
Classification Calibration Dimension for General Multiclass Losses

Harish G. Ramaswamy Shivani Agarwal
Department of Computer Science and Automation
Indian Institute of Science, Bangalore 560012, India
{harish_gurup, shivani}@csa.iisc.ernet.in

Abstract

We study consistency properties of surrogate loss functions for general multiclass classification problems, defined by a general loss matrix. We extend the notion of classification calibration, which has been studied for binary and multiclass 0-1 classification problems (and for certain other specific learning problems), to the general multiclass setting, and derive necessary and sufficient conditions for a surrogate loss to be classification calibrated with respect to a loss matrix in this setting. We then introduce the notion of *classification calibration dimension* of a multiclass loss matrix, which measures the smallest ‘size’ of a prediction space for which it is possible to design a convex surrogate that is classification calibrated with respect to the loss matrix. We derive both upper and lower bounds on this quantity, and use these results to analyze various loss matrices. In particular, as one application, we provide a different route from the recent result of Duchi et al. (2010) for analyzing the difficulty of designing ‘low-dimensional’ convex surrogates that are consistent with respect to pairwise subset ranking losses. We anticipate the classification calibration dimension may prove to be a useful tool in the study and design of surrogate losses for general multiclass learning problems.

1 Introduction

There has been significant interest and progress in recent years in understanding consistency of learning methods for various finite-output learning problems, such as binary classification, multiclass 0-1 classification, and various forms of ranking and multi-label prediction problems [1–15]. Such finite-output problems can all be viewed as instances of a general multiclass learning problem, whose structure is defined by a loss function, or equivalently, by a loss matrix. While the studies above have contributed to the understanding of learning problems corresponding to certain forms of loss matrices, a framework for analyzing consistency properties for a general multiclass learning problem, defined by a general loss matrix, has remained elusive.

In this paper, we analyze consistency of surrogate losses for general multiclass learning problems, building on the results of [3, 5–7] and others. We start in Section 2 with some background and examples that will be used as running examples to illustrate concepts throughout the paper, and formalize the notion of classification calibration with respect to a general loss matrix. In Section 3, we derive both necessary and sufficient conditions for classification calibration with respect to general multiclass losses; these are both of independent interest and useful in our later results. Section 4 introduces the notion of *classification calibration dimension* of a loss matrix, a fundamental quantity that measures the smallest ‘size’ of a prediction space for which it is possible to design a convex surrogate that is classification calibrated with respect to the loss matrix. We derive both upper and lower bounds on this quantity, and use these results to analyze various loss matrices. As one application, in Section 5, we provide a different route from the recent result of Duchi et al. [10] for analyzing the difficulty of designing ‘low-dimensional’ convex surrogates that are consistent with respect to certain pairwise subset ranking losses. We conclude in Section 6 with some future directions.

2 Preliminaries, Examples, and Background

Setup. We are given training examples $(X_1, Y_1), \dots, (X_m, Y_m)$ drawn i.i.d. from a distribution D on $\mathcal{X} \times \mathcal{Y}$, where \mathcal{X} is an instance space and $\mathcal{Y} = [n] = \{1, \dots, n\}$ is a finite set of *class labels*. We are also given a finite set $\mathcal{T} = [k] = \{1, \dots, k\}$ of *target labels* in which predictions are to be made, and a *loss function* $\ell : \mathcal{Y} \times \mathcal{T} \rightarrow [0, \infty)$, where $\ell(y, t)$ denotes the loss incurred on predicting $t \in \mathcal{T}$ when the label is $y \in \mathcal{Y}$. In many common learning problems, $\mathcal{T} = \mathcal{Y}$, but in general, these could be different (e.g. when there is an ‘abstain’ option available to a classifier, in which case $k = n + 1$).

We will find it convenient to represent the loss function ℓ as a *loss matrix* $\mathbf{L} \in \mathbb{R}_+^{n \times k}$ (here $\mathbb{R}_+ = [0, \infty)$), and for each $y \in [n], t \in [k]$, will denote by ℓ_{yt} the (y, t) -th element of \mathbf{L} , $\ell_{yt} = (\mathbf{L})_{yt} = \ell(y, t)$, and by ℓ_t the t -th column of \mathbf{L} , $\ell_t = (\ell_{1t}, \dots, \ell_{nt})^\top \in \mathbb{R}^n$. Some examples follow:

Example 1 (0-1 loss). Here $\mathcal{Y} = \mathcal{T} = [n]$, and the loss incurred is 1 if the predicted label t is different from the actual class label y , and 0 otherwise: $\ell^{0-1}(y, t) = \mathbf{1}(t \neq y)$, where $\mathbf{1}(\cdot)$ is 1 if the argument is true and 0 otherwise. The loss matrix \mathbf{L}^{0-1} for $n = 3$ is shown in Figure 1(a).

Example 2 (Ordinal regression loss). Here $\mathcal{Y} = \mathcal{T} = [n]$, and predictions t farther away from the actual class label y are penalized more heavily, e.g. using absolute distance: $\ell^{\text{ord}}(y, t) = |t - y|$. The loss matrix \mathbf{L}^{ord} for $n = 3$ is shown in Figure 1(b).

Example 3 (Hamming loss). Here $\mathcal{Y} = \mathcal{T} = [2^r]$ for some $r \in \mathbb{N}$, and the loss incurred on predicting t when the actual class label is y is the number of bit-positions in which the r -bit binary representations of $t - 1$ and $y - 1$ differ: $\ell^{\text{Ham}}(y, t) = \sum_{i=1}^r \mathbf{1}((t - 1)_i \neq (y - 1)_i)$, where for any $z \in \{0, \dots, 2^r - 1\}$, $z_i \in \{0, 1\}$ denotes the i -th bit in the r -bit binary representation of z . The loss matrix \mathbf{L}^{Ham} for $r = 2$ is shown in Figure 1(c). This loss is used in sequence labeling tasks [16].

Example 4 (‘Abstain’ loss). Here $\mathcal{Y} = [n]$ and $\mathcal{T} = [n + 1]$, where $t = n + 1$ denotes ‘abstain’. One possible loss function in this setting assigns a loss of 1 to incorrect predictions in $[n]$, 0 to correct predictions, and $\frac{1}{2}$ for abstaining: $\ell^{(3)}(y, t) = \mathbf{1}(t \neq y) \mathbf{1}(t \in [n]) + \frac{1}{2} \mathbf{1}(t = n + 1)$. The loss matrix $\mathbf{L}^{(3)}$ for $n = 3$ is shown in Figure 1(d).

The goal in the above setting is to learn from the training examples a function $h : \mathcal{X} \rightarrow [k]$ with low expected loss on a new example drawn from D , which we will refer to as the ℓ -risk of h :

$$\text{er}_D^\ell[h] \triangleq \mathbf{E}_{(X, Y) \sim D} \ell(Y, h(X)) = \mathbf{E}_X \sum_{y=1}^n p_y(X) \ell(y, h(X)) = \mathbf{E}_X \mathbf{p}(X)^\top \ell_{h(X)}, \quad (1)$$

where $p_y(x) = \mathbf{P}(Y = y \mid X = x)$ under D , and $\mathbf{p}(x) = (p_1(x), \dots, p_n(x))^\top \in \mathbb{R}^n$ denotes the conditional probability vector at x . In particular, the goal is to learn a function with ℓ -risk close to the *optimal* ℓ -risk, defined as

$$\text{er}_D^{\ell, *} \triangleq \inf_{h: \mathcal{X} \rightarrow [k]} \text{er}_D^\ell[h] = \inf_{h: \mathcal{X} \rightarrow [k]} \mathbf{E}_X \mathbf{p}(X)^\top \ell_{h(X)} = \mathbf{E}_X \min_{t \in [k]} \mathbf{p}(X)^\top \ell_t. \quad (2)$$

Minimizing the discrete ℓ -risk directly is typically difficult computationally; consequently, one usually employs a *surrogate loss function* $\psi : \mathcal{Y} \times \widehat{\mathcal{T}} \rightarrow \mathbb{R}_+$ operating on a *surrogate target space* $\widehat{\mathcal{T}} \subseteq \mathbb{R}^d$ for some appropriate $d \in \mathbb{N}$,¹ and minimizes (approximately, based on the training sample) the ψ -risk instead, defined for a (vector) function $\mathbf{f} : \mathcal{X} \rightarrow \widehat{\mathcal{T}}$ as

$$\text{er}_D^\psi[\mathbf{f}] \triangleq \mathbf{E}_{(X, Y) \sim D} \psi(Y, \mathbf{f}(X)) = \mathbf{E}_X \sum_{y=1}^n p_y(X) \psi(y, \mathbf{f}(X)). \quad (3)$$

The learned function $\mathbf{f} : \mathcal{X} \rightarrow \widehat{\mathcal{T}}$ is then used to make predictions in $[k]$ via some transformation $\text{pred} : \widehat{\mathcal{T}} \rightarrow [k]$: the prediction on a new instance $x \in \mathcal{X}$ is given by $\text{pred}(\mathbf{f}(x))$, and the ℓ -risk incurred is $\text{er}_D^\ell[\text{pred} \circ \mathbf{f}]$. As an example, several algorithms for multiclass classification with respect to 0-1 loss learn a function of the form $\mathbf{f} : \mathcal{X} \rightarrow \mathbb{R}^n$ and predict according to $\text{pred}(\mathbf{f}(x)) = \text{argmax}_{t \in [n]} f_t(x)$.

Below we will find it useful to represent the surrogate loss function ψ via n real-valued functions $\psi_y : \widehat{\mathcal{T}} \rightarrow \mathbb{R}_+$ defined as $\psi_y(\hat{\mathbf{t}}) = \psi(y, \hat{\mathbf{t}})$ for $y \in [n]$, or equivalently, as a vector-valued function $\boldsymbol{\psi} : \widehat{\mathcal{T}} \rightarrow \mathbb{R}_+^n$ defined as $\boldsymbol{\psi}(\hat{\mathbf{t}}) = (\psi_1(\hat{\mathbf{t}}), \dots, \psi_n(\hat{\mathbf{t}}))^\top$. We will also define the sets

$$\mathcal{R}_\psi \triangleq \{\boldsymbol{\psi}(\hat{\mathbf{t}}) : \hat{\mathbf{t}} \in \widehat{\mathcal{T}}\} \quad \text{and} \quad \mathcal{S}_\psi \triangleq \text{conv}(\mathcal{R}_\psi), \quad (4)$$

where for any $\mathcal{A} \subseteq \mathbb{R}^n$, $\text{conv}(\mathcal{A})$ denotes the convex hull of \mathcal{A} .

¹Equivalently, one can define $\psi : \mathcal{Y} \times \mathbb{R}^d \rightarrow \bar{\mathbb{R}}_+$, where $\bar{\mathbb{R}}_+ = \mathbb{R}_+ \cup \{\infty\}$ and $\psi(y, \hat{\mathbf{t}}) = \infty \forall \hat{\mathbf{t}} \notin \widehat{\mathcal{T}}$.

$$\begin{array}{cccc}
\begin{bmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix} & \begin{bmatrix} 0 & 1 & 2 \\ 1 & 0 & 1 \\ 2 & 1 & 0 \end{bmatrix} & \begin{bmatrix} 0 & 1 & 1 & 2 \\ 1 & 0 & 2 & 1 \\ 1 & 2 & 0 & 1 \\ 2 & 1 & 1 & 0 \end{bmatrix} & \begin{bmatrix} 0 & 1 & 1 & \frac{1}{2} \\ 1 & 0 & 1 & \frac{1}{2} \\ 1 & 1 & 0 & \frac{1}{2} \end{bmatrix} \\
\text{(a)} & \text{(b)} & \text{(c)} & \text{(d)}
\end{array}$$

Figure 1: Loss matrices corresponding to Examples 1-4: (a) \mathbf{L}^{0-1} for $n = 3$; (b) \mathbf{L}^{ord} for $n = 3$; (c) \mathbf{L}^{Ham} for $r = 2$ ($n = 4$); (d) $\mathbf{L}^{(?)}$ for $n = 3$.

Under suitable conditions, algorithms that approximately minimize the ψ -risk based on a training sample are known to be consistent with respect to the ψ -risk, i.e. to converge (in probability) to the *optimal ψ -risk*, defined as

$$\text{er}_D^{\psi,*} \triangleq \inf_{\mathbf{f}: \mathcal{X} \rightarrow \widehat{\mathcal{T}}} \text{er}_D^\psi[\mathbf{f}] = \inf_{\mathbf{f}: \mathcal{X} \rightarrow \widehat{\mathcal{T}}} \mathbf{E}_X \mathbf{p}(X)^\top \psi(\mathbf{f}(X)) = \mathbf{E}_X \inf_{\mathbf{z} \in \mathcal{R}_\psi} \mathbf{p}(X)^\top \mathbf{z} = \mathbf{E}_X \inf_{\mathbf{z} \in \mathcal{S}_\psi} \mathbf{p}(X)^\top \mathbf{z}. \quad (5)$$

This raises the natural question of whether, for a given loss ℓ , there are surrogate losses ψ for which consistency with respect to the ψ -risk also guarantees consistency with respect to the ℓ -risk, i.e. guarantees convergence (in probability) to the optimal ℓ -risk (defined in Eq. (2)). This question has been studied in detail for the 0-1 loss, and for square losses of the form $\ell(y, t) = a_y \mathbf{1}(t \neq y)$, which can be analyzed similarly to the 0-1 loss [6, 7]. In this paper, we consider this question for general multiclass losses $\ell: [n] \times [k] \rightarrow \mathbb{R}_+$, including rectangular losses with $k \neq n$. The only assumption we make on ℓ is that for each $t \in [k]$, $\exists \mathbf{p} \in \Delta_n$ such that $\text{argmin}_{t' \in [k]} \mathbf{p}^\top \ell_{t'} = \{t\}$ (otherwise the label t never needs to be predicted and can simply be ignored).²

Definitions and Results. We will need the following definitions and basic results, generalizing those of [5–7]. The notion of classification calibration will be central to our study; as Theorem 3 below shows, classification calibration of a surrogate loss ψ w.r.t. ℓ corresponds to the property that consistency w.r.t. ψ -risk implies consistency w.r.t. ℓ -risk. Proofs of these results are straightforward generalizations of those in [6, 7] and are omitted.

Definition 1 (Classification calibration). *A surrogate loss function $\psi: [n] \times \widehat{\mathcal{T}} \rightarrow \mathbb{R}_+$ is said to be classification calibrated with respect to a loss function $\ell: [n] \times [k] \rightarrow \mathbb{R}_+$ over $\mathcal{P} \subseteq \Delta_n$ if there exists a function $\text{pred}: \widehat{\mathcal{T}} \rightarrow [k]$ such that*

$$\forall \mathbf{p} \in \mathcal{P}: \quad \inf_{\hat{\mathbf{t}} \in \widehat{\mathcal{T}}: \text{pred}(\hat{\mathbf{t}}) \notin \text{argmin}_t \mathbf{p}^\top \ell_t} \mathbf{p}^\top \psi(\hat{\mathbf{t}}) > \inf_{\hat{\mathbf{t}} \in \widehat{\mathcal{T}}} \mathbf{p}^\top \psi(\hat{\mathbf{t}}).$$

Lemma 2. *Let $\ell: [n] \times [k] \rightarrow \mathbb{R}_+$ and $\psi: [n] \times \widehat{\mathcal{T}} \rightarrow \mathbb{R}_+$. Then ψ is classification calibrated with respect to ℓ over $\mathcal{P} \subseteq \Delta_n$ iff there exists a function $\text{pred}': \mathcal{S}_\psi \rightarrow [k]$ such that*

$$\forall \mathbf{p} \in \mathcal{P}: \quad \inf_{\mathbf{z} \in \mathcal{S}_\psi: \text{pred}'(\mathbf{z}) \notin \text{argmin}_t \mathbf{p}^\top \ell_t} \mathbf{p}^\top \mathbf{z} > \inf_{\mathbf{z} \in \mathcal{S}_\psi} \mathbf{p}^\top \mathbf{z}.$$

Theorem 3. *Let $\ell: [n] \times [k] \rightarrow \mathbb{R}_+$ and $\psi: [n] \times \widehat{\mathcal{T}} \rightarrow \mathbb{R}_+$. Then ψ is classification calibrated with respect to ℓ over Δ_n iff \exists a function $\text{pred}: \widehat{\mathcal{T}} \rightarrow [k]$ such that for all distributions D on $\mathcal{X} \times [n]$ and all sequences of random (vector) functions $\mathbf{f}_m: \mathcal{X} \rightarrow \widehat{\mathcal{T}}$ (depending on $(X_1, Y_1), \dots, (X_m, Y_m)$),³*

$$\text{er}_D^\psi[\mathbf{f}_m] \xrightarrow{P} \text{er}_D^{\psi,*} \quad \text{implies} \quad \text{er}_D^\ell[\text{pred} \circ \mathbf{f}_m] \xrightarrow{P} \text{er}_D^{\ell,*}.$$

Definition 4 (Positive normals). *Let $\psi: [n] \times \widehat{\mathcal{T}} \rightarrow \mathbb{R}_+$. For each point $\mathbf{z} \in \mathcal{S}_\psi$, the set of positive normals at \mathbf{z} is defined as⁴*

$$\mathcal{N}_{\mathcal{S}_\psi}(\mathbf{z}) \triangleq \{ \mathbf{p} \in \Delta_n : \mathbf{p}^\top (\mathbf{z} - \mathbf{z}') \leq 0 \forall \mathbf{z}' \in \mathcal{S}_\psi \}.$$

Definition 5 (Trigger probabilities). *Let $\ell: [n] \times [k] \rightarrow \mathbb{R}_+$. For each $t \in [k]$, the set of trigger probabilities of t with respect to ℓ is defined as*

$$\mathcal{Q}_t^\ell \triangleq \{ \mathbf{p} \in \Delta_n : \mathbf{p}^\top (\ell_t - \ell_{t'}) \leq 0 \forall t' \in [k] \} = \{ \mathbf{p} \in \Delta_n : t \in \text{argmin}_{t' \in [k]} \mathbf{p}^\top \ell_{t'} \}.$$

Examples of trigger probability sets for various losses are shown in Figure 2.

²Here Δ_n denotes the probability simplex in \mathbb{R}^n , $\Delta_n = \{ \mathbf{p} \in \mathbb{R}^n : p_i \geq 0 \forall i \in [n], \sum_{i=1}^n p_i = 1 \}$.

³Here \xrightarrow{P} denotes convergence in probability.

⁴The set of positive normals is non-empty only at points \mathbf{z} in the boundary of \mathcal{S}_ψ .




Figure 2: Trigger probability sets for (a) 0-1 loss ℓ^{0-1} ; (b) ordinal regression loss ℓ^{ord} ; and (c) ‘abstain’ loss $\ell^{(?)}$; all for $n = 3$, for which the probability simplex can be visualized easily. Calculations of these sets can be found in the appendix. We note that such sets have also been studied in [17, 18].

3 Necessary and Sufficient Conditions for Classification Calibration

We start by giving a necessary condition for classification calibration of a surrogate loss ψ with respect to any multiclass loss ℓ over Δ_n , which requires the positive normals of all points $\mathbf{z} \in \mathcal{S}_\psi$ to be ‘well-behaved’ w.r.t. ℓ and generalizes the ‘admissibility’ condition used for 0-1 loss in [7]. All proofs not included in the main text can be found in the appendix.

Theorem 6. *Let $\psi : [n] \times \widehat{\mathcal{T}} \rightarrow \mathbb{R}_+$ be classification calibrated with respect to $\ell : [n] \times [k] \rightarrow \mathbb{R}_+$ over Δ_n . Then for all $\mathbf{z} \in \mathcal{S}_\psi$, there exists some $t \in [k]$ such that $\mathcal{N}_{\mathcal{S}_\psi}(\mathbf{z}) \subseteq \mathcal{Q}_t^\ell$.*

We note that, as in [7], it is possible to give a necessary and sufficient condition for classification calibration in terms of a similar property holding for positive normals associated with projections of \mathcal{S}_ψ in lower dimensions. Instead, below we give a different sufficient condition that will be helpful in showing classification calibration of certain surrogates. In particular, we show that for a surrogate loss ψ to be classification calibrated with respect to ℓ over Δ_n , it is sufficient for the above property of positive normals to hold only at a finite number of points in \mathcal{R}_ψ , as long as their positive normal sets jointly cover Δ_n :

Theorem 7. *Let $\ell : [n] \times [k] \rightarrow \mathbb{R}_+$ and $\psi : [n] \times \widehat{\mathcal{T}} \rightarrow \mathbb{R}_+$. Suppose there exist $r \in \mathbb{N}$ and $\mathbf{z}_1, \dots, \mathbf{z}_r \in \mathcal{R}_\psi$ such that $\bigcup_{j=1}^r \mathcal{N}_{\mathcal{S}_\psi}(\mathbf{z}_j) = \Delta_n$ and for each $j \in [r]$, $\exists t \in [k]$ such that $\mathcal{N}_{\mathcal{S}_\psi}(\mathbf{z}_j) \subseteq \mathcal{Q}_t^\ell$. Then ψ is classification calibrated with respect to ℓ over Δ_n .*

Computation of $\mathcal{N}_{\mathcal{S}_\psi}(\mathbf{z})$. The conditions in the above results both involve the sets of positive normals $\mathcal{N}_{\mathcal{S}_\psi}(\mathbf{z})$ at various points $\mathbf{z} \in \mathcal{S}_\psi$. Thus in order to use the above results to show that a surrogate ψ is (or is not) classification calibrated with respect to a loss ℓ , one needs to be able to compute or characterize the sets $\mathcal{N}_{\mathcal{S}_\psi}(\mathbf{z})$. Here we give a method for computing these sets for certain surrogate losses ψ and points $\mathbf{z} \in \mathcal{S}_\psi$.

Lemma 8. *Let $\widehat{\mathcal{T}} \subseteq \mathbb{R}^d$ be a convex set and let $\psi : \widehat{\mathcal{T}} \rightarrow \mathbb{R}_+^n$ be convex.⁵ Let $\mathbf{z} = \psi(\hat{\mathbf{t}})$ for some $\hat{\mathbf{t}} \in \widehat{\mathcal{T}}$ such that for each $y \in [n]$, the subdifferential of ψ_y at $\hat{\mathbf{t}}$ can be written as $\partial\psi_y(\hat{\mathbf{t}}) = \text{conv}(\{\mathbf{w}_1^y, \dots, \mathbf{w}_{s_y}^y\})$ for some $s_y \in \mathbb{N}$ and $\mathbf{w}_1^y, \dots, \mathbf{w}_{s_y}^y \in \mathbb{R}^d$.⁶ Let $s = \sum_{y=1}^n s_y$, and let*

$$\mathbf{A} = [\mathbf{w}_1^1 \dots \mathbf{w}_{s_1}^1 \mathbf{w}_1^2 \dots \mathbf{w}_{s_2}^2 \dots \mathbf{w}_1^n \dots \mathbf{w}_{s_n}^n] \in \mathbb{R}^{d \times s}; \quad \mathbf{B} = [b_{yj}] \in \mathbb{R}^{n \times s},$$

where b_{yj} is 1 if the j -th column of \mathbf{A} came from $\{\mathbf{w}_1^y, \dots, \mathbf{w}_{s_y}^y\}$ and 0 otherwise. Then

$$\mathcal{N}_{\mathcal{S}_\psi}(\mathbf{z}) = \{\mathbf{p} \in \Delta_n : \mathbf{p} = \mathbf{B}\mathbf{q} \text{ for some } \mathbf{q} \in \text{Null}(\mathbf{A}) \cap \Delta_s\},$$

where $\text{Null}(\mathbf{A}) \subseteq \mathbb{R}^s$ denotes the null space of the matrix \mathbf{A} .

⁵A vector function is convex if all its component functions are convex.

⁶Recall that the subdifferential of a convex function $\phi : \mathbb{R}^d \rightarrow \mathbb{R}_+$ at a point $\mathbf{u}_0 \in \mathbb{R}^d$ is defined as $\partial\phi(\mathbf{u}_0) = \{\mathbf{w} \in \mathbb{R}^d : \phi(\mathbf{u}) - \phi(\mathbf{u}_0) \geq \mathbf{w}^\top(\mathbf{u} - \mathbf{u}_0) \forall \mathbf{u} \in \mathbb{R}^d\}$ and is a convex set in \mathbb{R}^d (e.g. see [19]).

We give an example illustrating the use of Theorem 7 and Lemma 8 to show classification calibration of a certain surrogate loss with respect to the ordinal regression loss ℓ^{ord} defined in Example 2:

Example 5 (Classification calibrated surrogate for ordinal regression loss). *Consider the ordinal regression loss ℓ^{ord} defined in Example 2 for $n = 3$. Let $\widehat{\mathcal{T}} = \mathbb{R}$, and let $\psi : \{1, 2, 3\} \times \mathbb{R} \rightarrow \mathbb{R}_+$ be defined as (see Figure 3)*

$$\psi(y, \hat{t}) = |\hat{t} - y| \quad \forall y \in \{1, 2, 3\}, \hat{t} \in \mathbb{R}. \quad (6)$$

Thus $\mathcal{R}_\psi = \{\psi(\hat{t}) = (|\hat{t} - 1|, |\hat{t} - 2|, |\hat{t} - 3|)^\top : \hat{t} \in \mathbb{R}\}$. We will show there are 3 points in \mathcal{R}_ψ satisfying the conditions of Theorem 7. Specifically, consider $\hat{t}_1 = 1$, $\hat{t}_2 = 2$, and $\hat{t}_3 = 3$, giving $\mathbf{z}_1 = \psi(\hat{t}_1) = (0, 1, 2)^\top$, $\mathbf{z}_2 = \psi(\hat{t}_2) = (1, 0, 1)^\top$, and $\mathbf{z}_3 = \psi(\hat{t}_3) = (2, 1, 0)^\top$ in \mathcal{R}_ψ . Observe that $\widehat{\mathcal{T}}$ here is a convex set and $\psi : \widehat{\mathcal{T}} \rightarrow \mathbb{R}^3$ is a convex function. Moreover, for $\hat{t}_1 = 1$, we have

$$\begin{aligned} \partial\psi_1(1) &= [-1, 1] = \text{conv}(\{+1, -1\}); \\ \partial\psi_2(1) &= \{-1\} = \text{conv}(\{-1\}); \\ \partial\psi_3(1) &= \{-1\} = \text{conv}(\{-1\}). \end{aligned}$$

Therefore, we can use Lemma 8 to compute $\mathcal{N}_{\mathcal{S}_\psi}(\mathbf{z}_1)$. Here $s = 4$, and

$$\mathbf{A} = \begin{bmatrix} +1 & -1 & -1 & -1 \end{bmatrix}; \quad \mathbf{B} = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

This gives

$$\begin{aligned} \mathcal{N}_{\mathcal{S}_\psi}(\mathbf{z}_1) &= \{\mathbf{p} \in \Delta_3 : \mathbf{p} = (q_1 + q_2, q_3, q_4) \text{ for some } \mathbf{q} \in \Delta_4, q_1 - q_2 - q_3 - q_4 = 0\} \\ &= \{\mathbf{p} \in \Delta_3 : \mathbf{p} = (q_1 + q_2, q_3, q_4) \text{ for some } \mathbf{q} \in \Delta_4, q_1 = \frac{1}{2}\} \\ &= \{\mathbf{p} \in \Delta_3 : p_1 \geq \frac{1}{2}\} \\ &= \mathcal{Q}_1^{\text{ord}}. \end{aligned}$$

A similar procedure yields $\mathcal{N}_{\mathcal{S}_\psi}(\mathbf{z}_2) = \mathcal{Q}_2^{\text{ord}}$ and $\mathcal{N}_{\mathcal{S}_\psi}(\mathbf{z}_3) = \mathcal{Q}_3^{\text{ord}}$. Thus, by Theorem 7, we get that ψ is classification calibrated with respect to ℓ^{ord} over Δ_3 .

We note that in general, computational procedures such as Fourier-Motzkin elimination [20] can be helpful in computing $\mathcal{N}_{\mathcal{S}_\psi}(\mathbf{z})$ via Lemma 8.

4 Classification Calibration Dimension

We now turn to the study of a fundamental quantity associated with the property of classification calibration with respect to a general multiclass loss ℓ . Specifically, in the above example, we saw that to develop a classification calibrated surrogate loss w.r.t. the ordinal regression loss for $n = 3$, it was sufficient to consider a surrogate target space $\widehat{\mathcal{T}} = \mathbb{R}$, with dimension $d = 1$; in addition, this yielded a convex surrogate $\psi : \mathbb{R} \rightarrow \mathbb{R}_+^3$ which can be used in developing computationally efficient algorithms. In fact the same surrogate target space with $d = 1$ can be used to develop a similar convex, classification calibrated surrogate loss w.r.t. the ordinal regression loss for *any* $n \in \mathbb{N}$. However not all losses ℓ have such ‘low-dimensional’ surrogates. This raises the natural question of what is the smallest dimension d that supports a convex classification calibrated surrogate for a given multiclass loss ℓ , and leads us to the following definition:

Definition 9 (Classification calibration dimension). *Let $\ell : [n] \times [k] \rightarrow \mathbb{R}_+$. Define the classification calibration dimension (CC dimension) of ℓ as*

$$\text{CCdim}(\ell) \triangleq \min \{d \in \mathbb{N} : \exists \text{ a convex set } \widehat{\mathcal{T}} \subseteq \mathbb{R}^d \text{ and a convex surrogate } \psi : \widehat{\mathcal{T}} \rightarrow \mathbb{R}_+^n \text{ that is classification calibrated w.r.t. } \ell \text{ over } \Delta_n\},$$

if the above set is non-empty, and $\text{CCdim}(\ell) = \infty$ otherwise.

From the above discussion, $\text{CCdim}(\ell^{\text{ord}}) = 1$ for all n . In the following, we will be interested in developing an understanding of the CC dimension for general losses ℓ , and in particular in deriving upper and lower bounds on this.




Figure 3: The surrogate ψ

4.1 Upper Bounds on the Classification Calibration Dimension

We start with a simple result that establishes that the CC dimension of any multiclass loss ℓ is finite, and in fact is strictly smaller than the number of class labels n .

Lemma 10. *Let $\ell : [n] \times [k] \rightarrow \mathbb{R}_+$. Let $\widehat{\mathcal{T}} = \{\hat{\mathbf{t}} \in \mathbb{R}_+^{n-1} : \sum_{j=1}^{n-1} \hat{t}_j \leq 1\}$, and for each $y \in [n]$, let $\psi_y : \widehat{\mathcal{T}} \rightarrow \mathbb{R}_+$ be given by*

$$\psi_y(\hat{\mathbf{t}}) = \mathbf{1}(y \neq \bar{n}) (\hat{t}_y - 1)^2 + \sum_{j \in [n-1], j \neq y} \hat{t}_j^2.$$

Then ψ is classification calibrated with respect to ℓ over Δ_n . In particular, since ψ is convex, $\text{CCdim}(\ell) \leq n - 1$.

It may appear surprising that the convex surrogate ψ in the above lemma is classification calibrated with respect to *all* multiclass losses ℓ on n classes. However this makes intuitive sense, since in principle, for any multiclass problem, if one can estimate the conditional probabilities of the n classes accurately (which requires estimating $n - 1$ real-valued functions on \mathcal{X}), then one can predict a target label that minimizes the expected loss according to these probabilities. Minimizing the above surrogate effectively corresponds to such class probability estimation. Indeed, the above lemma can be shown to hold for any surrogate that is a strictly proper composite multiclass loss [21].

In practice, when the number of class labels n is large (such as in a sequence labeling task, where n is exponential in the length of the input sequence), the above result is not very helpful; in such cases, it is of interest to develop algorithms operating on a surrogate target space in a lower-dimensional space. Next we give a different upper bound on the CC dimension that depends on the loss ℓ , and for certain losses, can be significantly tighter than the general bound above.

Theorem 11. *Let $\ell : [n] \times [k] \rightarrow \mathbb{R}_+$. Then $\text{CCdim}(\ell) \leq \text{rank}(\mathbf{L})$, the rank of the loss matrix \mathbf{L} .*

Proof. Let $\text{rank}(\mathbf{L}) = d$. We will construct a convex classification calibrated surrogate loss ψ for ℓ with surrogate target space $\widehat{\mathcal{T}} \subseteq \mathbb{R}^d$.

Let $\ell_{t_1}, \dots, \ell_{t_d}$ be linearly independent columns of \mathbf{L} . Let $\{\mathbf{e}_1, \dots, \mathbf{e}_d\}$ denote the standard basis in \mathbb{R}^d . We can define a linear function $\tilde{\psi} : \mathbb{R}^d \rightarrow \mathbb{R}^n$ by

$$\tilde{\psi}(\mathbf{e}_j) = \ell_{t_j} \quad \forall j \in [d].$$

Then for each \mathbf{z} in the column space of \mathbf{L} , there exists a unique vector $\mathbf{u} \in \mathbb{R}^d$ such that $\tilde{\psi}(\mathbf{u}) = \mathbf{z}$. In particular, there exist unique vectors $\mathbf{u}_1, \dots, \mathbf{u}_k \in \mathbb{R}^d$ such that for each $t \in [k]$, $\tilde{\psi}(\mathbf{u}_t) = \ell_t$. Let $\widehat{\mathcal{T}} = \text{conv}(\{\mathbf{u}_1, \dots, \mathbf{u}_k\})$, and define $\psi : \widehat{\mathcal{T}} \rightarrow \mathbb{R}_+^n$ as

$$\psi(\hat{\mathbf{t}}) = \tilde{\psi}(\hat{\mathbf{t}});$$

we note that the resulting vectors are always in \mathbb{R}_+^n , since by definition, for any $\hat{\mathbf{t}} = \sum_{t=1}^k \alpha_t \mathbf{u}_t$ for $\alpha \in \Delta_k$, $\psi(\hat{\mathbf{t}}) = \sum_{t=1}^k \alpha_t \ell_t$, and $\ell_t \in \mathbb{R}_+^n \forall t \in [k]$. The function ψ is clearly convex. To show ψ is classification calibrated w.r.t. ℓ over Δ_n , we will use Theorem 7. Specifically, consider the k points $\mathbf{z}_t = \psi(\mathbf{u}_t) = \ell_t \in \mathcal{R}_\psi$ for $t \in [k]$. By definition of ψ , we have $\mathcal{S}_\psi = \text{conv}(\{\ell_1, \dots, \ell_k\})$; from the definitions of positive normals and trigger probabilities, it then follows that $\mathcal{N}_{\mathcal{S}_\psi}(\mathbf{z}_t) = \mathcal{N}_{\mathcal{S}_\psi}(\ell_t) = \mathcal{Q}_t^\ell$ for all $t \in [k]$. Thus by Theorem 7, ψ is classification calibrated w.r.t. ℓ over Δ_n . \square

Example 6 (CC dimension of Hamming loss). *Consider the Hamming loss ℓ^{Ham} defined in Example 3, for $n = 2^r$. For each $i \in [r]$, define $\sigma_i \in \mathbb{R}^n$ as*

$$\sigma_{iy} = \begin{cases} +1 & \text{if } (y-1)_i, \text{ the } i\text{-th bit in the } r\text{-bit binary representation of } (y-1), \text{ is } 1 \\ -1 & \text{otherwise.} \end{cases}$$

Then the loss matrix \mathbf{L}^{Ham} satisfies

$$\mathbf{L}^{\text{Ham}} = \frac{r}{2} \mathbf{e} \mathbf{e}^\top - \frac{1}{2} \sum_{i=1}^r \sigma_i \sigma_i^\top,$$

where \mathbf{e} is the $n \times 1$ all ones vector. Thus $\text{rank}(\mathbf{L}^{\text{Ham}}) \leq r + 1$, giving us $\text{CCdim}(\ell^{\text{Ham}}) \leq r + 1$. For $r \geq 3$, this is a significantly tighter upper bound than the bound of $2^r - 1$ given by Lemma 10.

We note that the upper bound of Theorem 11 need not always be tight: for example, for the ordinal regression loss, for which we already know $\text{CCdim}(\ell^{\text{ord}}) = 1$, the theorem actually gives an upper bound of n , which is even weaker than that implied by Lemma 10.

4.2 Lower Bound on the Classification Calibration Dimension

In this section we give a lower bound on the CC dimension of a loss function ℓ and illustrate it by using it to calculate the CC dimension of the 0-1 loss. Section 5 we will explore consequences of the lower bound for classification calibrated surrogates for certain types of ranking losses. We will need the following definition:

Definition 12. *The feasible subspace dimension of a convex set \mathcal{C} at $\mathbf{p} \in \mathcal{C}$, denoted by $\mu_{\mathcal{C}}(\mathbf{p})$, is defined as the dimension of the subspace $\mathcal{F}_{\mathcal{C}}(\mathbf{p}) \cap (-\mathcal{F}_{\mathcal{C}}(\mathbf{p}))$, where $\mathcal{F}_{\mathcal{C}}(\mathbf{p})$ is the cone of feasible directions of \mathcal{C} at \mathbf{p} .⁷*

The following gives a lower bound on the CC dimension of a loss ℓ in terms of the feasible subspace dimension of the trigger probability sets \mathcal{Q}_t^ℓ at certain points $\mathbf{p} \in \mathcal{Q}_t^\ell$:

Theorem 13. *Let $\ell : [n] \times [k] \rightarrow \mathbb{R}_+$. Then for all $\mathbf{p} \in \text{relint}(\Delta_n)$ and $t \in \arg \min_{t'} \mathbf{p}^\top \ell_{t'}$ (i.e. such that $\mathbf{p} \in \mathcal{Q}_t^\ell$):⁸*

$$\text{CCdim}(\ell) \geq n - \mu_{\mathcal{Q}_t^\ell}(\mathbf{p}) - 1.$$

The proof requires extensions of the definition of positive normals and the necessary condition of Theorem 6 to sequences of points in \mathcal{S}_ψ and is quite technical. In the appendix, we provide a proof in the special case when $\mathbf{p} \in \text{relint}(\Delta_n)$ is such that $\inf_{\mathbf{z} \in \mathcal{S}_\psi} \mathbf{p}^\top \mathbf{z}$ is achieved in \mathcal{S}_ψ , which does not require these extensions. Full proof details will be provided in a longer version of the paper. Both the proof of the lower bound and its applications make use of the following lemma, which gives a method to calculate the feasible subspace dimension for certain convex sets \mathcal{C} and points $\mathbf{p} \in \mathcal{C}$:

Lemma 14. *Let $\mathcal{C} = \{\mathbf{u} \in \mathbb{R}^n : \mathbf{A}^1 \mathbf{u} \leq \mathbf{b}^1, \mathbf{A}^2 \mathbf{u} \leq \mathbf{b}^2, \mathbf{A}^3 \mathbf{u} = \mathbf{b}^3\}$. Let $\mathbf{p} \in \mathcal{C}$ be such that $\mathbf{A}^1 \mathbf{p} = \mathbf{b}^1, \mathbf{A}^2 \mathbf{p} < \mathbf{b}^2$. Then $\mu_{\mathcal{C}}(\mathbf{p}) = \text{nullity} \left(\begin{bmatrix} \mathbf{A}^1 \\ \mathbf{A}^3 \end{bmatrix} \right)$, the dimension of the null space of $\begin{bmatrix} \mathbf{A}^1 \\ \mathbf{A}^3 \end{bmatrix}$.*

The above lower bound allows us to calculate precisely the CC dimension of the 0-1 loss:

Example 7 (CC dimension of 0-1 loss). *Consider the 0-1 loss ℓ^{0-1} defined in Example 1. Take $\mathbf{p} = (\frac{1}{n}, \dots, \frac{1}{n})^\top \in \text{relint}(\Delta_n)$. Then $\mathbf{p} \in \mathcal{Q}_t^{0-1}$ for all $t \in [k] = [n]$ (see Figure 2); in particular, we have $\mathbf{p} \in \mathcal{Q}_1^{0-1}$. Now \mathcal{Q}_1^{0-1} can be written as*

$$\begin{aligned} \mathcal{Q}_1^{0-1} &= \{\mathbf{q} \in \Delta_n : q_1 \geq q_y \forall y \in \{2, \dots, n\}\} \\ &= \{\mathbf{q} \in \mathbb{R}^n : [-\mathbf{e}_{n-1} \quad \mathbf{I}_{n-1}] \mathbf{q} \leq \mathbf{0}, -\mathbf{q} \leq \mathbf{0}, \mathbf{e}_n^\top \mathbf{q} = 1\}, \end{aligned}$$

where $\mathbf{e}_{n-1}, \mathbf{e}_n$ denote the $(n-1) \times 1$ and $n \times 1$ all ones vectors, respectively, and \mathbf{I}_{n-1} denotes the $(n-1) \times (n-1)$ identity matrix. Moreover, we have $[-\mathbf{e}_{n-1} \quad \mathbf{I}_{n-1}] \mathbf{p} = \mathbf{0}, -\mathbf{p} < \mathbf{0}$. Therefore, by Lemma 14, we have

$$\mu_{\mathcal{Q}_1^{0-1}}(\mathbf{p}) = \text{nullity} \left(\begin{bmatrix} -\mathbf{e}_{n-1} & \mathbf{I}_{n-1} \\ \mathbf{e}_n^\top \end{bmatrix} \right) = \text{nullity} \left(\begin{bmatrix} -1 & 1 & 0 & \dots & 0 \\ -1 & 0 & 1 & \dots & 0 \\ & & \vdots & & \\ -1 & 0 & 0 & \dots & 1 \\ 1 & 1 & 1 & \dots & 1 \end{bmatrix} \right) = 0.$$

Thus by Theorem 13, we get $\text{CCdim}(\ell^{0-1}) \geq n - 1$. Combined with the upper bound of Lemma 10, this gives $\text{CCdim}(\ell^{0-1}) = n - 1$.

⁷For a set $\mathcal{C} \subseteq \mathbb{R}^n$ and point $\mathbf{p} \in \mathcal{C}$, the cone of feasible directions of \mathcal{C} at \mathbf{p} is defined as $\mathcal{F}_{\mathcal{C}}(\mathbf{p}) = \{\mathbf{v} \in \mathbb{R}^n : \exists \epsilon_0 > 0 \text{ such that } \mathbf{p} + \epsilon \mathbf{v} \in \mathcal{C} \forall \epsilon \in (0, \epsilon_0)\}$.

⁸Here $\text{relint}(\Delta_n)$ denotes the relative interior of Δ_n : $\text{relint}(\Delta_n) = \{\mathbf{p} \in \Delta_n : p_y > 0 \forall y \in [n]\}$.

5 Application to Pairwise Subset Ranking

We consider an application of the above framework to analyzing certain types of subset ranking problems, where each instance $x \in \mathcal{X}$ consists of a query together with a set of r documents (for simplicity, $r \in \mathbb{N}$ here is fixed), and the goal is to learn a predictor which given such an instance will return a ranking (permutation) of the r documents [8]. Duchi et al. [10] showed recently that for certain pairwise subset ranking losses ℓ , finding a predictor that minimizes the ℓ -risk is an NP-hard problem. They also showed that several common pairwise convex surrogate losses that operate on $\widehat{\mathcal{T}} = \mathbb{R}^r$ (and are used to learn scores for the r documents) fail to be classification calibrated with respect to such losses ℓ , even under some low-noise conditions on the distribution, and proposed an alternative convex surrogate, also operating on $\widehat{\mathcal{T}} = \mathbb{R}^r$, that is classification calibrated under certain conditions on the distribution (i.e. over a strict subset of the associated probability simplex).

Here we provide an alternative route to analyzing the difficulty of obtaining consistent surrogates for such pairwise subset ranking problems using the classification calibration dimension. Specifically, we will show that even for a simple setting of such problems, the classification calibration dimension of the underlying loss ℓ is greater than r , and therefore no convex surrogate operating on $\widehat{\mathcal{T}} \subseteq \mathbb{R}^r$ can be classification calibrated w.r.t. such a loss over the full probability simplex.

Formally, we will identify the set of class labels \mathcal{Y} with a set \mathcal{G} of ‘preference graphs’, which are simply directed acyclic graphs (DAGs) over r vertices; for each directed edge (i, j) in a preference graph $\mathbf{g} \in \mathcal{G}$ associated with an instance $x \in \mathcal{X}$, the i -th document in the document set in x is preferred over the j -th document. Here we will consider a simple setting where each preference graph has exactly one edge, so that $|\mathcal{Y}| = |\mathcal{G}| = r(r-1)$; in this setting, we can associate each $\mathbf{g} \in \mathcal{G}$ with the edge (i, j) it contains, which we will write as $\mathbf{g}_{(i,j)}$. The target labels consist of permutations over r objects, so that $\mathcal{T} = S_r$ with $|\mathcal{T}| = r!$. Consider now the following simple pairwise loss $\ell^{\text{pair}} : \mathcal{Y} \times \mathcal{T} \rightarrow \mathbb{R}_+$:

$$\ell^{\text{pair}}(\mathbf{g}_{(i,j)}, \sigma) = \mathbf{1}(\sigma(i) > \sigma(j)). \quad (7)$$

Let $\mathbf{p} = (\frac{1}{r(r-1)}, \dots, \frac{1}{r(r-1)})^\top \in \text{relint}(\Delta_{r(r-1)})$, and observe that $\mathbf{p}^\top \ell_\sigma^{\text{pair}} = \frac{1}{2}$ for all $\sigma \in \mathcal{T}$. Thus $\mathbf{p}^\top (\ell_{\sigma'}^{\text{pair}} - \ell_{\sigma''}^{\text{pair}}) = 0 \forall \sigma, \sigma' \in \mathcal{T}$, and so $\mathbf{p} \in \mathcal{Q}_\sigma^{\text{pair}} \forall \sigma \in \mathcal{T}$.

Let $(\sigma_1, \dots, \sigma_{r!})$ be any fixed ordering of the permutations in \mathcal{T} , and consider $\mathcal{Q}_{\sigma_1}^{\text{pair}}$, defined by the intersection of $r! - 1$ half-spaces of the form $\mathbf{q}^\top (\ell_{\sigma_1}^{\text{pair}} - \ell_{\sigma_t}^{\text{pair}}) \leq 0$ for $t = 2, \dots, r!$ and the simplex constraints $\mathbf{q} \in \Delta_{r(r-1)}$. Moreover, from the above observation, $\mathbf{p} \in \mathcal{Q}_{\sigma_1}^{\text{pair}}$ satisfies $\mathbf{p}^\top (\ell_{\sigma_1}^{\text{pair}} - \ell_{\sigma_t}^{\text{pair}}) = 0 \forall t = 2, \dots, r!$. Therefore, by Lemma 14, we get

$$\mu_{\mathcal{Q}_{\sigma_1}^{\text{pair}}}(\mathbf{p}) = \text{nullity}([\ell_{\sigma_1}^{\text{pair}} - \ell_{\sigma_2}^{\text{pair}}, \dots, \ell_{\sigma_1}^{\text{pair}} - \ell_{\sigma_{r!}}^{\text{pair}}, \mathbf{e}]^\top), \quad (8)$$

where \mathbf{e} is the $r(r-1) \times 1$ all ones vector. It is not hard to see that the set $\{\ell_\sigma^{\text{pair}} : \sigma \in \mathcal{T}\}$ spans a $\frac{r(r-1)}{2}$ dimensional space, and hence the nullity of the above matrix is at most $r(r-1) - (\frac{r(r-1)}{2} - 1)$. Thus by Theorem 13, we get that $\text{CCdim}(\ell^{\text{pair}}) \geq r(r-1) - (\frac{r(r-1)}{2} + 1) - 1 = \frac{r(r-1)}{2} - 2$. In particular, for $r \geq 5$, this gives $\text{CCdim}(\ell^{\text{pair}}) > r$, and therefore establishes that no convex surrogate ψ operating on a surrogate target space $\widehat{\mathcal{T}} \subseteq \mathbb{R}^r$ can be classification calibrated with respect to ℓ^{pair} on the full probability simplex $\Delta_{r(r-1)}$.

6 Conclusion

We developed a framework for analyzing consistency for general multiclass learning problems defined by a general loss matrix, introduced the notion of classification calibration dimension of a multiclass loss, and used this to analyze consistency properties of surrogate losses for various general multiclass problems. An interesting direction would be to develop a generic procedure for designing consistent convex surrogates operating on a ‘minimal’ surrogate target space according to the classification calibration dimension of the loss matrix. It would also be of interest to extend the results here to account for noise conditions as in [9, 10].

Acknowledgments

We would like to thank the anonymous reviewers for helpful comments. HG thanks Microsoft Research India for a travel grant. This research is supported in part by a Ramanujan Fellowship to SA from DST and an Indo-US Joint Center Award from the Indo-US Science & Technology Forum.

References

- [1] Gábor Lugosi and Nicolas Vayatis. On the bayes-risk consistency of regularized boosting methods. *Annals of Statistics*, 32(1):30–55, 2004.
- [2] Wenxin Jiang. Process consistency for AdaBoost. *Annals of Statistics*, 32(1):13–29, 2004.
- [3] Tong Zhang. Statistical behavior and consistency of classification methods based on convex risk minimization. *Annals of Statistics*, 32(1):56–134, 2004.
- [4] Ingo Steinwart. Consistency of support vector machines and other regularized kernel classifiers. *IEEE Transactions on Information Theory*, 51(1):128–142, 2005.
- [5] Peter Bartlett, Michael Jordan, and Jon McAuliffe. Convexity, classification and risk bounds. *Journal of the American Statistical Association*, 101(473):138–156, 2006.
- [6] Tong Zhang. Statistical analysis of some multi-category large margin classification methods. *Journal of Machine Learning Research*, 5:1225–1251, 2004.
- [7] Ambuj Tewari and Peter Bartlett. On the consistency of multiclass classification methods. *Journal of Machine Learning Research*, 8:1007–1025, 2007.
- [8] David Cossock and Tong Zhang. Statistical analysis of bayes optimal subset ranking. *IEEE Transactions on Information Theory*, 54(11):5140–5154, 2008.
- [9] Fen Xia, Tie-Yan Liu, Jue Wang, Wensheng Zhang, and Hang Li. Listwise approach to learning to rank: Theory and algorithm. In *International Conference on Machine Learning*, 2008.
- [10] John Duchi, Lester Mackey, and Michael Jordan. On the consistency of ranking algorithms. In *International Conference on Machine Learning*, 2010.
- [11] Pradeep Ravikumar, Ambuj Tewari, and Eunho Yang. On NDCG consistency of listwise ranking methods. In *International Conference on Artificial Intelligence and Statistics(AISTATS)*, volume 15. JMLR: W&CP, 2011.
- [12] David Buffoni, Clément Calauzènes, Patrick Gallinari, and Nicolas Usunier. Learning scoring functions with order-preserving losses and standardized supervision. In *International Conference on Machine Learning*, 2011.
- [13] Wei Gao and Zhi-Hua Zhou. On the consistency of multi-label learning. In *Conference on Learning Theory*, 2011.
- [14] Wojciech Kotłowski, Krzysztof Dembczynski, and Eyke Huellermeier. Bipartite ranking through minimization of univariate loss. In *International Conference on Machine Learning*, 2011.
- [15] Ingo Steinwart. How to compare different loss functions and their risks. *Constructive Approximation*, 26:225–287, 2007.
- [16] Ben Taskar, Carlos Guestrin, and Daphne Koller. Max-margin Markov networks. In *Neural Information Processing Systems*, 2003.
- [17] Deirdre O’Brien, Maya Gupta, and Robert Gray. Cost-sensitive multi-class classification from probability estimates. In *International Conference on Machine Learning*, 2008.
- [18] Nicolas Lambert and Yoav Shoham. Eliciting truthful answers to multiple-choice questions. In *ACM Conference on Electronic Commerce*, 2009.
- [19] Dimitri Bertsekas, Angelia Nedic, and Asuman Ozdaglar. *Convex Analysis and Optimization*. Athena Scientific, 2003.
- [20] Jean Gallier. Notes on convex sets, polytopes, polyhedra, combinatorial topology, Voronoi diagrams and Delaunay triangulations. Technical report, Department of Computer and Information Science, University of Pennsylvania, 2009.
- [21] Elodie Vernet, Robert C. Williamson, and Mark D. Reid. Composite multiclass losses. In *Neural Information Processing Systems*, 2011.