ranking loss in multi-label classification [19].
90 Therefore, directly optimizing them is usually difficult and even NP-hard [17].

91 In order to optimize the target risk efficiently, surrogate risk minimization is preferred that minimizing
92 the expectation of a continuous surrogate loss instead, e.g., the hinge loss in binary classification
93 and the cross entropy loss in multi-class classification. For the statistical consistency of learning,
94 *calibration* [52] is considered as a basic requirement for surrogate losses, which is a pointwise version
95 of consistency and means that the minimization of the surrogate loss yields that of the target loss for
96 each possible sample. A commonly adopted definition of the calibration of surrogates in multi-class
97 classification is given as follows:

Definition 2. (ℓ_{01} -Calibration [7, 53, 46]) For a K -class classification problem with target loss ℓ_{01} , we say $\Phi : \mathbb{R}^K \times \mathcal{Y} \rightarrow \mathbb{R}_+$ is ℓ_{01} -calibrated if for any $\mathbf{p} \in \Delta^K$:

$$\inf_{\mathbf{u} \in \mathbb{R}^K, \mathbf{u} \notin \operatorname{argmin}_{\mathbf{u}} \mathbf{p}^T \mathbf{L}_{01}(\mathbf{u})} \mathbf{p}^T \Phi(\mathbf{u}) > \inf_{\mathbf{u} \in \mathbb{R}^K} \mathbf{p}^T \Phi(\mathbf{u}),$$

98 where $\Phi(\mathbf{u}) = \{\Phi(\mathbf{u}, y)\}_{y=1}^K$, $\mathbf{L}_{01}(\mathbf{u}) = \{\ell_{01}(\operatorname{argmax}_{y' \in \mathcal{Y}} \mathbf{u}_{y'}, y)\}_{y=1}^K$.

99 The definition of ℓ_{01} -calibration requires that a surrogate loss should be able to distinguish between
100 optimal solutions and non-optimal ones *w.r.t.* any potential posterior distribution \mathbf{p} . This property

Table 1: Comparisons between our proposed method and previous works of multi-class classification with rejection. Since our method is induced from a $(K+1)$ -class classification problem, we can render a consistent learning guarantee with arbitrary surrogate losses that are calibrated *w.r.t.* the zero-one loss. Thanks to the abundant choices of losses, our proposed method can avoid CPE and the use of confidence thresholds.

Method	CPE-Free	Instance-Dependent Cost	Confidence Threshold-Free	Arbitrary Losses
[47]	✓	✓	✗	✗
[43]	✗	✓	✗	✗
[40]	✗	✓	✓	✗
[8]	✓	✗	✓	✗
Proposed	✓	✓	✓	✓

101 is shown to be a necessary and sufficient condition for the statistical consistency of surrogate risk
 102 minimization, and fruitful research on the verification of ℓ_{01} -calibrated surrogates has been conducted
 103 [7, 63, 53, 46, 45, 18].

104 Besides multi-class classification, the calibration of surrogate losses also has been studied in various
 105 aspects of statistical learning, including but not limited to, multi-label classification [19, 62, 29, 57],
 106 AUC optimization [20, 38], general linear-fractional utility maximization [3], cost-sensitive learning
 107 [11, 50], top- K classification [31, 60], and adversarially robust classification [4, 2, 1].

108 2.3 Consistency in Classification with Rejection

109 In the field of CwR, we are also interested in the consistency of surrogate losses. Let $\mathcal{C} \subset \mathbb{R}^d$ where
 110 $d \in \mathbb{N}$ and $\Phi : \mathcal{C} \times \mathcal{Y} \rightarrow \mathbb{R}_+$ is a surrogate loss, the consistency is defined as follows:

111 **Definition 3.** (ℓ_{01c} -Consistency) A surrogate loss $\Phi : \mathcal{C} \times \mathcal{Y} \rightarrow \mathbb{R}_+$ is ℓ_{01c} -consistent if there
 112 exists a function $\varphi : \mathcal{C} \rightarrow \mathcal{Y}^{\otimes}$ for all probability distributions and all the sequences of functions
 113 $\{\mathbf{g}_i\}_{i \in \mathbb{N}} : \mathcal{X} \rightarrow \mathcal{C}$:

$$R_{\Phi}(\mathbf{g}_i) \rightarrow R_{\Phi}^* \Rightarrow R_{01c}(\varphi \circ \mathbf{g}_i) \rightarrow R_{01c}^*, \quad (2)$$

114 where $R_{\Phi}(\mathbf{g}) = \mathbb{E}_{p(\mathbf{x}, y)}[\Phi(\mathbf{g}(\mathbf{x}), y)]$, $R_{\Phi}^* = \inf_{\mathbf{g}: \mathcal{X} \rightarrow \mathcal{C}} R_{\Phi}(\mathbf{g})$, and $R_{01c}^* = \inf_{f: \mathcal{X} \rightarrow \mathcal{Y}^{\otimes}} R_{01c}(f)$.

115 This definition is inspired by the problem of general multi-class classification [46]. For an ℓ_{01c} -
 116 consistent surrogate loss Φ , we can safely minimize the surrogate risk R_{Φ} instead while remaining
 117 the consistency guarantee of R_{01c} .

118 To ensure the consistency of Φ , it is routine to discuss the calibration of surrogate losses. However,
 119 unlike the classical multi-class classification problem, where φ is usually an argmax operator, the
 120 design of φ in the field of CwR can be quite complicated and hard to be unified, which makes it
 121 difficult to directly conduct calibration analysis on Φ . The flexibility of φ also limits the discussions
 122 to specific types of surrogate losses. In Ramaswamy et al. [47], the authors considered the multi-class
 123 extensions of the hinge-loss with a confidence threshold. Ni et al. [43] indicated that the confidence-
 124 based method is indispensable and only focuses on class probability estimation via surrogate risk
 125 minimization. Both of Mozannar and Sontag [40] and Charoenphakdee et al. [8] gave surrogate
 126 losses for the zero-one- c loss that does not depend on the accurate estimation of the class probability,
 127 while Mozannar and Sontag [40] focused on a variant of the cross entropy loss and Charoenphakdee
 128 et al. [8] constructed calibrated surrogate losses with the ensemble of K calibrated losses for binary
 129 classification.

130 In this paper, instead of directly discussing the calibration of surrogate Φ , we show that there is an
 131 equivalence between classical multi-class classification and CwR. Based on this equivalence, we
 132 show that it is sufficient for Φ to be ℓ_{01c} -consistent by letting it be a simple variant of **any** calibrated
 133 surrogate loss *w.r.t.* the traditional zero-one loss ℓ_{01} . The comparison of the proposed method and
 134 related works is shown in Table 1.

135 **3 Equivalence between Classification with Rejection and Ordinary**
 136 **Classification**

137 In this section, we first show that the risk $R_{01c}(f)$ can be formalized as a $(K+1)$ -class classification
 138 problem, and show that we can obtain ℓ_{01c} -consistent surrogates with a variant of any calibrated
 139 surrogate *w.r.t.* ℓ_{01} , which enables the use of $\mathcal{C} \subset \mathbb{R}^{K+1}$ and $\varphi(\cdot) = \operatorname{argmax}(\cdot)$ as in the traditional
 140 multi-class classification tasks. We also show that such equivalence also holds when the cost c
 141 depends on sample \mathbf{x} . The proof of the conclusions in this section can be found in Appendix A.

142 We start by considering the following distribution $\mathcal{D}_c^\circledast$ over $\mathcal{X} \times \mathcal{Y}^\circledast$ with probability density $\tilde{p}(\mathbf{x}, \tilde{y})$:

Definition 4. (Self-Augmented Distribution) A distribution $\mathcal{D}_c^\circledast$ is called a c -self-augmented distribu-
 tion *w.r.t.* \mathcal{D} if its probability density meets the following conditions:

$$\tilde{p}(\mathbf{x}, \tilde{y}) = \begin{cases} \frac{p(\mathbf{x}, y)}{2-c}, & \tilde{y} \in \{1, 2, \dots, K\}, \\ \frac{(1-c)p(\mathbf{x})}{2-c}, & \tilde{y} = \circledast. \end{cases}$$

143 It can be seen that distribution $\mathcal{D}_c^\circledast$ shares the same marginal density of \mathbf{x} as the original distribution
 144 \mathcal{D} while $\mathcal{D}_c^\circledast$ has an augmented class \circledast with class possibility determined by the rejection cost c .
 145 Based on the connection between $\mathcal{D}_c^\circledast$ and \mathcal{D} , we can further explore the relation between the two
 146 tasks: classification on $\mathcal{D}_c^\circledast$ and CwR on \mathcal{D} .

Theorem 1. For any classifier $f : \mathcal{X} \rightarrow \mathcal{Y}^\circledast$, the following equation holds:

$$R_{01c}(f) - R_{01c}^* = (2 - c) \left(\tilde{R}_{01}(f) - \tilde{R}_{01}^* \right),$$

147 where $\tilde{R}_{01}(f) = \mathbb{E}_{\tilde{p}(\mathbf{x}, \tilde{y})}[\ell_{01}(f(\mathbf{x}), \tilde{y})]$ and $\tilde{R}_{01}^* = \inf_{f: \mathcal{X} \rightarrow \mathcal{Y}^\circledast} \tilde{R}_{01}(f)$.

148 This equation reveals the equivalence between the two tasks in a straightforward manner. Since
 149 the multiplication of the classification risk on $\mathcal{D}_c^\circledast$ with a positive constant is equal to $R_{01c}(f)$, the
 150 minimization of $\tilde{R}_{01}(f)$ immediately yields the minimization of $R_{01c}(f)$ and vice versa. Furthermore,
 151 according to the linear correlation between $\tilde{R}_{01}(f)$ and $R_{01c}(f)$, we can directly quantify the excess
 152 error $R_{01c}(f) - R_{01c}^*$ by bounding $\tilde{R}_{01}(f) - \tilde{R}_{01}^*$, which is an easier work thanks to the existing
 153 research of multi-class classification. In conclusion, risk minimization with $\tilde{R}_{01}(f)$ can also give a
 154 classifier with a rejection option with the optimality guarantee, and then we can consider a surrogate
 155 risk minimization problem for multi-class classification instead of CwR.

156 When the cost $c(\mathbf{x})$ is an instance-dependent function, we show that such equivalence still holds with a
 157 minor modification. Considering the reweighted zero-one loss: $\bar{\ell}_{01}(f(\mathbf{x}), y) = (2 - c(\mathbf{x})) \mathbb{I}[f(\mathbf{x}) \neq y]$
 158 and its expectation $\bar{R}_{01}(f)$ on $\mathcal{D}_c^\circledast$, we have the following conclusion:

Corollary 1. For any classifier $f : \mathcal{X} \rightarrow \mathcal{Y}^\circledast$, the following inequalities holds:

$$\bar{R}_{01}(f) - \bar{R}_{01}^* = R_{01c}(f) - R_{01c}^*.$$

159 It is obvious that Lemma 1 is a special case of Lemma 1 with constant cost functions. Though here
 160 we consider a reweighted classification task, the calibration result of multi-class surrogate losses can
 161 still be applied without any modification since the minimization of $\bar{R}_{01}(f)$ can be seen as ordinary
 162 classification risk minimization with a slightly different marginal density $p'(\mathbf{x})$, which does not affect
 163 the calibration result since the class-posterior possibilities remain unchanged. All the conclusions in
 164 the rest of this paper can be extended to the scenario of instance-dependent cost and we provide them
 165 in Appendix G.

166 **4 ℓ_{01c} -Consistent Surrogates with Arbitrary ℓ_{01} -Calibrated Losses**

167 According to the discussions in Section 3, CwR can be safely replaced by multi-class classification
 168 on a special distribution $\mathcal{D}_c^\circledast$. Following the practice of surrogate risk minimization in multi-class
 169 classification, we can replace the zero-one loss ℓ_{01} with a surrogate risk $\Phi : \mathbb{R}^{K+1} \times \mathcal{Y} \cup \{K+1\} \rightarrow$

170 \mathbb{R}_+ and minimizing the surrogate risk with a score-based classifier $\mathbf{g} : \mathcal{X} \rightarrow \mathbb{R}^{K+1}$ instead, which is
 171 defined as follows:

$$\tilde{R}_\Phi(\mathbf{g}) = \mathbb{E}_{\tilde{p}(\mathbf{x}, \tilde{y})}[\Phi(\mathbf{g}(\mathbf{x}), t(\tilde{y}))], \quad (3)$$

where $t(\tilde{y}) = K + 1$ if $\tilde{y} = \textcircled{0}$ and $t(\tilde{y}) = \tilde{y}$ otherwise. (3) is a typical formulation of the multi-class classification risk and we can asymptotically minimize it following the ERM framework [54]. After the risk minimization process, the prediction is generated with the following link function $\varphi : \mathbb{R}^{K+1} \rightarrow \mathcal{Y}^{\textcircled{0}}$:

$$\varphi(\mathbf{u}) = \begin{cases} \textcircled{0}, & \text{argmax}_{y \in \mathcal{Y} \cup \{K+1\}} \mathbf{u}_y(\mathbf{x}) = K + 1, \\ \text{argmax}_{y \in \mathcal{Y} \cup \{K+1\}} \mathbf{u}_y(\mathbf{x}), & \text{else.} \end{cases}$$

172 With a properly chosen surrogate Φ , the minimization of $\tilde{R}_\Phi(\mathbf{g})$ can lead to that of $\tilde{R}_{01}(\varphi(\mathbf{g}))$, which
 173 indicates the minimization of $R_{01c}(\varphi\mathbf{g})$ according to Lemmas 1 and 1. The theory of how to find
 174 such surrogates has been thoroughly studied in the field of the classification-calibration of multi-class
 175 surrogates [63, 53, 46].

176 However, we do not have direct access toward $\mathcal{D}_c^{\textcircled{0}}$ though it is closely related to the available data
 177 distribution \mathcal{D} . In this section, we propose a family of surrogate losses based on the conclusions
 178 in the previous section, which allows the use of any multi-class classification surrogates. With this
 179 formulation of surrogates, we can recover the classification risk of $\tilde{R}_\Phi(\mathbf{g})$ without access to $\mathcal{D}_c^{\textcircled{0}}$ by
 180 taking its expectation in \mathcal{D} . Based on the loss formulation, we also provide the estimation error bound
 181 to show the validity of ERM.

182 4.1 Formulation of Surrogates

183 Here, we begin with the definition of a family of surrogates for the zero-one- c loss, and then show how
 184 it can relate \mathcal{D} and $\mathcal{D}_c^{\textcircled{0}}$. With any multi-class classification loss Φ , we have the following formulation
 185 of surrogate losses for ℓ_{01c} :

186 **Definition 5.** Given a pre-defined rejection cost c , we have the following formulation of surrogate
 187 $L_c^\Phi : \mathbb{R}^{K+1} \times \mathcal{Y} \rightarrow \mathbb{R}_+$ for CwR:

$$L_c^\Phi(\mathbf{u}, y) = \Phi(\mathbf{u}, y) + (1 - c)\Phi(\mathbf{u}, K + 1), \quad (4)$$

188 where $\Phi : \mathbb{R}^{K+1} \times \mathcal{Y} \cup \{K + 1\} \rightarrow \mathbb{R}_+$ and $\mathbf{u} \in \mathbb{R}^{K+1}$.

189 The proposed surrogate loss is the linear combination of a $(K + 1)$ dimensional multi-class classifica-
 190 tion loss with coefficient determined by the predefined cost c . It can also be learned from Appendix
 191 A.2 of Charoenphakdee et al. [8] that when Φ is the softmax cross entropy loss, (4) is equivalent to
 192 Mozannar and Sontag [40]. The following theorem reveals the connection between $\tilde{R}_\Phi(\mathbf{g})$ and the
 193 expectation of L_c^Φ on \mathcal{D} :

194 **Theorem 2.** For any $\mathbf{g} : \mathcal{X} \rightarrow \mathbb{R}^{K+1}$ and $R_{L_c^\Phi}(\mathbf{g}) = \mathbb{E}_{p(\mathbf{x}, y)}[L_c^\Phi(\mathbf{g}(\mathbf{x}), y)]$:

$$R_{L_c^\Phi}(\mathbf{g}) = (2 - c)\tilde{R}_\Phi(\mathbf{g}).$$

195 The proof is provided in Appendix B. From Theorem 2, we can obtain the risk $\tilde{R}_\Phi(\mathbf{g})$ without access
 196 to $\mathcal{D}_c^{\textcircled{0}}$ with the use of the proposed surrogate (4). Following the common practice, we can finally
 197 conduct ERM [54] that minimizes the unbiased estimator of $R_{L_c^\Phi}(\mathbf{g})$, which is also that of $\tilde{R}_\Phi(\mathbf{g})$
 198 according to Theorem 2:

$$\hat{R}_{L_c^\Phi}(\mathbf{g}) = \frac{1}{n} \sum_{i=1}^n L_c^\Phi(\mathbf{g}(\mathbf{x}_i), y_i) \quad (5)$$

199 After minimizing $\hat{R}_{L_c^\Phi}(\mathbf{g})$ and obtaining the empirically optimal $\hat{\mathbf{g}}$, we can use it for predicting with
 200 the link function $\varphi \circ \hat{\mathbf{g}}$, where $\varphi \circ \hat{\mathbf{g}}(\mathbf{x}) = \varphi(\hat{\mathbf{g}}(\mathbf{x}))$.

201 According to the unbiasedness of (3), it is promising that the induced prediction rule $\varphi \circ \hat{g}$ can
 202 approximate Chow’s rule (Definition 1). To quantify such approximation, there remains two questions:
 203 what is the relation between the minimization of empirical risk $\hat{R}_{L_c^\Phi}(\mathbf{g})$ and $R_{L_c^\Phi}(\mathbf{g})$, and whether
 204 the minimization of $R_{L_c^\Phi}(\mathbf{g})$ yields that of $R_{01c}(\varphi \circ \mathbf{g})$. We will answer the two problems in Section
 205 4.2 and Section 5, respectively.

206 4.2 Estimation Error Bound

207 In Section 4.1, we proposed a family of surrogates that can recover the surrogate risk on \mathcal{D}_c^\otimes with
 208 only \mathcal{D} and provided an ERM framework to learn the empirically optimal \hat{g} . Here we further justify
 209 the use of ERM by showing that the minimization of $\hat{R}_{L_c^\Phi}$ can also result in that of $R_{L_c^\Phi}$ with the
 210 following estimation error bound.

211 **Theorem 3.** For any $\delta \in (0, 1)$, suppose the model class of g_y is \mathcal{G}_y and $\mathbf{g} \in \mathcal{G}$, where $\mathcal{G}_y \subset \mathcal{X} \rightarrow \mathbb{R}$
 212 and $\mathcal{G} \subset \mathcal{X} \rightarrow \mathbb{R}^{K+1}$ is composed of $\{\mathcal{G}_y\}_{y=1}^{K+1}$. $\Phi(\cdot, y)$ is ρ -Lipschitz continuous and is bounded
 213 by $C_\Phi > 0$. Assume that the identifiable condition holds, i.e., $\min_{\mathbf{g} \in \mathcal{G}} R_{L_c^\Phi}(\mathbf{g}) = R_{L_c^\Phi}^*$, then the
 214 following inequality holds with probability at least $1 - \delta$:

$$R_{L_c^\Phi}(\hat{\mathbf{g}}) - R_{L_c^\Phi}^* \leq 4\sqrt{2}(2-c)\rho \sum_{y=1}^{K+1} \mathfrak{R}_n(\mathcal{G}_y) + (2-c)C_\Phi \sqrt{\frac{2 \log 2/\delta}{n}}, \quad (6)$$

215 where $\mathfrak{R}_n(\mathcal{G}_y)$ is the Rademacher complexity [5] w.r.t. \mathcal{G}_y on the distribution with density $p(\mathbf{x})$ that
 216 often decays in the rate of $O(\frac{1}{\sqrt{n}})$.

217 We prove this conclusion in Appendix C. From the theorem above, we can learn that with the
 218 identifiable condition which is a common assumption with the use of complex models [4, 26, 33],
 219 $R_{L_c^\Phi}(\hat{\mathbf{g}})$ converges to $R_{L_c^\Phi}^*$ in $O_p(1/\sqrt{n})$, which is the optimal parametric convergence rate without
 220 additional assumptions [37]. According to Theorem 2, it is straightforward that $\tilde{R}_\Phi(\mathbf{g}) \xrightarrow{P} \tilde{R}_\Phi^*$ also
 221 holds. Nevertheless, the relation between the minimization of surrogate risk $\tilde{R}_\Phi(\mathbf{g})$ and that of the
 222 target risk $\tilde{R}_{01}(\varphi \circ \mathbf{g})$ is still unknown. According to Lemma 1, the minimization of $\tilde{R}_{01}(\varphi \circ \mathbf{g})$ is
 223 equivalent to zero-one- c risk minimization, which is the goal of CwR. We answer this question in the
 224 next section by giving a necessary and sufficient condition for the ℓ_{01c} -consistency for L_c^Φ .

225 5 Theoretical Analysis

226 In this section, we first point out the necessary and sufficient condition for L_c^Φ to be ℓ_{01c} -calibrated.
 227 Then we further specify the regret transfer bounds for a family of CPE-free surrogates [64], which
 228 has not been provided with theoretical analysis before.

229 5.1 Necessary and Sufficient Condition for ℓ_{01c} -Consistency

230 Given the loss formulation (4), a natural idea is to construct surrogate L_c^Φ with commonly used
 231 multi-class loss functions. However, the ℓ_{01c} -consistency of such surrogates still remains unchecked.
 232 Here, we show that we can borrow the calibration analyses of multi-class surrogates and set Φ to any
 233 $(K+1)$ -class ℓ_{01} -calibrated surrogates according to the following necessary and sufficient condition:

234 **Theorem 4.** L_c^Φ is ℓ_{01c} -consistent for any $c \in [0, 1]$ if and only if Φ is an ℓ_{01} -calibrated surrogate
 235 loss.

236 The complete proof is shown in Appendix D and here we provide its sketch. The equivalence between
 237 CwR and multi-class classification on \mathcal{D}_c shown in Lemmas 1, 1, and Theorem 2 directly yields the
 238 sufficiency of this condition. Though the equivalent classification problem is limited on \mathcal{D}_c , $\tilde{p}(y|\mathbf{x})$
 239 can be any valid class-posterior probabilities due to the arbitrariness of c and thus the calibration of Φ
 240 is necessary.

241 As a result, we can use any Φ in an off-the-shelf manner, i.e., to the consistency of different L_c^Φ , we
 242 only have to check if Φ is ℓ_{01} -calibrated, which has been studied thoroughly [7, 53, 46], instead of

243 tedious case-based discussions. Furthermore, there is also no need for the consideration of any other
 244 potential Φ since ℓ_{01} -calibration is also necessary.

245 5.2 Calibration Result for Generalized Cross Entropy Loss

246 Given the necessary and sufficient condition for ℓ_{01c} -consistency, we can construct L_c^Φ with any
 247 ℓ_{01} -calibrated surrogates. However, it has been shown in Charoenphakdee et al. [8] that it can lead
 248 to a model that rejects more data than necessary if the cross entropy (CE) loss is used as Φ , which
 249 is a popular choice as a surrogate. Another common surrogate is the mean absolute error (MAE).
 250 Though it can avoid CPE and only focus on the crucial class with the maximum posterior possibility,
 251 it usually takes more training epochs before convergence [64], which can be costly in practical use.

252 Here, we consider the *generalized cross entropy* (GCE) loss [64] that can take the advantages of the
 253 CE loss and MAE, which is defined as below:

Definition 6. (Generalized cross entropy losses) For any $\gamma \in (0, 1]$, the GCE loss is defined as below:

$$\Phi_\gamma(\mathbf{g}(\mathbf{x}), y) = (1 - S(\mathbf{g})_y^\gamma) / \gamma,$$

254 where $S(\cdot)$ is the softmax-transformation.

255 It can be seen that the loss formulation is equivalent to MAE if $\gamma = 1$ and it is also reported in Zhang
 256 and Sabuncu [64] that the GCE loss can approximate the CE loss if $\gamma \rightarrow 0$. Though the GCE loss has
 257 proved to be effective in practical use, to the best of our knowledge, its calibration results remain
 258 unknown, and thus it is unsafe directly combining it with L_c^Φ .

259 **Theorem 5.** The GCE loss Φ^γ is ℓ_{01} -calibrated for any $\gamma \in (0, 1]$. For the optimal model
 260 \mathbf{g}^* , $S(\mathbf{g}^*)_y = \eta_y^{\frac{1}{1-\gamma}} / \sum_{y'=1}^K \eta_{y'}^{\frac{1}{1-\gamma}}$ for all the $\mathbf{x} \in \mathcal{X}$ almost surely if $\gamma \in (0, 1)$. If $\gamma = 1$,
 261 $S(\mathbf{g}^*)_{\arg\max_y \eta_y} = 1$.

262 The proof can be found in Appendix E. After verifying the calibration result of the GCE loss,
 263 we can combine it with the loss formulation L_c^Φ and obtain an ℓ_{01c} -consistent surrogate. We will
 264 experimentally demonstrate its effectiveness in the next section.

265 6 Experiments

266 In this section, we provide the experiment results of CwR with deep models, which are evaluated by
 267 the zero-one- c loss following the common practice [43, 8]. We also show the misclassification rate
 268 of the accepted data and the ratio of the rejected data. Details of the setup and the experiments for
 269 instance-dependent cost can be found in Appendix F and G, respectively.

270 **Datasets and Models.** In the experiments, we evaluate the proposed methods and baselines on three
 271 widely-used benchmarks Fashion-MNIST [58], SVHN [42], CIFAR-10 [30] with cost c selected from
 272 $\{0.05, 0.06, 0.07, 0.08, 0.09, 0.10\}$ for Fashion-MNIST and $\{0.05, 0.10, 0.15, 0.20, 0.25, 0.30\}$ the
 273 other two. We conduct data augmentation for CIFAR-10 and use the original datasets of Fashion-
 274 MNIST and SVHN in the experiments. For Fashion-MNIST, we use a CNN defined in Charoen-
 275 phakdee et al. [8], and ResNet-18 and ResNet-34 [25] are used for SVHN and CIFAR-10, respectively.

276 **Baselines.** We compare our method with state-of-the-art methods in CwR, including confidence-
 277 based cross entropy loss (CE) [43], learning to defer (DEFER) [40], and cost-sensitive learning-based
 278 method with sigmoid loss (CS) [8], in which DEFER is a special case of our method that use cross
 279 entropy loss as Φ . For CE, we also conduct the temperature scaling [24] to alleviate overconfidence .
 280 For the proposed method, we use GCE with default parameter $\gamma = 0.7$ as suggested in Zhang and
 281 Sabuncu [64] and pairwise-sigmoid (Sigmoid) loss [63] to construct the surrogate L_c^Φ .

282 We implemented all the methods by Pytorch [44], and conducted all the experiments on NVIDIA
 283 GeForce 3090 GPUs.

Table 2: The mean and standard error of the zero-one- c losses (rescaled to 0-100), rejection ratio, and missclassification rates of the accepted data for 5 trails. The best and comparable methods based on the paired t-test at the significance level 5% are highlighted in boldface.

Method	Cost	CE			CS			DEFER			GCE			Sigmoid		
		01c	Rej	01	01c	Rej	01	01c	Rej	01	01c	Rej	01	01c	Rej	01
FMNIST	0.05	2.30 (0.07)	25.17 (3.17)	1.39 (0.11)	2.93 (0.25)	34.95 (1.94)	1.81 (0.48)	3.79 (0.28)	50.461 (2.51)	2.58 (0.46)	3.22 (0.07)	50.47 (2.49)	1.39 (0.30)	2.23 (0.01)	30.98 (0.62)	0.99 (0.05)
	0.06	2.58 (0.07)	22.92 (1.45)	1.56 (0.09)	3.37 (0.15)	33.13 (1.27)	2.07 (0.28)	4.63 (0.10)	56.45 (3.69)	2.84 (0.42)	3.78 (0.17)	50.46 (1.24)	1.53 (0.30)	2.62 (0.08)	26.76 (3.37)	1.37 (0.21)
	0.07	2.73 (0.14)	21.17 (2.23)	1.58 (0.31)	3.45 (0.17)	35.77 (2.62)	1.47 (0.04)	5.18 (0.47)	56.46 (6.85)	2.86 (0.41)	4.23 (0.21)	48.05 (5.24)	1.66 (0.25)	2.94 (0.07)	29.87 (0.85)	1.21 (0.17)
	0.08	3.12 (0.11)	20.71 (1.68)	1.85 (0.07)	4.13 (0.36)	33.68 (0.32)	2.17 (0.52)	5.86 (0.30)	54.08 (3.47)	3.36 (0.29)	4.50 (0.06)	45.66 (2.36)	1.55 (0.25)	3.14 (0.17)	26.10 (0.23)	1.43 (0.25)
	0.09	3.55 (0.21)	23.64 (1.82)	1.86 (0.18)	4.20 (0.21)	31.90 (1.74)	1.96 (0.15)	6.31 (0.40)	54.62 (4.33)	3.09 (0.49)	4.95 (0.06)	44.05 (1.74)	1.77 (0.23)	3.50 (0.05)	23.71 (.28)	1.79 (0.18)
	0.10	3.59 (0.16)	18.32 (1.56)	2.15 (0.32)	4.45 (0.20)	28.96 (0.13)	2.18 (0.41)	6.72 (0.07)	52.69 (0.74)	3.08 (0.18)	5.06 (0.23)	39.01 (4.87)	1.89 (0.26)	3.73 (0.05)	23.96 (1.90)	1.76 (0.20)
SVHN	0.05	3.33 (0.14)	14.37 (0.94)	3.05 (0.13)	4.42 (0.13)	12.81 (0.14)	4.33 (0.12)	4.19 (0.29)	33.05 (1.59)	3.80 (0.37)	2.68 (0.17)	19.79 (0.72)	2.10 (0.24)	2.70 (0.14)	29.56 (1.16)	1.73 (0.17)
	0.10	4.66 (0.20)	10.91 (0.57)	4.01 (0.21)	4.48 (0.14)	12.85 (0.42)	3.67 (0.11)	5.55 (0.56)	30.72 (2.64)	3.58 (0.52)	4.13 (0.11)	14.83 (0.54)	3.10 (0.10)	4.13 (0.39)	19.16 (1.94)	2.74 (0.43)
	0.15	5.40 (0.09)	8.52 (0.15)	4.50 (0.07)	5.14 (0.10)	13.21 (0.62)	3.64 (0.19)	6.37 (0.21)	21.19 (0.94)	4.05 (0.25)	4.66 (0.06)	11.47 (0.41)	3.31 (0.07)	4.83 (0.44)	18.38 (1.37)	2.54 (0.61)
	0.20	6.16 (0.13)	7.74 (0.26)	4.99 (0.09)	5.51 (0.20)	12.78 (1.03)	3.19 (0.24)	5.99 (0.17)	12.33 (0.51)	4.02 (0.16)	5.44 (0.04)	10.02 (0.25)	3.82 (0.03)	6.39 (0.45)	15.86 (0.72)	3.82 (0.48)
	0.25	7.08 (0.32)	6.51 (1.06)	5.83 (0.36)	6.77 (0.16)	12.96 (0.97)	4.06 (0.18)	6.69 (0.16)	9.18 (0.35)	4.33 (0.16)	5.75 (0.14)	8.64 (0.20)	3.93 (0.12)	6.74 (0.13)	13.79 (0.33)	3.82 (0.14)
	0.30	7.12 (0.16)	5.31 (0.36)	5.83 (0.18)	7.26 (0.33)	13.21 (1.20)	3.80 (0.41)	7.07 (0.31)	12.35 (2.31)	4.55 (0.34)	6.30 (0.09)	8.72 (0.11)	4.04 (0.09)	7.69 (0.22)	10.79 (0.76)	5.00 (0.13)
CIFAR-10	0.05	4.43 (0.23)	29.93 (1.85)	4.18 (0.33)	6.59 (0.27)	20.20 (0.51)	7.00 (0.35)	4.62 (0.47)	44.97 (5.24)	4.30 (0.88)	3.80 (0.20)	34.52 (2.77)	3.16 (0.35)	3.67 (0.03)	42.69 (8.74)	2.63 (0.49)
	0.10	7.13 (0.11)	21.13 (0.81)	6.35 (0.18)	7.68 (0.32)	20.31 (0.66)	7.08 (0.42)	6.56 (0.26)	26.21 (1.12)	5.34 (0.39)	5.84 (0.12)	25.47 (0.98)	4.41 (0.15)	6.11 (0.13)	31.66 (2.17)	4.30 (0.30)
	0.15	9.03 (0.32)	7.76 (0.39)	7.74 (0.37)	8.35 (0.29)	21.83 (0.92)	6.49 (0.45)	8.39 (0.19)	20.39 (1.59)	6.69 (0.35)	7.56 (0.14)	20.43 (0.60)	5.65 (0.23)	8.18 (0.10)	23.39 (0.82)	6.10 (0.18)
	0.20	10.45 (0.29)	14.53 (0.47)	8.82 (0.38)	9.32 (0.21)	21.86 (0.46)	6.33 (0.33)	9.65 (0.14)	17.16 (1.04)	7.50 (0.11)	9.09 (0.14)	18.45 (1.93)	6.62 (0.42)	9.69 (0.15)	19.54 (1.55)	7.20 (0.07)
	0.25	11.64 (0.26)	11.20 (0.30)	9.96 (0.32)	10.46 (0.24)	22.02 (0.40)	6.35 (0.35)	10.85 (0.08)	14.22 (1.35)	8.50 (0.30)	10.31 (0.23)	15.39 (1.47)	7.64 (0.38)	10.96 (0.11)	14.99 (1.71)	8.48 (0.40)
	0.30	12.20 (0.18)	10.02 (0.53)	10.89 (0.15)	11.43 (0.23)	22.23 (0.81)	6.13 (0.24)	11.90 (0.17)	11.48 (0.75)	9.55 (0.31)	11.23 (0.16)	12.52 (0.122)	8.55 (0.14)	12.14 (0.12)	11.08 (0.60)	9.91 (0.25)

284 **Experimental Results.** As can be seen from the experimental results reported in Table 2, our
285 proposed method (i.e., either GCE or Sigmoid) significantly outperforms other compared methods
286 in most cases. Obviously, for all the datasets and cost c , our GCE method outperforms the baseline
287 DEFER method, which indicates that CwR cannot be simply solved by the methods used for learning
288 to defer. It can be also seen that confidence-based CE is only comparable to the proposed method
289 on FMNIST with a simple CNN. When complex models are used, the effect of overconfidence is
290 inevitable even with the use of temperature scaling, which can be induced from the fact that CE often
291 rejects less data than GCE on SVHN and CIFAR-10. Though CS is comparable to GCE on CIFAR-10
292 when the rejection cost is high, its performance degrades drastically when the classification cost
293 decreases, which shows that it is not the best choice in highly error-critical tasks. When ResNet-18
294 and ResNet-34 are used on SVHN and CIFAR-10 respectively, our GCE method outperforms or
295 is comparable to all the baselines, which shows that GCE is more stable on complex models. Our
296 proposed Sigmoid method performs better than most baselines and is comparable to CE with the
297 use of a simple CNN model, which aligns with the existing observations that pairwise losses are
298 often effective with simple models [55, 14]. These results show that our method can benefit from the
299 flexibility of the choices of loss functions.

300 7 Conclusion

301 In this paper, we studied the problem of classification with rejection, which can refrain from making
302 a prediction to avoid critical misclassification. We derived a novel formulation for CwR that can
303 be equipped with arbitrary loss functions while maintaining the theoretical guarantees, making
304 them highly adaptive to the dataset in practical use. First, we showed the equivalence between
305 K -class CwR and a $(K+1)$ -class classification problem, and proposed an empirical risk minimization
306 formulation to solve this problem with an estimation error bound. Then, we pointed out a necessary
307 and sufficient condition for the learning consistency of the surrogates constructed on our proposed
308 formulation equipped with any classification-calibrated multi-class losses. Finally, experimental
309 results demonstrated the effectiveness of our proposed method.

References

- [1] Pranjal Awasthi, Natalie Frank, Anqi Mao, Mehryar Mohri, and Yutao Zhong. Calibration and consistency of adversarial surrogate losses. *CoRR*, abs/2104.09658, 2021. URL <https://arxiv.org/abs/2104.09658>.
- [2] Pranjal Awasthi, Anqi Mao, Mehryar Mohri, and Yutao Zhong. A finer calibration analysis for adversarial robustness. *CoRR*, abs/2105.01550, 2021. URL <https://arxiv.org/abs/2105.01550>.
- [3] Han Bao and Masashi Sugiyama. Calibrated surrogate maximization of linear-fractional utility in binary classification. In *AISTATS*, volume 108 of *Proceedings of Machine Learning Research*, pages 2337–2347. PMLR, 2020.
- [4] Han Bao, Clayton Scott, and Masashi Sugiyama. Calibrated surrogate losses for adversarially robust classification. In *COLT*, volume 125 of *Proceedings of Machine Learning Research*, pages 408–451. PMLR, 2020.
- [5] Peter L. Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *J. Mach. Learn. Res.*, 3:463–482, 2002. URL <http://jmlr.org/papers/v3/bartlett02a.html>.
- [6] Peter L. Bartlett and Marten H. Wegkamp. Classification with a reject option using a hinge loss. *J. Mach. Learn. Res.*, 9:1823–1840, 2008.
- [7] Peter L. Bartlett, Michael I. Jordan, and Jon D. McAuliffe. Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473):138–156, 2006.
- [8] Nontawat Charoenphakdee, Zhenghang Cui, Yivan Zhang, and Masashi Sugiyama. Classification with rejection based on cost-sensitive classification. In *ICML*, volume 139 of *Proceedings of Machine Learning Research*, pages 1507–1517. PMLR, 2021.
- [9] Nontawat Charoenphakdee, Jayakorn Vongkulbhisal, Nuttapong Chairatanakul, and Masashi Sugiyama. On focal loss for class-posterior probability estimation: A theoretical perspective. In *CVPR*, pages 5202–5211, 2021.
- [10] C. Chow. On optimum recognition error and reject tradeoff. *IEEE Transactions on Information Theory*, 16(1):41–46, 1970. doi: 10.1109/TIT.1970.1054406.
- [11] Scott Clayton. Calibrated asymmetric surrogate losses. *Electronic Journal of Stats*, 6:238–238, 2012.
- [12] Corinna Cortes, Giulia DeSalvo, and Mehryar Mohri. Boosting with abstention. In *NeurIPS*, pages 1660–1668, 2016.
- [13] Corinna Cortes, Giulia DeSalvo, and Mehryar Mohri. Learning with rejection. In *ALT*, volume 9925, pages 67–82, 2016.
- [14] Ürün Dogan, Tobias Glasmachers, and Christian Igel. A unified view on multi-class support vector classification. *J. Mach. Learn. Res.*, 17:45:1–45:32, 2016.
- [15] Ran El-Yaniv and Yair Wiener. On the foundations of noise-free selective classification. *J. Mach. Learn. Res.*, 11:1605–1641, 2010.
- [16] Charles Elkan. The foundations of cost-sensitive learning. In *IJCAI*, pages 973–978. Morgan Kaufmann, 2001.
- [17] Vitaly Feldman, Venkatesan Guruswami, Prasad Raghavendra, and Yi Wu. Agnostic learning of monomials by halfspaces is hard. *SIAM J. Comput.*, 41(6):1558–1590, 2012.

- 352 [18] Jessica Finocchiaro, Rafael M. Frongillo, and Bo Waggoner. An embedding
353 framework for consistent polyhedral surrogates. In *NeurIPS*, pages 10780–10790,
354 2019. URL [https://proceedings.neurips.cc/paper/2019/hash/
355 9ec51f6eb240fb631a35864e13737bca-Abstract.html](https://proceedings.neurips.cc/paper/2019/hash/9ec51f6eb240fb631a35864e13737bca-Abstract.html).
- 356 [19] Wei Gao and Zhi-Hua Zhou. On the consistency of multi-label learning. *Artif. Intell.*, 199-200:
357 22–44, 2013.
- 358 [20] Wei Gao and Zhi-Hua Zhou. On the consistency of AUC pairwise optimization. In *IJCAI*, pages
359 939–945. AAAI Press, 2015.
- 360 [21] Yonatan Geifman and Ran El-Yaniv. Selective classification for deep neural networks. In
361 *NeurIPS*, pages 4878–4887, 2017.
- 362 [22] Gustavo L. Gilardoni. On pinsker’s and vajda’s type inequalities for csiszár’s f-divergences.
363 *IEEE Trans. Inf. Theory*, 56(11):5377–5386, 2010. doi: 10.1109/TIT.2010.2068710.
- 364 [23] Yves Grandvalet, Alain Rakotomamonjy, Joseph Keshet, and Stéphane Canu. Support vector
365 machines with a reject option. In *NeurIPS*, pages 537–544, 2008.
- 366 [24] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. On calibration of modern neural
367 networks. In *ICML*, volume 70 of *Proceedings of Machine Learning Research*, pages 1321–1330.
368 PMLR, 2017.
- 369 [25] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image
370 recognition. In *CVPR*, pages 463–469, 2016.
- 371 [26] Yen-Chang Hsu, Zhaoyang Lv, Joel Schlosser, Phillip Odom, and Zsolt Kira. Multi-class
372 classification without multi-class labels. In *ICLR*, 2019.
- 373 [27] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training
374 by reducing internal covariate shift. In *ICML*, volume 37, pages 448–456. JMLR.org, 2015.
- 375 [28] D. E. Knuth. Two notes on notation. *American Mathematical Monthly*, 99(5):403–422, 1992.
- 376 [29] Oluwasanmi Koyejo, Nagarajan Natarajan, Pradeep Ravikumar, and Inderjit S. Dhillon. Consis-
377 tent multilabel classification. In *NeurIPS*, pages 3321–3329, 2015.
- 378 [30] Alex Krizhevsky. Learning multiple layers of features from tiny images. *University of Toronto*,
379 05 2012.
- 380 [31] Maksim Lapin, Matthias Hein, and Bernt Schiele. Analysis and optimization of loss functions
381 for multiclass, top-k, and multilabel classification. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40
382 (7):1533–1554, 2018.
- 383 [32] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. Focal loss for
384 dense object detection. In *ICCV*, pages 2999–3007, 2017.
- 385 [33] Nan Lu, Shida Lei, Gang Niu, Issei Sato, and Masashi Sugiyama. Binary classification from
386 multiple unlabeled datasets via surrogate set classification. In *ICML*, 2021.
- 387 [34] Naresh Manwani, Kalpit Desai, Sanand Sasidharan, and Ramasubramanian Sundararajan.
388 Double ramp loss based reject option classifier. In *PAKDD*, volume 9077, pages 151–163, 2015.
- 389 [35] Andreas Maurer. A vector-contraction inequality for rademacher complexities. In Ronald
390 Ortner, Hans Ulrich Simon, and Sandra Zilles, editors, *ALT*, pages 3–17, 2016.
- 391 [36] C. Mcdiarmid. *On the method of bounded differences*. Surveys in Combinatorics, 1989.

- 392 [37] Shahar Mendelson. Lower bounds for the empirical minimization algorithm. *IEEE Trans. Inf.*
393 *Theory*, 54(8):3797–3803, 2008. doi: 10.1109/TIT.2008.926323. URL [https://doi.org/](https://doi.org/10.1109/TIT.2008.926323)
394 [10.1109/TIT.2008.926323](https://doi.org/10.1109/TIT.2008.926323).
- 395 [38] Aditya Krishna Menon and Robert C. Williamson. In *COLT*, volume 35 of *JMLR Workshop and*
396 *Conference Proceedings*, pages 68–106. JMLR.org, 2014. URL [http://proceedings.](http://proceedings.mlr.press/v35/menon14.html)
397 [mlr.press/v35/menon14.html](http://proceedings.mlr.press/v35/menon14.html).
- 398 [39] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of Machine Learning.*
399 Adaptive computation and machine learning. MIT Press, 2012. ISBN 978-0-262-01825-8.
- 400 [40] Hussein Mozannar and David A. Sontag. Consistent estimators for learning to defer to an expert.
401 In *ICML*, volume 119 of *Proceedings of Machine Learning Research*, pages 7076–7087. PMLR,
402 2020.
- 403 [41] Hussein Mozannar, Arvind Satyanarayan, and David A. Sontag. Teaching humans when to
404 defer to a classifier via exemplars. *CoRR*, abs/2111.11297, 2021. URL [https://arxiv.](https://arxiv.org/abs/2111.11297)
405 [org/abs/2111.11297](https://arxiv.org/abs/2111.11297).
- 406 [42] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Ng. Reading
407 digits in natural images with unsupervised feature learning. *NeurIPS Workshop on Deep*
408 *Learning and Unsupervised Feature Learning*, 1 2011.
- 409 [43] Chenri Ni, Nontawat Charoenphakdee, Junya Honda, and Masashi Sugiyama. On the calibration
410 of multiclass classification with rejection. In *NeurIPS*, pages 2582–2592, 2019.
- 411 [44] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan,
412 Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas
413 Köpf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy,
414 Benoît Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style,
415 high-performance deep learning library. In *NeurIPS*, pages 8024–8035, 2019.
- 416 [45] Bernardo Ávila Pires and Csaba Szepesvári. Multiclass classification calibration functions.
417 *CoRR*, abs/1609.06385, 2016. URL <http://arxiv.org/abs/1609.06385>.
- 418 [46] Harish G. Ramaswamy and Shivani Agarwal. Convex calibration dimension for multiclass loss
419 matrices. *J. Mach. Learn. Res.*, 17:14:1–14:45, 2016. URL [http://jmlr.org/papers/](http://jmlr.org/papers/v17/14-316.html)
420 [v17/14-316.html](http://jmlr.org/papers/v17/14-316.html).
- 421 [47] Harish G Ramaswamy, Ambuj Tewari, and Shivani Agarwal. Consistent algorithms for mul-
422 ticlass classification with an abstain option. *Electronic Journal of Statistics*, 12(1):530–554,
423 2018.
- 424 [48] Alexey E. Rastegin. Bounds of the pinsker and fannes types on the tsallis relative entropy.
425 *Mathematical Physics Analysis & Geometry*, 16(3):213–228, 2013.
- 426 [49] Mark D. Reid and Robert C. Williamson. Composite binary losses. *J. Mach. Learn. Res.*, 11:
427 2387–2422, 2010.
- 428 [50] Clayton Scott. Surrogate losses and regret bounds for cost-sensitive classification with example-
429 dependent costs. In Lise Getoor and Tobias Scheffer, editors, *ICML*, pages 153–160. Omnipress,
430 2011. URL https://icml.cc/2011/papers/138_icmlpaper.pdf.
- 431 [51] Song-Qing Shen, Bin-Bin Yang, and Wei Gao. AUC optimization with a reject option. In *AAAI*,
432 pages 5684–5691, 2020.
- 433 [52] Ingo Steinwart. How to compare different loss functions and their risks. *Constructive Approxi-*
434 *mation*, 26:225–287, 2007.

- 435 [53] Ambuj Tewari and Peter L. Bartlett. On the consistency of multiclass classification methods. *J.*
436 *Mach. Learn. Res.*, 8:1007–1025, 2007.
- 437 [54] Vladimir Vapnik. *Statistical learning theory*. Wiley, 1998. ISBN 978-0-471-03003-4.
- 438 [55] Yutong Wang and Clayton Scott. Weston-watkins hinge loss and ordered partitions. In *NeurIPS*,
439 2020.
- 440 [56] Robert C. Williamson, Elodie Vernet, and Mark D. Reid. Composite multiclass losses. *J. Mach.*
441 *Learn. Res.*, 17:223:1–223:52, 2016.
- 442 [57] Guoqiang Wu, Chongxuan Li, Kun Xu, and Jun Zhu. Rethinking and reweighting the univariate
443 losses for multi-label ranking: Consistency and generalization. *CoRR*, abs/2105.05026, 2021.
444 URL <https://arxiv.org/abs/2105.05026>.
- 445 [58] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for
446 benchmarking machine learning algorithms. *CoRR*, abs/1708.07747.
- 447 [59] F. Y. Yang. Maximum lq-likelihood estimation. *Annals of Statistics*, 2010.
- 448 [60] Forest Yang and Sanmi Koyejo. On the consistency of top-k surrogate losses. In *ICML*, volume
449 119 of *Proceedings of Machine Learning Research*, pages 10727–10735. PMLR, 2020.
- 450 [61] Ming Yuan and Marten H. Wegkamp. Classification methods with reject option based on convex
451 risk minimization. *J. Mach. Learn. Res.*, 11:111–130, 2010.
- 452 [62] Mingyuan Zhang, Harish Guruprasad Ramaswamy, and Shivani Agarwal. Convex calibrated
453 surrogates for the multi-label f-measure. In *ICML*, volume 119 of *Proceedings of Machine*
454 *Learning Research*, pages 11246–11255. PMLR, 2020.
- 455 [63] Tong Zhang. Statistical analysis of some multi-category large margin classification methods.
456 *Journal of Machine Learning Research*, 5(Oct):1225–1251, 2004.
- 457 [64] Zhilu Zhang and Mert R. Sabuncu. Generalized cross entropy loss for training deep neural
458 networks with noisy labels. In *NeurIPS*, pages 8792–8802, 2018.

459 **Checklist**

- 460 1. For all authors...
- 461 (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s
462 contributions and scope? [Yes]
- 463 (b) Did you describe the limitations of your work? [Yes] See Appendix H.
- 464 (c) Did you discuss any potential negative societal impacts of your work? [Yes] See
465 Appendix H.
- 466 (d) Have you read the ethics review guidelines and ensured that your paper conforms to
467 them? [Yes]
- 468 2. If you are including theoretical results...
- 469 (a) Did you state the full set of assumptions of all theoretical results? [Yes]
- 470 (b) Did you include complete proofs of all theoretical results? [Yes] See Appendix.
- 471 3. If you ran experiments...
- 472 (a) Did you include the code, data, and instructions needed to reproduce the main ex-
473 perimental results (either in the supplemental material or as a URL)? [Yes] See the
474 supplemental material.
- 475 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they
476 were chosen)? [Yes] See Section 6 and Appendix.
- 477 (c) Did you report error bars (e.g., with respect to the random seed after running experi-
478 ments multiple times)? [Yes] See Table 2.
- 479 (d) Did you include the total amount of compute and the type of resources used (e.g., type
480 of GPUs, internal cluster, or cloud provider)? [Yes] See Section 6.
- 481 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 482 (a) If your work uses existing assets, did you cite the creators? [Yes] See Section 6.
- 483 (b) Did you mention the license of the assets? [N/A] The used datasets are open bench-
484 marks.
- 485 (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
486 Please refer to the supplemental materials.
- 487 (d) Did you discuss whether and how consent was obtained from people whose data you’re
488 using/curating? [N/A]
- 489 (e) Did you discuss whether the data you are using/curating contains personally identifiable
490 information or offensive content? [N/A]
- 491 5. If you used crowdsourcing or conducted research with human subjects...
- 492 (a) Did you include the full text of instructions given to participants and screenshots, if
493 applicable? [N/A]
- 494 (b) Did you describe any potential participant risks, with links to Institutional Review
495 Board (IRB) approvals, if applicable? [N/A]
- 496 (c) Did you include the estimated hourly wage paid to participants and the total amount
497 spent on participant compensation? [N/A]

498 **A Proof of Theorem 1 and Corollary 1**

499 We begin with the proof of Corollary 1 and show that Theorem 1 is its special case.

Proof. First of all, we prove that the Bayes optimal solution on $\tilde{p}(\mathbf{x}, \tilde{y})$ coincide with the Chow's rule of $p(\mathbf{x}, y)$ with cost c . According to the optimality condition of multi-class classification, the optimal classifier $f^*(\mathbf{x})$ on $\tilde{p}(\mathbf{x}, \tilde{y})$ should fulfill the following condition almost surely:

$$f^*(\mathbf{x}) = \operatorname{argmax}_{\tilde{y}} \tilde{p}(\tilde{y}|\mathbf{x}), \tilde{y} \in \{1, \dots, K, \textcircled{\ast}\}.$$

According to the definition of \tilde{p} , we can further rewrite it as:

$$f^*(\mathbf{x}) = \begin{cases} \textcircled{\ast}, & \max_{\tilde{y} \in \{1, \dots, K\}} \frac{p(\tilde{y}|\mathbf{x})}{2-c(\mathbf{x})} \leq \frac{1-c(\mathbf{x})}{2-c(\mathbf{x})}, \\ \operatorname{argmax}_{\tilde{y} \in \{1, \dots, K\}} \frac{p(\tilde{y}|\mathbf{x})}{2-c(\mathbf{x})}, & \text{else,} \end{cases}$$

500 which coincides with the Chow's rule. Then we have the following conclusions:

$$\begin{aligned} \bar{R}_{01}(f) &= \mathbb{E}_{\tilde{p}(\mathbf{x}, \tilde{y})}[(2-c(\mathbf{x}))\ell_{01}(f(\mathbf{x}), \tilde{y})] \\ &= \int_{\mathbf{x}} \sum_{\tilde{y}=1}^K (2-c(\mathbf{x}))\ell_{01}(f(\mathbf{x}), \tilde{y}) \frac{p(\mathbf{x}, y)}{2-c(\mathbf{x})} d\mathbf{x} + \int_{\mathbf{x}} (2-c(\mathbf{x}))\ell_{01}(f(\mathbf{x}), K+1) \frac{p(\mathbf{x})}{2-c(\mathbf{x})} d\mathbf{x} \\ &= \int_{\mathbf{x}} \sum_{\tilde{y}=1}^K (2-c(\mathbf{x}))\ell_{01}(f(\mathbf{x}), \tilde{y}) \frac{p(\mathbf{x}, y)}{2-c(\mathbf{x})} d\mathbf{x} + \int_{\mathbf{x}} (2-c(\mathbf{x}))\ell_{01}(f(\mathbf{x}), \textcircled{\ast}) \frac{(1-c(\mathbf{x}))p(\mathbf{x})}{2-c(\mathbf{x})} d\mathbf{x} \\ &= \int_{\mathbf{x}} \sum_{\tilde{y}=1}^K \ell_{01}(f(\mathbf{x}), \tilde{y}) p(\mathbf{x}, y) d\mathbf{x} + \int_{\mathbf{x}} \ell_{01}(f(\mathbf{x}), \textcircled{\ast}) (1-c(\mathbf{x})) p(\mathbf{x}) d\mathbf{x} \end{aligned}$$

501 Suppose $C(f(\mathbf{x})) = \sum_{\tilde{y}=1}^K \ell_{01}(f(\mathbf{x}), \tilde{y}) p(y|\mathbf{x}) + \ell_{01}(f(\mathbf{x}), \textcircled{\ast}) (1-c(\mathbf{x}))$ is the inner risk and f^*
502 is the Chow's rule, we have that

- If $f^*(\mathbf{x}) = \textcircled{\ast}$ and $f(\mathbf{x}) \neq f^*(\mathbf{x})$:

$$C(f(\mathbf{x})) - C(f^*(\mathbf{x})) = 1 - c(\mathbf{x}) - p(f(\mathbf{x})|\mathbf{x}).$$

- If $f^*(\mathbf{x}) \in \{1, \dots, K\}$ and $f(\mathbf{x}) = \textcircled{\ast}$:

$$C(f(\mathbf{x})) - C(f^*(\mathbf{x})) = p(f^*(\mathbf{x})|\mathbf{x}).$$

- If $f^*(\mathbf{x}), f(\mathbf{x}) \in \{1, \dots, K\}$ and $f(\mathbf{x}) \neq f^*(\mathbf{x})$:

$$C(f(\mathbf{x})) - C(f^*(\mathbf{x})) = p(f^*(\mathbf{x})|\mathbf{x}) - p(f(\mathbf{x})|\mathbf{x}).$$

These conclusion shows that

$$C(f(\mathbf{x})) - C(f^*(\mathbf{x})) = \mathbb{E}_{p(y|\mathbf{x})}[\ell_{01c}(f(\mathbf{x}), y)] - \mathbb{E}_{p(y|\mathbf{x})}[\ell_{01c}(f^*(\mathbf{x}), y)].$$

503 We can conclude the proof by taking the expectation over $p(\mathbf{x})$ on both sides of the equation. \square

504 It can be seen that when $c(\mathbf{x})$ is constant, we can divide each side of Corollary 1 to get the proof of
505 Theorem 1.

506 **B Proof of Theorem 2**

Proof.

$$\begin{aligned} R_{L_c^\Phi}(\mathbf{g}) &= \mathbb{E}_{p(\mathbf{x}, y)}[L_c^\Phi(\mathbf{g}(\mathbf{x}), y)] \\ &= \mathbb{E}_{p(\mathbf{x}, y)}[\Phi(\mathbf{g}(\mathbf{x}), y)] + (1-c)\mathbb{E}_{p(\mathbf{x})}[\Phi(\mathbf{g}(\mathbf{x}), K+1)] \\ &= (2-c)\tilde{R}_\Phi(\mathbf{g}) \end{aligned}$$

507 \square

508 **C Proof of Theorem 3**

509 We first give the definition of Rademacher complexity:

510 **Definition 7.** (Rademacher complexity [5]) Let Z_1, \dots, Z_n be *n i.i.d.* random variables drawn from
 511 a probability distribution μ and $\mathcal{F} = \{f : Z \rightarrow \mathbb{R}\}$ be a class of measurable functions. Then the
 512 expected Rademacher complexity of function class \mathcal{F} is given as follow:

$$\mathfrak{R}_n(\mathcal{F}) = \mathbb{E}_{Z_1, \dots, Z_n \sim \mu} \mathbb{E}_{\sigma} \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n \sigma_i f(Z_i) \right], \quad (7)$$

513 where $\sigma_1, \dots, \sigma_n$ are the Rademacher variables that take the value from $\{-1, +1\}$ uniformly.

514 Then we can begin proving Theorem 3.

515 *Proof.* According to the conditions in Theorem 3, we can learn that L_c^Φ is $(2-c)\rho$ -Lipschitz
 516 continuous and is bounded by $(2-c)C_\Phi$. By applying the McDiarmid's inequality [36], it is routine
 517 [39] to show that the following inequalities holds with probability at least $1 - \frac{\delta}{2}$, respectively:

$$\begin{aligned} \sup_{g \in \mathcal{G}} \left(R_{L_c^\Phi}(g) - \hat{R}_{L_c^\Phi}(g) \right) &\leq \mathbb{E}_{\mathbf{x}_1, \dots, \mathbf{x}_n} \left[\sup_{g \in \mathcal{G}} \left(R_{L_c^\Phi}(g) - \hat{R}_{L_c^\Phi}(g) \right) \right] + (2-c)C_\Phi \sqrt{\frac{\log \frac{2}{\delta}}{2n}} \\ \sup_{g \in \mathcal{G}} \left(\hat{R}_{L_c^\Phi}(g) - R_{L_c^\Phi}(g) \right) &\leq \mathbb{E}_{\mathbf{x}_1, \dots, \mathbf{x}_n} \left[\sup_{g \in \mathcal{G}} \left(\hat{R}_{L_c^\Phi}(g) - R_{L_c^\Phi}(g) \right) \right] + (2-c)C_\Phi \sqrt{\frac{\log \frac{2}{\delta}}{2n}} \end{aligned}$$

518 By applying Talagrand's contraction lemma [35], we can learn that:

$$\mathbb{E}_{\mathbf{x}_1, \dots, \mathbf{x}_n} \left[\sup_{g \in \mathcal{G}} \left(R_{L_c^\Phi}(g) - \hat{R}_{L_c^\Phi}(g) \right) \right] \leq \sqrt{2}(2-c)\rho \sum_{y=1}^{K+1} \mathfrak{R}_n(\mathcal{G}_y)$$

519 and this conclusion also holds for another direction. Plugging this conclusion into the former
 520 inequalities and using the union bound, we can learn:

$$\sup_{g \in \mathcal{G}} \left| R_{L_c^\Phi}(g) - \hat{R}_{L_c^\Phi}(g) \right| \leq \sqrt{2}(2-c)\rho \sum_{y=1}^{K+1} \mathfrak{R}_n(\mathcal{G}_y) + (2-c)C_\Phi \sqrt{\frac{\log \frac{2}{\delta}}{2n}}$$

521 According to the definition of empirical risk minimization and identifiable condition, we can get the
 522 following conclusion, where g^* is the optimal solution among all the measurable functions:

$$\begin{aligned} R_{L_c^\Phi}(\hat{g}) - R_{L_c^\Phi}^* &= \left(R_{L_c^\Phi}(\hat{g}) - \hat{R}_{L_c^\Phi}(\hat{g}) \right) + \left(\hat{R}_{L_c^\Phi}(\hat{g}) - \hat{R}_{L_c^\Phi}(g^*) \right) + \left(\hat{R}_{L_c^\Phi}(g^*) - R_{L_c^\Phi}^* \right) \\ &\leq \left(R_{L_c^\Phi}(\hat{g}) - \hat{R}_{L_c^\Phi}(\hat{g}) \right) + \left(\hat{R}_{L_c^\Phi}(g^*) - R_{L_c^\Phi}^* \right) \\ &\leq 2 \sup_{g \in \mathcal{G}} \left| R_{L_c^\Phi}(g) - \hat{R}_{L_c^\Phi}(g) \right| \end{aligned}$$

523 which concludes the proof. □

524 **D Proof of Theorem 4**

525 *Proof.* According to Theorem 1, Theorem 2, and Theorem 3 in Ramaswamy and Agarwal [46], we
 526 can immediately learn the sufficiency of this condition.

We complete the proof of the necessity of the calibration of Φ by contradiction. Suppose there are
 some $\mathbf{u} \in \Delta^{K+1}$ that:

$$\inf_{\mathbf{u} \in \mathbb{R}^K, \mathbf{u} \notin \arg \min_{\mathbf{u}} \mathbf{p}^T \mathbf{L}_{01}(\mathbf{u})} \mathbf{p}^T \Phi(\mathbf{u}) = \inf_{\mathbf{u} \in \mathbb{R}^K} \mathbf{p}^T \Phi(\mathbf{u}).$$

527 It is easy to learn that any re-permutation of \mathbf{u} also fulfill the equation above, and we define the
528 collection of these vectors as \mathcal{U} . Then we can construct a distribution over $\mathcal{X} \times \{1, \dots, K+1\}$
529 whose posterior possibility is $\mathbf{u}' \in \mathcal{U}$ for all \mathbf{x} , on which Φ is not ℓ_{01} -consistent. However, in our
530 scenario, we only focus on a special distribution with density $\tilde{p}(\mathbf{x}, \tilde{y})$ over $\mathcal{X} \times \{1, \dots, K+1\}$,
531 where $\tilde{p}(K+1|\mathbf{x}) = \frac{1-c}{2-c}$ and $\tilde{p}(\tilde{y}|\mathbf{x}) = p(\tilde{y}|\mathbf{x})/(2-c)$ if $\tilde{y} \neq K+1$. A natural idea is that according
532 to the particularity of \tilde{p} , there may not be overlap between \mathcal{U} and all the potential $\{\tilde{p}(\tilde{y}|\mathbf{x})\}_{\tilde{y}=1}^{K+1}$.
533 However, according to the arbitrariness of c , this idea is not true, i.e., there always exists a distribution
534 $\{p(y|\mathbf{x})\}_{y=1}^K$ and c that $\{\tilde{p}(\tilde{y}|\mathbf{x})\}_{\tilde{y}=1}^{K+1} \in \mathcal{U}$. Then we can easily define a distribution based on
535 $\{\tilde{p}(\tilde{y}|\mathbf{x})\}_{\tilde{y}=1}^{K+1}$, on which Φ is not ℓ_{01} -consistent. According to the equivalence shown in Theorem
536 1 and 2, this observation indicates that L_c^Φ is not ℓ_{01c} -consistent *w.r.t.* to this distribution, which
537 shows the necessity of the ℓ_{01} -calibration of Φ .

538

□

539 E Proof of Theorem 5

540 *Proof.* According to [59], we can directly get the formulation of the optimal solution of GCE. Based
541 on this formulation, we prove the classification-calibration of GCE constructively by giving an regret
542 transfer bound.

543 First of all, we show that the excess error of GCE loss for any \mathbf{x} is a reweighted version of the Tsallis
544 relative entropy [22, 48] in actual. Denote by $S(\mathbf{g}^*)_y = \mathbf{q}_y^*$, $S(\mathbf{g})_y = \mathbf{q}_y$ for any \mathbf{g} , and $p(y|\mathbf{x}) = \eta_y$.
545 We substitute γ with r in the proof for simplicity:

$$\begin{aligned} Ex(\mathbf{q}, \mathbf{x}) &= \sum_{y=1}^y \eta_y \frac{(1 - \mathbf{q}_y^r)}{r} - \sum_{y=1}^y \eta_y \frac{(1 - \mathbf{q}_y^{*r})}{r} \\ &= \frac{\sum_{y=1}^K \eta_y (\mathbf{q}_y^{*r} - \mathbf{q}_y^r)}{r} \\ &= \left(\sum_{y=1}^K \eta_y^{\frac{1}{1-r}} \right)^{1-r} \frac{\left(1 - \sum_{y=1}^K \mathbf{q}_y^{*(1-r)} \mathbf{q}_y^r \right)}{r} \end{aligned}$$

546 It can be seen that the second term of the last equation is the Tsallis relative entropy between discrete
547 possibilities \mathbf{q}^* and \mathbf{q} . According to the Corollary 9 of [22] and (4.13) of [48], we can lower bound
548 the excess error with the total variation distance between \mathbf{q}^* and \mathbf{q} and get a Pinsker's type inequality:

$$Ex(\mathbf{q}, \mathbf{x}) \geq \left(\sum_{y=1}^K \eta_y^{\frac{1}{1-r}} \right)^{1-r} \frac{1-r}{2} \|\mathbf{q}^* - \mathbf{q}\|_1^2$$

549 Then we have to connect the *r.h.s.* of the inequality to the excess error *w.r.t.* 0-1 loss. When
550 $\operatorname{argmax}_y \mathbf{q}_y(\mathbf{x}) \neq \operatorname{argmax}_y \eta_y$, denote by $\operatorname{argmax}_y \mathbf{q}_y(\mathbf{x}) = \mathit{pred}$ and $\operatorname{argmax}_y \eta_y = \mathit{max}$:

$$\begin{aligned} \|\mathbf{q}^* - \mathbf{q}\|_1 &= \sum_{y=1}^K |q_y^* - q_y| \\ &\geq |q_{\mathit{max}}^* - q_{\mathit{max}}| + |q_{\mathit{pred}}^* - q_{\mathit{pred}}| \\ &\geq |q_{\mathit{max}}^* - q_{\mathit{pred}}^* + q_{\mathit{pred}} - q_{\mathit{max}}| \end{aligned}$$

551 According to the formulation of the optimal solution of GCE, we can learn that $q_{max}^* \geq q_{pred}^*$. Since
 552 $\operatorname{argmax}_y q_y(\mathbf{x}) \neq \operatorname{argmax}_y \eta_y$, we can learn that $q_{pred} \geq q_{max}$. Then we can further learn that:

$$\begin{aligned}
 \|\mathbf{q}^* - \mathbf{q}\|_1 &\geq |q_{max}^* - q_{pred}^*| \\
 &= \left(\sum_{y=1}^K \eta_y^{\frac{1}{1-r}} \right)^{-1} |\eta_{max}^{\frac{1}{1-r}} - \eta_{pred}^{\frac{1}{1-r}}| \\
 &= \left(\sum_{y=1}^K \eta_y^{\frac{1}{1-r}} \right)^{-1} (\eta_{max}^{\frac{1}{1-r}} - \eta_{pred}^{\frac{1}{1-r}}) \\
 &= \left(\sum_{y=1}^K \eta_y^{\frac{1}{1-r}} \right)^{-1} (\eta_{max} * \eta_{max}^{\frac{r}{1-r}} - \eta_{pred} * \eta_{pred}^{\frac{r}{1-r}}) \\
 &\geq \left(\sum_{y=1}^K \eta_y^{\frac{1}{1-r}} \right)^{-1} (\eta_{max} * \eta_{max}^{\frac{r}{1-r}} - \eta_{pred} * \eta_{max}^{\frac{r}{1-r}}) \\
 &= \left(\sum_{y=1}^K \eta_y^{\frac{1}{1-r}} \right)^{-1} \eta_{max}^{\frac{r}{1-r}} (\eta_{max} - \eta_{pred})
 \end{aligned}$$

553 Then we can learn that:

$$\begin{aligned}
 Ex(\mathbf{q}, \mathbf{x}) &\geq \left(\sum_{y=1}^K \eta_y^{\frac{1}{1-r}} \right)^{-1-r} \eta_{max}^{\frac{2r}{1-r}} * \frac{1-r}{2} (\eta_{max} - \eta_{pred})^2 \\
 &\geq \frac{1-r}{2K^{\frac{2r}{1-r} + r + r^2}} (\eta_{max} - \eta_{pred})^2
 \end{aligned}$$

Then we have the following regret transfer bound:

$$R_{01}(\operatorname{argmax}_y \mathbf{g}_y) - R_{01}^* \leq \sqrt{C(R_G(\mathbf{g}) - R_G^*)},$$

554 where $C = \frac{2K^{\frac{2r}{1-r} + r + r^2}}{1-r}$, R_G is the expected version of GCE loss, and R_G^* and R_{01}^* are the opti-
 555 mal value of the expected version of GCE loss and 0-1 loss, respectively. From this bound, we
 556 constructively prove the classification-calibration of GCE loss with $r \in (0, 1)$. \square

557 It is noticeable that the bound does not hold for $r = 1$, e.g., the case of MAE loss, and the regret
 558 transfer bound becomes less compact when r increases. We prove the classification-calibration of
 559 MAE loss by showing its regret transfer bound.

Corollary 2. Suppose the expected version of MAE loss is $R_M(\mathbf{g})$ and its minimal value is R_M^* . Then we have:

$$R_{01}(\operatorname{argmax}_y \mathbf{g}_y) - R_{01}^* \leq K(R_M(\mathbf{g}) - R_M^*).$$

560 *Proof.* Given the formulation of the optimal solution \mathbf{q}^* of expected MAE loss in Theorem 5, for any
 561 \mathbf{x} , the excess error can be written as:

$$\begin{aligned}
 Ex(\mathbf{q}, \mathbf{x}) &= \sum_{y=1}^K \eta_y (1 - q_y) - \sum_{y=1}^K \eta_y (1 - q_y^*) \\
 &= \sum_{y=1}^K \eta_y (q_y^* - q_y) \\
 &= \eta_{max} - \sum_{y=1}^K \eta_y q_y
 \end{aligned}$$

562 When $\operatorname{argmax}_y q_y(\mathbf{x}) \neq \operatorname{argmax}_y \eta_y$:

$$\begin{aligned}
 \eta_{max} - \sum_{y=1}^K \eta_y q_y &= \eta_{max} - \eta_{pred} q_{pred} - \sum_{y \neq pred}^K \eta_y q_y \\
 &\geq \eta_{max} - \eta_{pred} q_{pred} - \eta_{max} (1 - q_{pred}) \\
 &= q_{pred} (\eta_{max} - \eta_{pred}) \\
 &\geq \frac{1}{K} (\eta_{max} - \eta_{pred}),
 \end{aligned}$$

563 which concludes the proof by taking the expectation on both sides. \square

564 Combine the conclusions above and we can conclude the proof. Though the bound for GCE becomes
 565 less tight when r increases, the MAE loss has a better regret transfer bound, which indicates that
 566 the regret transfer bound of GCE for $r \in (0, 1)$ may not be good enough. A potential reason is that
 567 [22, 48] considered the general case of Tsallis relative entropy while we only need the case that \mathbf{q} is a
 568 probability distribution. It is promising to further tighten this bound by modifying the conclusions in
 569 [22, 48] and limiting \mathbf{q} to a $K - 1$ -dimensional probability simplex.

570 F Details of the Experiment Setup

571 F.1 Detailed Information of Benchmark Datasets

572 In the experiments, we used 3 widely-used benchmark datasets. Here, we report the sources of these
 573 datasets and the way we split them.

- 574 • Fashion-MNIST [58]. It is a 10-class dataset of fashion items. Each instance is
 575 a 28*28 grayscale image. Source: [https://github.com/zalandoresearch/
 576 fashion-mnist](https://github.com/zalandoresearch/fashion-mnist).
- 577 • SVHN [42] It is a 10-class dataset for 10 different digits and each instance is a
 578 32*32*3 colored image in RGB format. Source: [http://ufldl.stanford.edu/
 579 housenumbers/](http://ufldl.stanford.edu/housenumbers/).
- 580 • CIFAR-10 [30]. It is a 10-class dataset for 10 different objects and each instance is a
 581 32*32*3 colored image in RGB format. Source: [https://www.cs.toronto.edu/
 582 ~kriz/cifar.html](https://www.cs.toronto.edu/~kriz/cifar.html).

583 For Fashion-MNIST and SVHN, we trained models on the whole training dataset. For CIFAR-10,
 584 we splitted 10% of the training dataset as the validation set and conducted random crop and flips
 585 for data augmentation. The cost c is less than 0.5 as suggested in [47] and further decreased on
 586 Fashion-MNIST since it is a less difficult dataset.

587 F.2 Detailed Information of the Models and Optimization Algorithm

588 For Fashion-MNIST, we used the model defined in [8] for the experiments. For SVHN and CIFAR-
 589 10, ResNet-18 and ResNet-34 is used, respectively. For the cost-sensitive method [8], we use
 590 batch normalization [27] at the output layer as suggested in [8] since it fails to work without this
 591 modification.

592 Adam with default momentum was used for optimization in this paper. For Fashion-MNIST, the
 593 epoch number, batch size, learning rate, and weight decay are set to 20, 256, 1e-3, and 1e-4. For
 594 SVHN, the epoch number, batch size, learning rate, and weight decay are set to 20, 1024, 1e-3, and
 595 1e-4. For CIFAR-10, the epoch number, batch size, learning rate, and weight decay are set to or
 596 selected from 200, 1024, {1e-3, 2e-3, 3e-3}, and 1e-4. For Fashion-MNIST and SVHN, we use the
 597 model after the 20th epoch for performance evaluation. For CIFAR-10, we report the performance
 598 of the model with the best performance on the validation dataset. Temperature scaling is further
 599 conducted for CE on CIFAR-10.

600 **G Details of Instance-dependent Rejection Cost**

601 In practical applications, it can be beneficial letting the rejection cost $c(\mathbf{x})$ vary among different
 602 samples. For example, when constructing a system to automatically prescribe for users, a wrong
 603 prescription can be fatal for users of advanced ages or with underlying diseases. To prevent such
 604 wrong prescriptions, the cost for this type of users can be decreased to encourage rejection. However,
 605 it is not suitable encouraging rejection for all the users, which makes the system meaningless. An
 606 acceptable choice is to increase the cost for rejection instead for users of low risk.

607 In this appendix, we expand the Theorem 2 and propose a surrogate for instance dependent cost based
 608 on Corollary 1, whose estimation error bound and calibration analysis can be derived almost symmet-
 609 rically thanks to the equivalence shown in Corollary 3. Then we further evaluate its performance on
 610 SVHN dataset.

611 **G.1 Expansion of Theorem 2**

612 Theorem 2 tells the equivalence between surrogate risk minimization of L_c^Φ on $p(\mathbf{x}, y)$ and surrogate
 613 risk minimization of Φ on $\tilde{p}(\mathbf{x}, \tilde{y})$. Here we expand it to the case of instance-dependent cost.

Given the cost function $c(\mathbf{x})$ and any function $\Phi(\cdot) : \mathbb{R}^{K+1} \times \{1, \dots, K+1\} \rightarrow \mathbb{R}^+$:

$$L_{c(\mathbf{x})}^\Phi(\mathbf{u}, y) = (\Phi(\mathbf{u}, y) + (1 - c(\mathbf{x}))\Phi(\mathbf{u}, K + 1))/(2 - c(\mathbf{x})).$$

614 Then we have the following conclusion:

Corollary 3. For any $\mathbf{g} : \mathcal{X} \rightarrow \mathbb{R}^{K+1}$ and $R_{L_{c(\mathbf{x})}^\Phi}(\mathbf{g}) = \mathbb{E}_{p(\mathbf{x}, y)}[L_{c(\mathbf{x})}^\Phi(\mathbf{g}(\mathbf{x}), y)]$:

$$R_{L_{c(\mathbf{x})}^\Phi}(\mathbf{g}) = \tilde{R}_\Phi(\mathbf{g})$$

Proof.

$$\begin{aligned} R_{L_{c(\mathbf{x})}^\Phi}(\mathbf{g}) &= \mathbb{E}_{p(\mathbf{x}, y)}[L_{c(\mathbf{x})}^\Phi(\mathbf{g}(\mathbf{x}), y)] \\ &= \mathbb{E}_{p(\mathbf{x}, y)}[(\Phi(\mathbf{g}(\mathbf{x}), y) + (1 - c(\mathbf{x}))\Phi(\mathbf{g}(\mathbf{x}), K + 1))/(2 - c(\mathbf{x}))] \\ &= \int_{\mathbf{x}} \sum_{y=1}^K \Phi(\mathbf{g}(\mathbf{x}), y) \frac{p(\mathbf{x}, y)}{2 - c(\mathbf{x})} d\mathbf{x} + \int_{\mathbf{x}} \frac{(1 - c(\mathbf{x}))p(\mathbf{x})}{2 - c(\mathbf{x})} \Phi(\mathbf{g}(\mathbf{x}), K + 1) d\mathbf{x} \\ &= \tilde{R}_\Phi(\mathbf{g}) \end{aligned}$$

615 □

616 The derivation of its estimation error bound is similar to that of Theorem 3 by modifying the upper
 617 bound and Lipschitz constant, and the necessity and sufficiency of the ℓ_{01} -calibration of Φ can also
 618 be proved by utilizing the arbitrariness of $\tilde{p}(\mathbf{x}, y)$ as in Appendix D.

619 **G.2 Experiments on SVHN**

620 In this section, we compare our proposed surrogate $L_{c(\mathbf{x})}^\Phi$ with CE and DEFER on SVHN. The
 621 cost-sensitive learning-based method [8] is not compared since it cannot tackle the case of instance-
 622 dependent cost.

623 In the experiments, we use SVHN [42] to demonstrate the effectiveness of $L_{c(\mathbf{x})}^\Phi$. To generate
 624 instance-dependent costs, we split 10% of the training dataset and manually corrupt it into a
 625 binary dataset by aggregating the 10 classes into ['0', '2', '3', '5', '6', '8', '9'] and ['1', '4', '7'].
 626 We train a binary classifier with on the corrupted dataset with 10 epochs. Then we further use the
 627 obtained classifier on training and testing set to split them into 2 parts. For any \mathbf{x} that is classified as
 628 ['0', '2', '3', '5', '6', '8', '9'], we set $c(\mathbf{x}) = c_1$ and c_2 otherwise. In the experiments, Adam with
 629 default momentum is used with learning rate, batch size and weight decay set to 1e-3, 1024, and 1e-4,
 630 respectively. The model used is ResNet-18.

Table 3: The mean and standard error of the zero-one- c losses (rescaled to 0-100), rejection ratio, and missclassification rates of the accepted data for 5 trails. The best and comparable methods based on the paired t-test at the significance level 5% are highlighted in boldface.

Method	(c_1, c_2)	CE			DEFER			GCE		
		01c	Rej	01	01c	Rej	01	01c	Rej	01
SVHN	(0.50, 0.10)	8.03 (0.16)	4.46 (0.54)	7.60 (0.01)	8.00 (0.30)	9.20 (0.72)	5.07 (0.25)	7.20 (0.17)	6.73 (6.73)	5.13 (5.13)
	(0.45, 0.15)	7.80 (0.26)	4.36 (0.31)	7.03 (0.23)	9.10 (0.46)	9.93 (0.41)	5.07 (0.42)	6.93 (0.31)	7.00 (0.35)	4.70 (0.26)
		(0.40, 0.20)	7.70 (0.10)	4.50 (0.50)	6.83 (0.25)	7.80 (0.26)	11.13 (1.27)	5.00 (0.44)	7.03 (0.21)	7.93 (0.55)
	(0.35, 0.25)	7.76 (0.12)	4.90 (0.20)	6.67 (0.15)	7.70 (0.26)	11.93 (0.45)	4.80 (0.10)	6.83 (0.20)	8.43 (0.31)	4.63 (0.15)

631 The experimental results are reported in the table above. It can be seen that in the scenario of
632 instance-dependent cost, the proposed surrogate with GCE loss still outperforms baseline methods,
633 which aligns with the observations in Section 6.

634 H Limitations and Potential Negative Social Impacts

635 **Limitations:** This framework is used for multi-class classification with rejection, while there are
636 also other scenarios for learning with rejection, e.g., AUC optimization with rejection [51]. We
637 believe that extensions to CwR with complex evaluation is a promising future direction.

638 **Potential Negative Social Impacts:** Though classification with rejection can be useful in risk-
639 critical missions, it can lead to inefficient services once abused, i.e., used in risk-insensitive missions.
640 This is also the potential negative social impact of all the methods for CwR.