

# Multiclass ROC

**Liang Wang**

*Department of Mathematics and Statistics  
Boston University  
Boston, MA, 02215 USA*

LEONWANG@BU.EDU

**Luis Carvalho**

*Department of Mathematics and Statistics  
Boston University  
Boston, MA, 02215 USA*

LECARVAL@MATH.BU.EDU

**Editor:**

## Abstract

Model evaluation is of crucial importance in modern statistics application. The construction of ROC and calculation of AUC have been widely used for binary classification evaluation. Recent research generalizing the ROC/AUC analysis to multi-class classification has problems in at least one of the four areas: 1. failure to provide sensible plots 2. being sensitive to imbalanced data 3. unable to specify mis-classification cost and 4. unable to provide evaluation uncertainty quantification. Borrowing from a binomial matrix factorization model, we provide an evaluation metric summarizing the pair-wise multi-class True Positive Rate (TPR) and False Positive Rate (FPR) with one-dimensional vector representation. Visualization on the representation vector measures the relative speed of increment between TPR and FPR across all the classes pairs, which in turns provides a ROC plot for the multi-class counterpart. An integration over those factorized vector provides a binary AUC-equivalent summary on the classifier performance. Mis-classification weights specification and bootstrapped confidence interval are also enabled to accommodate a variety of evaluation criteria. To support our findings, we conducted extensive simulation studies and compared our method to the pair-wise averaged AUC statistics on benchmark datasets.

**Keywords:** ROC curve, AUC, classification, model evaluation, matrix factorization.

## 1. Introduction

For machine learning classification problems, the evaluation of a model is of crucial importance since it provides guidance on the choice of an ultimate models. In many fields of the evolving research, multi-class classification is of the interests. According to the its inputs, the evaluation metric can be generally divided into two categories.

### 1.1 Hard Classification Metric

The first sets of naive evaluation metrics is based upon "hard classification" metrics where the input is typically how many number of observations are classified correctly/incorrectly:

- True Positive(TP): classification result is **positive** and the ground truth is **positive**.
- False Positive(FP): classification result is **positive** but the ground truth is **negative**.

- True Negative(TN): classification result is **negative** and the ground truth is **negative**.
- False Negative(FN): classification result is **negative** and the ground truth is **positive**.

From which, one can define the True Positive Rate (TPR), True Negative Rate (TNR), False Positive Rate (FPR) and False Negative Rate Rate(FNR):

$$\begin{aligned} \text{TPR} &= \frac{\text{Total True Positive}}{\text{Total Positive}}, \text{FPR} = \frac{\text{Total False Positive}}{\text{Total Negative}} \\ \text{TNR} &= \frac{\text{Total True Negative}}{\text{Total Negative}}, \text{FNR} = \frac{\text{Total False Negative}}{\text{Total Positive}} \end{aligned}$$

For a balanced classification problem, practitioners can highlight the model’s capability to correctly identified either positive labels or the negative labels. With this criteria, the TPR and TNR are popularly adopted. Focusing on the percentage of correctly identified positive labels, the TPR is also termed as recall and sensitivity. Focusing on the percentage of correctly identified negative labels, the TNR is also termed as specificity. However, when we are dealing with imbalanced problem, focusing solely on the positive samples or the negative samples will mislead the evaluation to favour the majority classifier.

Under the imbalanced classification criteria, each of the above number captures part of the classifier performance, the commonly used classifier evaluations thus usually consider a combination of the four sub-evaluation metrics with various weighting scheme. For example, the Cohen’s Kappa (Cohen, 1960) and Matthews Correlation Coefficient (MCC) (Matthews, 1975) can be considered respectively as the geometric mean and harmonic mean of

$$\frac{TP \cdot TN - FP \cdot FN}{(TP + FP) \cdot (FP + TN)} \quad \text{and} \quad \frac{TP \cdot TN - FP \cdot FN}{(TP + FN) \cdot (FN + TN)}$$

The  $F_1$  score can be considered as the harmonic mean of the precision and recall, i.e:

$$\frac{TP}{TP + FP} \quad \text{and} \quad \frac{TP}{TP + FN}$$

Empirical comparisons of those commonly used evaluation metrics for binary classification have been demonstrated in (Chicco and Jurman, 2020), and (Chicco et al., 2021). With numerical examples, it is also concluded in those literature that MCC is more informative than Cohen’s Kappa and Brier score.

(Powers, 2010) also shown that the MCC can also be understood as the geometric mean of two regression coefficient of Markedness and Informedness. (Gorodkin, 2004) exploits this correlation coefficient relationship implied by binary MCC regression coefficients to generalize the MCC to a general mutli-class classification evaluation metrics. Although MCC can be considered as an ultimate choice of the "hard classification" metrics, it has a few shortcomings. The first one is that since the MCC is based upon the count number instead of the percentage rate, its performance deteriorates seriously when the classification labels are imbalanced(Zhu, 2020). Secondly, despite the fact that practitioner not only wants to know the classification result but also hope to understand how confident is the evaluation conclusion, the MCC statistics ignores the uncertainty behind the classification evaluation by simply counting the number of the incorrectness outputted by the model.

Lastly, summarizing the classifier’s performance with a single number overlooks many other useful information on the classifier that are available with visualization plot. For example, one might also want to get a sense of TPR and FPR performance with different level of confidence(threshold) from the evaluation as it is emphasized in (Drummond and Holte, 2006).

## 1.2 Soft Classification Metric

Taking probabilities as the input to the metrics, a better alternative ”Receiver Operating Characteristic (ROC) (Metz, 1978)” conducts judgement based on the ”soft” classification result. The method not only provides the access to a comprehensive visualization against the random classifier but also circumvents the class skewness problem by directly taking the percentage rate (TPR and FPR) into consideration. If one wants a comprehensive score to summarize the model performance, the Area Under the Curve (AUC) score can be computed using numerical integration. From a statistics perspective, the AUC score has also been demonstrated to be equivalent to the probability of the event that a positive labeled observation will receive a higher probability ranking as it is compared to the negative labeled observation (Provost and Domingos, 2000). This probability interpretation has embarked on a series of follow-up research in Mann-Whitney U-Statistic.

Though remains the benchmark metric for modern binary classification problems, one disadvantage of the AUC method is its plausible equal assumption on the misclassification costs. In practice, it is of a rarity and thus often problematic to assume that nothing about the relative severity of misclassification is known. An alternative AUC taking advantage of the relative severity of the binary classification costs has been proposed in (Adams and Hand, 1999). Other methods such as (Provost and Fawcett, 1998) overcome this issue by constructing the ROC convex hull after the ROC analysis, through which the slope of isocost analysis can be applied to conduct the classifier evaluation. However, both methods dealing with the misclassification cost are within the binary classification setup.

Generalizing the AUC/ROC construction to a multi-class classification is challenging. To summarize some recent attempts, (Mossman, 1999) and (Ferri et al., 2003) generalize the multi-class AUC through Volume Under the Surface (VUS) but the methods suffer from expensive computational budget and none of those methods can provide sensible visualization. Despite those drawbacks, there are generalizations of AUC statistics that received great popularity. For example, (Provost and Domingos, 2000) approaches the generalization problem by taking a one-versus-all-wise average of the binary AUC. (Hand and Till, 2001) adopt a similar average idea on the pair-wise binary AUC and demonstrated that the pair-wise AUC average preserves the class-skewness invariance property of the binary AUC. Employing the Mann-Whitney U-Statistic setup, (Ross and Page, 2019) generalizes the AUC statistics through a partition definition of the order statistics. Although both (Ross and Page, 2019)’s and (Provost and Domingos, 2000)’s methods are shown to be invariant to class-skewness, their methods discarded the ROC visualization and thus can only summarize the classifier performance with a scalar value for absolute comparison. Additionally, both of the methods lack of the capability of specifying the relative severity of misclassification costs, which is absolutely needed in practice to determine optimal classifier(Adams and Hand, 1999). Lastly, neither of those methods nor the binary AUC provides an uncer-

tainty quantification on the evaluation conclusion. All those shortcomings substantiate a conclusion that better multi-class AUC method that is as well-accepted as the binary AUC method is yet to be discovered.

### 1.3 Overall Contribution and Plans of the Paper

Using a binomial DMF, we propose a new multi-class equivalent ROC and AUC for model evaluation. Our method borrows from the pair-wise AUC design (Provost and Domingos, 2000) by constructing a True Positive Rate Matrix  $\mathbb{M}^{tp}$  and False Positive Rate Matrix  $\mathbb{M}^{fp}$  among all pairs of the classes. Then we center the factorized components by taking the mean of the two factorized matrices. Due to a flexible factorization weight, the factorization model also entitles the relative severity specification through factorization weights. Moreover, we demonstrate that the rank-one factorized components can be re-normalized to be of range  $[0, 1]$  and thus be interpreted as the shared threshold effects implied by the pair-wise averaged AUC. The normalized vector provides not only a multi-class equivalent ROC plot but also an AUC score that summarizes the comprehensive performance of the classifiers. As a summary of the advantages of our factorization evaluation metrics over the existing multi-class evaluation metric:

- The new metric provides access to a visualization plot that compares the classifier under different confidence(threshold) circumstances.
- The new metric is invariant to the label imbalance within the classification problem.
- The new metric enables the specification of relative severity for misclassification.
- The new metric quantifies the evaluation uncertainty with a reasonable confidence interval

Our paper is organized as follow: In Section 2, we provide some background information on the pair-wise multi-class AUC analysis. In Section 3, we formally introduce our evaluation metrics and provide some analyses on its desiring property. In Section 4, we summarize the implementation into an algorithm to facilitate applications. In Section 5, we provides both simulation and real-data examples to demonstrate the desired properties mentioned previously. In Section 6 we conclude the paper with a discussion and potential improvements.

## 2. Backgrounds on Pair-wise AUC

The pair-wise AUC design firstly connects the numerical integration to a ranking statistics and then generalizes the ranking idea toward a multi-class setup.

### 2.1 AUC and Mann Whitney U Statistics

For a binary classification outcome with index  $i$  and class label  $Y_i$  taking binary value of 0 or 1, a classifier usually takes some covariates  $X_i$  as input and outputs a simplex score vector  $(S_i, 1 - S_i)$ . The score  $S_i$  is usually taken to be a specific probability:  $S_i \equiv \mathbb{P}_i \equiv \mathbb{P}(Y_i = 0)$ , which represents the probability for the  $i$ -th observation to belong to class 0. Given a

specific probability threshold  $t$ , one can thus classify the  $i$ -th observation to class 0 if  $\mathbf{S}_i$  is **greater than** the specified threshold  $t$ . For an evaluation result with  $n$  observations, the specifications of threshold  $t$  will determine a "hard" classification assignment, which can then be used to compute evaluation metrics such as True Positive Rate (TPR denoted by  $f(t)$ ) and a False Positive Rate (FPR denoted by  $g(t)$ ). Both  $f(t)$  and  $g(t)$  are non-increasing functions with respect to threshold  $t$ , which implies  $dg(t) \leq 0, df(t) \leq 0$  ( $d$  here stands for differentiation operation). To summarize the performance of classifiers under different threshold circumstances, a ROC curve plots the TPR  $f(t)$  against FPR  $g(t)$  with a decreasing threshold  $t$ .

Based upon the constructed ROC curve, one can compute a more comprehensive score through numerical integration of  $f(t)$  over  $g(t)$  and obtain the Area Under the Curve (AUC). A perfect classifier will thus have a constant TPR  $f(t) = 1$  and increasing FPR  $g(t)$  with respect to decreasing threshold  $t$ , giving an AUC value equal to 1. While a random classifier will have TPR  $f(t)$  and FPR  $g(t)$  increases at the same pace with respect to decreasing threshold  $t$ , giving an AUC value of 0.5. The relationship between  $f(t)$  and  $g(t)$  thus determine the quality of the classification model.

The AUC statistics has an interesting connection to Mann-Whitney U statistics. To observe the connection, denote  $\{i \in I_k\}$  as the sets that the  $i$ -th observation belongs to class  $k$ , we could express the TPR and FPR with the following probability representation:

$$f(t) = \mathbb{P}(S_i > t | i \in I_0) = 1 - \mathbb{P}(S_i \leq t | i \in I_0)$$

$$g(t) = \mathbb{P}(S_i > t | i \in I_1) = 1 - \mathbb{P}(S_i \leq t | i \in I_1)$$

The AUC statistics as a numerical integration can then be written in the following form:

$$\begin{aligned} A(0,1) &= \int_{g(t)=0}^{g(t)=1} f(t) dg(t) \\ &= \int_{g(t)=0}^{g(t)=1} \mathbb{P}(S_i \leq t | i \in I_1) d(1 - \mathbb{P}(S_i \leq t | i \in I_0)) \\ &= - \int_{t=1}^{t=0} \mathbb{P}(S_i \leq t | i \in I_1) \mathbb{P}(S_i = t | i \in I_0) dt \\ &= \mathbb{P}[S_i 1_{\{X_i \in I_1\}} \leq S_i 1_{\{X_i \in I_0\}}] \end{aligned} \tag{1}$$

Where the integration interval of the third equality is ranged from  $[1, 0]$  due to the fact  $\frac{dg(t)}{dt} \leq 0$ . The mathematical interpretation in Eq (1) provides us with two methods of computing the AUC statistics:

- If we interpret the AUC statistics as the probability that the class 0 score of label 0 samples is higher than the class 0 score of label 1 samples, the AUC statistics can be estimated by using non-parametric Mann-Whitney U statistic.
- If we instead model parametrically the relationship between  $\mathbb{P}(S_i \leq t | i \in I_1)$  and  $\mathbb{P}(S_i \leq t | i \in I_0)$ , we could numerically integrate the AUC score with a fine numerical grid on threshold  $t \in [0, 1]$ .

Specifically by following the first probability interpretation, we can calculate the AUC score using the Mann-Whitney U statistic. That is we can denote the probability of being class 0 of  $n$  different samples using a vector  $\vec{S} = (S_1, S_2, \dots, S_n)$ . Then we can sort the score in an ascending order and obtain its order statistics  $\vec{R} = (R_1, R_2, \dots, R_n)$  with true class label  $\vec{Y} = (Y_1, Y_2, \dots, Y_n)$ . For the  $i$ -th observation from class 0, the number of points from class 1 with a score lower than the points from class 0 is given by:

$$U(0, 1) = \sum_{i \in I_0} (R_i - i) \quad (2)$$

Denote the number of points from class 0 ( $|\{i \in I_0\}|$ ) as  $n_0$ , the number of points from class 1 ( $|\{i \in I_1\}|$ ) as  $n_1$ . It is obvious to conclude that there are in total  $n_0 \times n_1$  possible sample elements of statistics  $U(0, 1)$ , which gives the following estimate on the probability of the ranking:

$$\hat{A}(0, 1) = \frac{U(0, 1)}{n_0 \times n_1} \quad (3)$$

The AUC score  $A(0, 1)$  is indexed with  $(0, 1)$  to emphasize that we are using the class 0 score (probability of the observation being class 0) to compare the observations from class 1 against the observations from class 0. For binary classification, the order of the index does not matter by having  $A(0, 1) = A(1, 0)$ , but this equality in general does not hold if we generalize to a multi-class classification problem. (Hand and Till, 2001) adopted a pair-wise design to reduce the multi-class classification to multiple binary classifications.

## 2.2 Generalizing AUC to Multiclass via Pair-wise Specification

The design of the sub-binary problem is constructed by focusing only on two out of the multiple classification labels. To illustrate the pair-wise AUC concept for multi-class classification problem, we provide the following diagram:

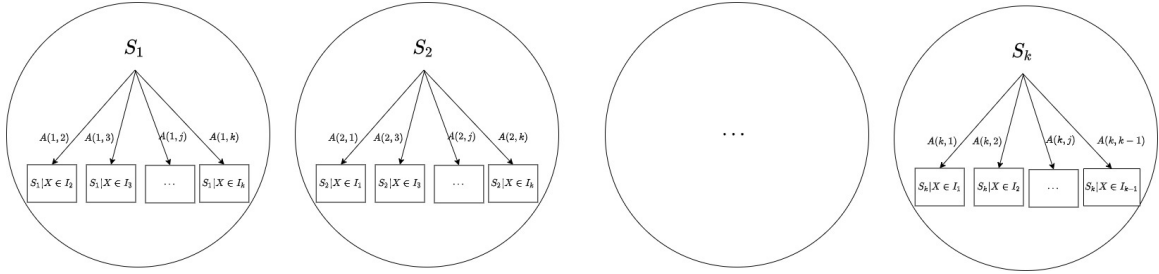


Figure 1: **Pair-wise AUC partition**

As it is indicated in Figure 1, if  $S_l$  is interpreted as classification probability for label  $l$ , we can potentially have  $k$  different scores being defined with a simplex relationship:  $\sum_{l=1}^k S_l = 1$ . Then for each of the  $k$  different scores, we can construct  $k - 1$  pairs of score comparison, concluding  $K = k(k - 1)$  pairs of the pair-wise AUC score for comparison. Consequently, for each index  $l$  with range  $1, \dots, k(k - 1)$ , it corresponds to a comparison index  $(i, j), i \neq j, i, j \in [1, \dots, k]$ . Given a pair-wise index  $l$  or equivalently the pair  $(i, j)$ , we can formally define the cardinality of class 0 and class 1 as:

$$n_0(l) \equiv n_0(i, j) = |\{\text{obs} \in I_i\}| \quad n_1(l) \equiv n_1(i, j) = |\{\text{obs} \in I_j\}| \quad (4)$$

where obs stands for observations,  $|\cdot|$  stands for the cardinality of a set, and  $I_i$  stands for class  $i$  assignment using the ground truth label. For each pair index  $(i, j)$ , its score can either be computed using numerical integration given in Eq (1) or using the Mann-Whitney U formula given in Eq (3). (Hand and Till, 2001) proposed the following score to measure the classifier performance:

$$\hat{\mathcal{M}} = \frac{1}{k(k-1)} \sum_{i=1}^k \sum_{j \neq i, j=1}^k \hat{A}(i, j) \quad (5)$$

Eq (5) is invariant to class-skewness since it can be considered as a pair-wise average of all binary AUC statistics, which are themselves invariant to class skewness. Equivalently, the  $\mathcal{M}$  statistics can be considered as an equal-weighted expectation of

$$\mathcal{M} = \mathbb{E}_{\{(i,j) \in \mathcal{K}\}}[A(i, j)] = \mathbb{E}_k \left[ \int f_k(t) \Delta g_k(t) dt \right] = \mathbb{E}_k \mathbb{E}_t [f_k(t) \Delta g_k(t)]$$

The statistics implicitly assume that  $f_k(t)$  and  $g_k(t)$  are given deterministic numbers and thus provide a point-wise estimate. However, similar to the soft and hard classification from MCC to AUC or from k-means to the Gaussian mixture, it is always of interest to investigate the behavior of the statistics if we generalize  $f_k(t)$  and  $g_k(t)$  to follow a random setup. This random generalization can additionally provide us with an uncertainty quantification through bootstrapped confidence interval. For the binary evaluation, (Balaswamy and Vishnu Vardhan, 2015) imposed a Generalized Half Normal assumption on  $f_k(t)$  and  $g_k(t)$  and proposed a corresponding confidence interval to the AUC analysis. In the next section, we will explore the possibility of generalizing  $f_k(t), g_k(t)$  to be random under binomial factorization setup.

### 3. A Joint Binomial Factorization Model

In this section, we demonstrate that step-by-step how we construct our multi-class ROC curve and how to compute its corresponding AUC statistics.

#### 3.1 An Example of Binary Matrix Construction

Although we aim at a generalization toward multi-class ROC curve, it is intuitive for us to start with the description of the binary counterpart. Given a decreasing threshold  $\mathbb{T} = \{t_1, t_2, \dots, t_T\}$  and unitary space  $\mathbb{Z} = [0, 1]$ , we start by computing the True Positive Rate (TPR) vector  $\mathbb{M}^{tp}(0, 1) \in \mathbb{Z}^{T \times 1}$  and False Positive Rate (FPR) vector  $\mathbb{M}^{fp}(0, 1) \in \mathbb{Z}^{T \times 1}$  by referencing the 1st class(class 0) as the positive class and the 2nd class(class 1) as the negative class:

$$\mathbb{M}^{tp}(0, 1) = \begin{pmatrix} f_{11} \\ f_{21} \\ \vdots \\ f_{T1} \end{pmatrix} \quad \mathbb{M}^{fp}(0, 1) = \begin{pmatrix} g_{11} \\ g_{21} \\ \vdots \\ g_{T1} \end{pmatrix} \quad (6)$$

In contract, if we reference the 2nd class(class 1) as the negative class and the 1st class(class 0) as the positive class, we can compute another set of vectors  $\mathbb{M}^{fp}(0, 1) \in \mathbb{Z}^{T \times 1}$  and

$\mathbb{M}^{tp}(0, 1) \in \mathbb{Z}^{T \times 1}$ :

$$\mathbb{M}^{tp}(1, 0) = \begin{pmatrix} f_{12} \\ f_{22} \\ \vdots \\ f_{T2} \end{pmatrix} \quad \mathbb{M}^{fp}(1, 0) = \begin{pmatrix} g_{12} \\ g_{22} \\ \vdots \\ g_{T2} \end{pmatrix} \quad (7)$$

Due to the label symmetry in binary classification problem, we in fact have

$$\begin{aligned} 1 &= f_{t1} + g_{t2}, \forall t \in \mathbb{T} \\ 1 &= g_{t1} + f_{t2}, \forall t \in \mathbb{T} \end{aligned} \quad (8)$$

We can concatenate  $\mathbb{M}^{tp}(0, 1)$  and  $\mathbb{M}^{tp}(1, 1)$  together into matrices  $\mathbb{M}^{tp} \in \mathbb{Z}^{T \times 2}$  and  $\mathbb{M}^{fp} \in \mathbb{Z}^{T \times 2}$  and obtain:

$$\begin{aligned} \mathbb{M}^{tp} &= \left( \mathbb{M}^{tp}(0, 1), \mathbb{M}^{tp}(1, 0) \right) = \begin{pmatrix} f_{11} & f_{12} \\ f_{21} & f_{22} \\ \vdots & \vdots \\ f_{T1} & f_{T2} \end{pmatrix} \\ \mathbb{M}^{fp} &= \left( \mathbb{M}^{fp}(0, 1), \mathbb{M}^{fp}(1, 0) \right) = \begin{pmatrix} g_{11} & g_{12} \\ g_{21} & g_{22} \\ \vdots & \vdots \\ g_{T1} & g_{T2} \end{pmatrix} \end{aligned} \quad (9)$$

Due to this special relationship of binary classification problem in Eq (8), the ROC curve plots  $f_{t1}$  against  $g_{t1}$  to measure the model performance, which is equivalent to plotting  $1 - g_{t2}$  against  $1 - f_{t2}$ . Hence, researchers have avoid the redundancy of computing two *FPR* and *TPR* by focusing on one choice of positive label assignment. This interesting property is however not preserved as we generalize toward the multi-class setup. Based upon  $\mathbb{M}^{tp}$  and  $\mathbb{M}^{fp}$ , we can take either the first column vector or the second column vector of  $\mathbb{M}^{tp}$  and  $\mathbb{M}^{fp}$  and further concatenate them together row-wisely to obtain  $\mathbb{M}$ :

$$\mathbb{M}_i \equiv \begin{pmatrix} \mathbb{M}_{*0}^{tp} \\ \mathbb{M}_{*0}^{fp} \end{pmatrix} \equiv \left( f_{11}, f_{21}, \dots, f_{T1}, g_{11}, g_{21}, \dots, g_{T1} \right)^\top \in \mathbb{Z}^{2T \times 1}$$

and assume:

$$\begin{pmatrix} \mathbb{M}_{i0}^{tp} \\ \mathbb{M}_{i1}^{fp} \end{pmatrix} \circ \begin{pmatrix} n_1 \\ n_0 \end{pmatrix} \stackrel{ind}{\sim} \text{Bin} \left( \mu_i = \begin{pmatrix} m(\Lambda_{i0}^{tp}) \\ m(\Lambda_{i0}^{fp}) \end{pmatrix}, \begin{pmatrix} n_1 \\ n_0 \end{pmatrix} \right)$$

Such a model is saturated with  $2T$  parameterization vector via  $(\Lambda_0^{tp}, \Lambda_0^{fp})$ , which provides the equivalence of  $m(\Lambda_{i0}) = \mathbb{M}_{i0}$ . The plotting of  $m(\Lambda_{i0}^{tp})$  against  $m(\Lambda_{i0}^{fp})$  is hence equivalent to the ROC plotting of  $m(\Lambda_{i0}^{tp})$  against  $m(\Lambda_{i0}^{fp})$ .

### 3.2 Construction of a Pair-wise Matrix

To generalize the binary case toward the multi-class set up, we operate on the same decreasing threshold  $\mathbb{T} = \{t_1, t_2, \dots, t_T\}$  and binary space  $\mathbb{Z} = [0, 1]$ . We start by building two matrices  $\mathbb{M}^{tp} \in \mathbb{Z}^{T \times k(k-1)}$  and  $\mathbb{M}^{fp} \in \mathbb{Z}^{T \times k(k-1)}$  where the two matrices store all pairs-wise (specified in Figure 1) TPR and FPR with respect to a decreasing threshold smoothed by the binary probability quantile  $\mathbb{T}$ . That is, for each pair-wise binary problem, we filter the rows the predicted probability matrix by focusing on observations whose ground truth label is either class  $i$  or class  $j$ , based upon which we compute a smoothed threshold grids using the quantiles from the predicted probability matrix of those two classes:

$$\mathbb{M}^{tp} = \begin{pmatrix} \boxed{f_{11}} & f_{12} & \cdots & f_{1K} \\ \boxed{f_{21}} & f_{22} & \cdots & f_{2K} \\ \vdots & \vdots & \ddots & \vdots \\ \boxed{f_{T1}} & f_{T2} & \cdots & f_{TK} \end{pmatrix} \xrightarrow[\text{AUC}(0,1)]{\mathbb{M}^{fp}} \begin{pmatrix} \boxed{g_{11}} & g_{12} & \cdots & g_{1K} \\ \boxed{g_{21}} & g_{22} & \cdots & g_{2K} \\ \vdots & \vdots & \ddots & \vdots \\ \boxed{g_{T1}} & g_{T2} & \cdots & g_{TK} \end{pmatrix} \quad (10)$$

Row-wisely, for a given  $i = 1, 2, \dots, T$ ,  $f_{i*}, g_{i*} \in [0, 1]^{k(k-1)}$  correspond to the TPR and FPR given threshold  $t_i$  across all the  $k(k-1)$  pair-wise construction. Column-wisely, for a given  $j = 1, 2, \dots, k(k-1) \equiv K$ ,  $f_{*j}, g_{*j} \in [0, 1]^T$  correspond to the  $j$ -th pair-wise TPR and FPR.  $\text{AUC}(0, 1)$  is thus a numerical integration of the first  $\mathbb{M}^{tp}$  column with respect to the first  $\mathbb{M}^{fp}$  column. Those two matrices contain valuable information for classification evaluation. To effectively summarize  $\mathbb{M}^{tp}$  and  $\mathbb{M}^{fp}$  using low dimensional vectors, we firstly emphasize the following observed properties of those two matrices:

- Row-wisely, both  $\mathbb{M}^{tp}$  and  $\mathbb{M}^{fp}$  are monotonic in the sense that its value is increasing with respect to the decrease of the threshold  $\mathbb{T}$ .
- Row-wisely,  $\mathbb{M}^{tp}$  and  $\mathbb{M}^{fp}$  increases at different speeds. How fast does  $\mathbb{M}^{tp}$  increases according to the increases  $\mathbb{M}^{fp}$  determines the quality of the classifier.
- Column-wisely, both  $\mathbb{M}^{tp}$  and  $\mathbb{M}^{fp}$  share the same pair-wise specification and thus should have similar column effects.

Due to this simple monotonic property of the  $\mathbb{M}^{tp}$  and  $\mathbb{M}^{fp}$ , it is naturally hopeful to summarize the "speed of increment" across the pair-wise columns through an appropriately designed matrix factorization.

### 3.3 The Factorization Model

The Deviance Matrix Factorization(DMF)(Wang and Carvalho, 2023) supports a binomial parametric assumption on the matrix. A follow-up research (Wang and Carvalho, 2024) studied the bayesian version of such a model with computational improvement. With  $Z^{tp}$  and  $Z^{fp}$  representing the count of true positive and the count of false positive, we can exploit the following data generating assumption to parameterize the multi-class ROC curve:

$$\begin{aligned} Z_{ij}^{tp} &\equiv \mathbb{M}_{ij}^{tp} W_j^{pos} \stackrel{ind}{\sim} \text{Bin}(\mu_{ij}^{tp} = m(\eta_{ij}^{tp}) | \eta_{ij}^{tp} = \Lambda_{i0}^{tp} + (\Lambda_i^{tp})^\top V_j, W_j^{pos}) \\ Z_{ij}^{fp} &\equiv \mathbb{M}_{ij}^{fp} W_j^{neg} \stackrel{ind}{\sim} \text{Bin}(\mu_{ij}^{fp} = m(\eta_{ij}^{fp}) | \eta_{ij}^{fp} = \Lambda_{i0}^{fp} + (\Lambda_i^{fp})^\top V_j, W_j^{neg}) \end{aligned} \quad (11)$$

Concatenate those two matrices and abbreviate the notations with

$$Z \equiv \begin{pmatrix} \mathbb{Z}^{tp} \\ \mathbb{Z}^{fp} \end{pmatrix}, \mathbb{M} \equiv \begin{pmatrix} \mathbb{M}^{tp} \\ \mathbb{M}^{fp} \end{pmatrix}, \Lambda_0 \equiv \begin{pmatrix} \Lambda_0^{tp} \\ \Lambda_0^{fp} \end{pmatrix}, \Lambda \equiv \begin{pmatrix} \Lambda^{tp} \\ \Lambda^{fp} \end{pmatrix}, \eta \equiv \begin{pmatrix} \eta^{tp} \\ \eta^{fp} \end{pmatrix}, W \equiv \begin{pmatrix} W^{pos} \\ W^{neg} \end{pmatrix}$$

The data generating process is equivalently:

$$Z_{ij} \equiv \mathbb{M}_{ij} W_j \stackrel{ind}{\sim} \text{Bin}(\mu_{ij} = m(\eta_{ij}) | \eta_{ij} = \Lambda_{i0} + \Lambda_i^\top V_j, W_j)$$

To appropriately account for the observations we made in Section 3.2, we impose the following constraint on the factorized components  $\Lambda$  and  $V$ .

- Row-wisely, The increment  $\mathbb{M}^{tp}$  and  $\mathbb{M}^{fp}$  can be decomposed into shared threshold effect (modeled by  $\Lambda_{i0}$ ) and columns specific effect  $\Lambda_i^\top V_j$ .  $\Lambda_{i0}$ s are monotonic with respect to the decrease of the threshold and are independent (constructed to be orthogonal) of the columns-specific effect  $\Lambda_i^\top V_j$ .
- With factorization rank  $q = 1$ ,  $\mathbb{M}^{tp}$  and  $\mathbb{M}^{fp}$  increases at different speed with the implicit relationship  $\eta_i^{tp} = \alpha_i + \beta_i \eta_i^{fp}$  where  $\alpha_i = \Lambda_{i0}^{tp} - \frac{\Lambda_i^{tp}}{\Lambda_i^{fp}} \Lambda_{i0}^{fp}$  and  $\beta_i = \frac{\Lambda_i^{tp}}{\Lambda_i^{fp}}$
- Column-wisely,  $\mathbb{M}^{tp}$  and  $\mathbb{M}^{fp}$  share the same column effect by having the same column matrix  $V$ . But within the columns of  $\mathbb{M}^{tp}$  and  $\mathbb{M}^{fp}$ , the increment should be independent.

Mathematically, this is equivalent to solving the low-rank representation  $(\hat{\Lambda}, \hat{V})$  by maximizing the binomial likelihood  $\mathcal{L}$  with link  $g(\cdot)$  assumption:

$$\underset{\substack{\Lambda_0 \in \mathbb{R}^T, V^\top \mathbf{1}_K = 0 \\ (\Lambda, V) \in \{\mathcal{V}_{T,K}(q)\}}}{\text{argmax}} W \circ \mathcal{L}(\mathbb{M}, \Lambda_0 \mathbf{1}_K^\top + \Lambda V^\top) = \underset{\substack{\Lambda_0 \in \mathbb{R}^T, V^\top \mathbf{1}_K = 0 \\ \Lambda \in \tilde{\mathcal{S}}_{T,q}, V \in \mathcal{S}_{K,q}}}{\text{argmax}} W \circ \mathcal{L}(\mathbb{M}, \Lambda_0 \mathbf{1}_K^\top + \Lambda V^\top), \quad (12)$$

where

- $q$  is the reduced dimension and is chosen to be one to maintain  $\eta_i^{tp} = \alpha_i + \beta_i \eta_i^{fp}$ :
- $W$  is the binomial weight for the observation of true positive count  $Z^{tp}$  and false positive count  $Z^{fp}$ .
- $V$  has orthogonal columns and are orthogonal to  $\mathbf{1}_K$ :  $V^\top V = I_q, V^\top \mathbf{1}_K = 0$ , i.e  $V$  belongs to the centered Stiefel manifold  $\mathcal{S}_{T,q}$ ;
- $\Lambda$  has pairwise orthogonal columns, that is,  $\Lambda = UD$  with  $U \in \mathcal{S}_{T,q}$  and  $D = \text{Diag}\{d_1, \dots, d_q\}$  with  $d_1 \geq \dots \geq d_q$  so that  $\Lambda^\top \Lambda = D^2$ . This space for  $\Lambda$  is denoted as  $\tilde{\mathcal{S}}_{T,q}$ .

This model can thus be solved through a centering trick after obtaining a full rank factorization of  $\Lambda \in \tilde{\mathcal{S}}_{T,K}, V \in \mathcal{S}_{K,K}$  with weight  $W_{ij}$ . That is we center the factorized components through a projector matrix  $H_0$  defined according to  $V_0$  with  $V_0$  being the vector of one.

$$\begin{aligned} \Lambda V^\top &= \Lambda V^\top H_0 + \Lambda V^\top (I_K - H_0) \\ &= \underbrace{\Lambda V^\top V_0 (V_0^\top V_0)^{-1} V_0^\top}_{\Lambda_0} + \underbrace{\Lambda V^\top (I_K - H_0)}_{\tilde{\Lambda} \tilde{V}^\top} \end{aligned} \quad (13)$$

The reduced component  $\Lambda_0, V_0$  are of rank one and are orthogonal to the newly constructed  $\tilde{\Lambda}\tilde{V}^\top$  through the construction of projector  $H_0$ . In the case of  $q = 1$ , we have  $\Lambda_0 = \sum_{j=1}^K \frac{1}{K} \Lambda$ , which aims at averaging the pair-wise TPR and FPR across  $K$  different classes on the centered Stiefel manifold given by  $(\mathbf{1}_K, V^\top)$ .  $\Lambda_0$  thus has the interpretation as the averaged TPR/FPR across the  $k(k-1)$  pair-wise groups. In the special case of  $K = 1$ , we have precisely the binary ROC equivalence by having  $V = \mathbb{I}$  as shown in Section 3.1. The model setup in fact is closely related to summary ROC analysis in the literature with details shown in Section 3.5.

### 3.4 Specifying Pair-wise Misclassification Cost

As demonstrated in (Metz, 1978), (Provost and Fawcett, 1998), and (Adams and Hand, 1999), ignoring the misclassification cost can sometimes provide misleading results for model selection. In fact, it has been shown in the same references that an optimal classification threshold can be obtained only if the misclassification cost is specified. For the binary classification evaluation, the choice of optimal threshold on the ROC curve is formulated as an optimization problem given both the classification cost and the prior probability of class occurrence. Specifically, given the misclassification costs  $C(0), C(1)$ , the prior probability of class occurrence  $\pi(0), \pi(1)$  and a threshold  $t$ , one can compute the probability of misclassifying class  $k$  as  $f_t(k)$  and define a threshold  $t$  dependent evaluation loss function:

$$L_t = C(0)f_t(0)\pi(0) + C(1)f_t(1)\pi(1)$$

To minimize the loss with respect to threshold  $t$ , the slope defined by the iso-cost line  $\frac{\pi(0)C(0)}{\pi(1)C(1)}$  is shown to be tangent to the optimal threshold  $t$  on the ROC curve. Since only the tangent points are the possible optimal, the analysis is also termed ROC Convex Hull (ROCCH). However, for a pair-wise multi-class AUC setup, the convex hull analysis is hardly applicable due to the absence of an equivalent ROC curve.

Notice that by default, this optimization setup in Eq (12) weight the observation of  $\mathbb{M}_{ij}^{pos}$  and  $\mathbb{M}_{ij}^{neg}$  according to the cardinality of the positive and negative samples defined by the binary pair  $W_j = n_0(j) \times n_1(j)$  where index  $j$  is defined as the  $j$ -th pair-wise index as illustrated in Figure 1. In some of the existing literature, the cardinality of positive samples and negative samples are also referenced as the prior (occurrence) probability of the pair  $(i, j)$ .

Although the default cardinality weighting is consistent with the intuition that  $\mathbb{M}_{ij}^{pos}$  and  $\mathbb{M}_{ij}^{neg}$  are observed with higher confidence due to large bernoulli sample size, this weighting scheme favours the majority classifier by putting more factorization weights on the majority binary pairs. To flexibly allow more intuitive misclassification cost, we introduce additional weight  $Q_{ij}$  specification to the factorization, which changes the optimization setup from Eq (12) to Eq (14):

$$\hat{\Lambda}\hat{V}^\top(Q) = \underset{\substack{\Lambda_0 \in \mathbb{R}^T, V^\top \mathbf{1}_K = 0 \\ \Lambda \in \tilde{\mathcal{S}}_{T,q}, V \in \mathcal{S}_{K,q}}}{\operatorname{argmax}} \sum_{i=1}^{2T} \sum_{j=1}^K Q_{ij} W_{ij} \mathcal{L}_{ij}(M, \Lambda_0 \mathbf{1}_K^\top + \Lambda V^\top), \quad (14)$$

With this additionally factorization weight  $Q_{ij}$ , the evaluation metric can be specified as proportional to the relative severity of the misclassification. For example, one can specify

$Q_j = \frac{1}{W_j}$  if the practitioners want the evaluation metric to be absolutely independent from label imbalance. This special choice of  $Q_j = \frac{1}{W_j}$  is referenced as unweighted DMF-ROC later since it weights each binary pair equally. As for the difference between  $Q_{ij}$  and  $W_{ij}$ , we emphasize that  $Q_{ij}$  is introduced only in the optimization while  $W_{ij}$  is assumed to appear in the data generating process of TPR/FPR.

This specification enables the third advantage of our AUC metrics because it circumvents the construction of multiple ROC convex hulls by directly taking the weights specification into TPR-like and FPR-like factorization. As we show in Section 5.3, this weight specification  $Q_{ij}$  flexibly allows misclassification cost to effectively affect the model evaluation results as we expected.

### 3.5 Connection to sROC

In fact, if we adopt a logit link function and a factorization rank  $q = 1$ , our factorization model can be considered as a multi-class generalization of the binary *summary ROC* (sROC) analysis (Moses et al., 1993; Arends et al., 2008), which has been additionally shown to connect with ordinal regression in (Tosteson and Begg, 1988).

The sROC basically transforms the original TPR and FPR to a linear continuous space through logit link function and analyzes the incremental speed of TPR relative to FPR through a regression coefficient. For given threshold  $i$  of a binary classification problem  $j = K = 1$ , the sROC estimated  $a, \beta$  according to the following setup:

$$\begin{aligned} D_{ij} &= \text{logit}(\mathbb{M}_{ij}^{tp}) - \text{logit}(\mathbb{M}_{ij}^{fp}) \\ E_{ij} &= \text{logit}(\mathbb{M}_{ij}^{tp}) + \text{logit}(\mathbb{M}_{ij}^{fp}) \\ D_{ij} &= \alpha_j + \beta E_{ij} + e_{ij} \end{aligned} \tag{15}$$

where  $e_{ij}$  is a random error and is orthogonal to the linear predictor. After estimating the coefficient  $b$  based upon discretely observed  $\mathbb{M}_{ij}^{tp}$  and  $\mathbb{M}_{ij}^{fp}$ , we can reconstruct  $\mathbb{M}_{ij}^{tp}$  and  $\mathbb{M}_{ij}^{fp}$  according to a finer numerical grid of  $E_{ij}$  and thus construct a smoothed ROC curve. Moreover, once the distribution of  $e_{ij}$  is known and its variance parameter is estimated, one can construct multiple ROC curves and thus its confidence interval through a naive bootstrapping on estimated parameters.

In (Reitsma et al., 2005), it is emphasized that  $\mathbb{M}_{ij}^{tp}$  and  $\mathbb{M}_{ij}^{fp}$  are negatively correlated across index threshold  $i$  because those two metrics share the same threshold specification. Lowering the cutoff value will then lead to more observation with a positive result, thereby increasing the number of true positives but also the number of false positive results. To take this potential negative correlation into consideration, a bivariate model is proposed for binary classification problem  $j = K = 1$ :

$$\begin{aligned} \begin{pmatrix} \mu_{ij}^{tp} \\ \mu_{ij}^{fp} \end{pmatrix} &\overset{i.i.d}{\sim} N\left(\begin{pmatrix} \mu_j^{tp} \\ \mu_j^{fp} \end{pmatrix}, \Sigma_1\right) \\ \begin{pmatrix} \text{logit}(\mathbb{M}_{ij}^{tp} | \mu_{ij}^{tp}) \\ \text{logit}(\mathbb{M}_{ij}^{fp} | \mu_{ij}^{tp}) \end{pmatrix} &\overset{i.i.d}{\sim} N\left(\begin{pmatrix} \mu_{ij}^{tp} \\ \mu_{ij}^{fp} \end{pmatrix}, \Sigma_2\right) \end{aligned} \tag{16}$$

where the parameter  $\Sigma_1, \Sigma_2$  are used to model the negative correlation. Alternatively (Rutter and Gatsonis, 2001) parametrized the ROC curve via binomial random effect model:

$$\begin{aligned} \text{logit} \begin{pmatrix} \mu_{ij}^{tp} \\ \mu_{ij}^{fp} \end{pmatrix} &\sim N \left( \begin{pmatrix} \eta_j^{tp} \\ \eta_j^{fp} \end{pmatrix}, \Sigma_1 \right) \\ \begin{pmatrix} M_{ij}^{tp} \\ M_{ij}^{fp} \end{pmatrix} \circ \begin{pmatrix} W_{ij}^{pos} \\ W_{ij}^{neg} \end{pmatrix} &\sim \text{Bin} \left( \begin{pmatrix} \mu_{ij}^{tp} \\ \mu_{ij}^{fp} \end{pmatrix}, \begin{pmatrix} W_{ij}^{pos} \\ W_{ij}^{neg} \end{pmatrix} \right) \end{aligned} \quad (17)$$

where the correlation between  $M_{ij}^{tp}$  and  $M_{ij}^{fp}$  are modeled through  $\Sigma_1$  but the computation is more expensive as it is compared to the Gaussian assumption in Eq (16). The setup is more close to our model in Eq (11).

Notice that all the sROC introduced concentrates on binary classification problems. Hence the sROC parameterization models the correlation of the threshold effect row-wisely since there is no pair effect. If we generalize the sROC curve parameterization to multi-class classification problem, the column-wise pair-wise correlation should be far more important to model than the row-wise threshold effect. The intuition is that we have  $K = k(k-1)$  columns composed of  $k$  classes, which makes the ability to distinguish different binary pairs at the same threshold to be highly correlated. This observation implies an index change of the binomial ROC index from Eq (17) to:

$$\begin{aligned} \text{logit} \begin{pmatrix} \mu_{ij}^{tp} \\ \mu_{ij}^{fp} \end{pmatrix} &\sim N \left( \begin{pmatrix} \eta_i^{tp} \\ \eta_i^{fp} \end{pmatrix}, \Sigma_1 \right) \\ \begin{pmatrix} M_{ij}^{tp} \\ M_{ij}^{fp} \end{pmatrix} \circ \begin{pmatrix} W_{ij}^{pos} \\ W_{ij}^{neg} \end{pmatrix} &\sim \text{Bin} \left( \begin{pmatrix} \mu_{ij}^{tp} \\ \mu_{ij}^{fp} \end{pmatrix}, \begin{pmatrix} W_{ij}^{pos} \\ W_{ij}^{neg} \end{pmatrix} \right), \forall j \in [K] \end{aligned} \quad (18)$$

The setup however still suffers computationally, our factorization model setup in Eq (11) can then be considered as an approximation to Eq (18) because conditional on threshold  $i$ , the pair-wise ( $j \in [1, \dots, K]$ ) correlation is modeled through covariance structure  $\Sigma_1$ .

More specifically, Eq (11) assumes a common  $V_j$  for both  $\eta_{ij}^{tp}$  and  $\eta_{ij}^{fp}$ :

$$\text{logit}(\mu_{ij}^{tp}) = \Lambda_{i0}^{tp} + (\Lambda_i^{tp})^\top V_j \quad (19)$$

$$\text{logit}(\mu_{ij}^{fp}) = \Lambda_{i0}^{fp} + (\Lambda_i^{fp})^\top V_j \quad (20)$$

with factorization rank  $q = 1$ , we have scalar  $\Lambda_i, V_j \in \mathbb{R}$ , which implies:

$$\begin{aligned} \text{logit}(\mu_{ij}^{tp}) &= \underbrace{\frac{\Lambda_i^{tp}}{\Lambda_i^{fp}} \text{logit}(\mu_{ij}^{fp})}_{\beta_i} + \underbrace{\Lambda_{i0}^{tp} - \frac{\Lambda_i^{tp}}{\Lambda_i^{fp}} \Lambda_{i0}^{fp}}_{\alpha_i}, \quad \forall j \in [K] \\ \text{logit}(\mu_i^{tp}) &= \alpha_i + \beta_i^\top \text{logit}(\mu_i^{fp}) \end{aligned} \quad (21)$$

That is our setup implies a linear relationship between  $\text{logit}(\mu_{ij}^{tp})$  and  $\text{logit}(\mu_{ij}^{fp})$  whose mean are modeled via  $(\eta_i^{tp}, \eta_i^{fp})$  and whose covariance  $\Sigma_1$  is modeled through parameter  $\Lambda_i, V^\top$ .

### 3.6 Visualization and Confidence Interval

With the interpretation in Section 3.5, we know that how fast does  $\Lambda_0^{tp}$  increase with respect to  $\Lambda_0^{fp}$  represent the quality of the classifier. As a result, after summarizing  $\mathbb{M}^{tp}$  and  $\mathbb{M}^{fp}$  through a binomial rank one representation, we could simply plot the center of the factorized components, i. e, plot  $\Lambda_0^{tp}$  against  $\Lambda_0^{fp}$  to visually inspect how fast on average does the row of  $\mathbb{M}^{tp}$  increases with respect to the row of  $\mathbb{M}^{fp}$ . To mimic the ROC plot in the binary classification problem, we further normalize the factorized components  $\Lambda_0^{tp}$  and  $\Lambda_0^{fp}$  through a monotonic link function  $m(\cdot) = \text{logit}^{-1}(\cdot)$  to provide a multi-class equivalent ROC plot.

The plot thus naturally determines the quality of the classifier with extensive empirical evidence given in Section 5. Like the construction of the AUC statistics, we could also summarize the classifier's performance with statistics  $\mathcal{D}$  through integration on the centered components:

$$\mathcal{D} \equiv \int \text{logit}^{-1}(\Lambda_0^{tp}) d\text{logit}^{-1}(\Lambda_0^{fp}) \quad (22)$$

Additionally, since  $\mathbb{M}_{ij} \stackrel{\text{ind}}{\sim} \text{Bin}(\mu_{ij} = m(\eta_{ij}) | \eta_{ij} = \Lambda_i V_j^\top, W_{ij})$  is implicitly assumed from the factorization, we can simulate  $\mathbb{M}^{tp}$  and  $\mathbb{M}^{fp}$  after obtaining an estimation of  $\eta^{tp}$  and  $\eta^{fp}$  with weight  $W_j$ . The simulation naturally provides the access to a confidence interval around the ROC curve and thus the confidence interval of  $\mathcal{D}$  statistics.

## 4. Algorithm ROC-DMF

Since the analysis in the previous section indicates a solid statistical connection to sROC using a factorization rank of 1 ( $q = 1$ ), we provide efficient vectorization considering the normal equation reduces to a scalar update when  $q = 1$ .

The factorization is conducted by concatenating the matrix  $\mathbb{M} = \begin{pmatrix} \mathbb{M}^{tp} \\ \mathbb{M}^{fp} \end{pmatrix}$  with  $\mathbb{M} \in [0, 1]^{2T \times k(k-1)}$ . The factorized latent space is of the same dimension with  $\eta \in \mathbb{R}^{2T \times k(k-1)}$ . The factorized components  $\Lambda$  is a vector of dimension  $2T$ :  $\Lambda \in \tilde{\mathcal{S}}_{2T,1}$  and the basis matrix  $V$  is also a vector of dimension  $k(k-1)$ :  $V \in \mathcal{S}_{k(k-1),1}$ . We denote the  $t$ -th scalar of  $\Lambda$  as  $\lambda_t$  and the  $j$ -th scalar of  $V$  as  $v_j$ ;  $\forall i = 1, \dots, 2T, j = 1, \dots, k(k-1)$ . With those notations, we can iterate the  $\lambda$  and  $v$  update by solving the following system of the equation:

$$\begin{aligned} V^\top D_i V \lambda_i^{(t+1)} &= V^\top D_i (V \lambda_i^{(t)} + D_i^{-1} G_i) \doteq V^\top D_i Z_i^{(t)} \\ \Lambda^\top D_{\cdot j} \Lambda v_j^{(t+1)} &= \Lambda^\top D_{\cdot j} (\Lambda v_j^{(t)} + D_{\cdot j}^{-1} G_{\cdot j}) \doteq \Lambda^\top D_{\cdot j} Z_{\cdot j}^{(t)}, \end{aligned} \quad (23)$$

where  $S_{ij}$ ,  $G_{ij}$  and  $Z_{ij}$  are defined as:

- $W_j = n_0(j) \times n_1(j) \times Q_j$ ,  $\eta_{ij} = \Lambda_i V_j$
- $g(\cdot) = \text{logit}(\cdot)$ ,  $\mathcal{V}(\mu_{ij}) = \mu_{ij}(1 - \mu_{ij})$ ,  $\mu_{ij} = g^{-1}(\eta_{ij})$
- $S_{ij} = w_{ij} \frac{g^{-1'}(\eta_{ij})^2}{\mathcal{V}(\mu_{ij})}$  and  $G_{ij} = \frac{g^{-1'}(\eta_{ij})}{\mathcal{V}(\mu_{ij})} w_{ij} (X_{ij} - \mu_{ij})$
- $D_i \doteq \text{Diag}\{S_i^{(t)}\}$  with  $S_i$  denotes the  $i$ -th row of  $S$ . Similarly,  $\text{Diag}\{S_{\cdot j}^{(t)}\}$  with  $S_{\cdot j}$  denotes the  $j$ -th column of  $S$ .

- $Z^{(t)}$  is the working response:

$$Z_{ij}^{(t)} = \eta_{ij}^{(t)} + \frac{G_{ij}^{(t)}}{S_{ij}^{(t)}} = \eta_{ij}^{(t)} + \frac{X_{ij} - \mu_{ij}^{(t)}}{g^{-1'}(\eta_{ij}^{(t)})}. \quad (24)$$

Vectorizing Eq (23) to obtain the following definition:

$$\begin{aligned} \underline{V} &= \sqrt{S} \circ \left( \underbrace{[V, \dots, V]}_{2T \text{ times}} \right)^\top \\ \underline{\Lambda} &= \sqrt{S} \circ \left[ \underbrace{\Lambda, \dots, \Lambda}_{k(k-1) \text{ times}} \right] \\ \underline{Z} &= Z \circ S \end{aligned} \quad (25)$$

$\underline{V}$ ,  $\underline{\Lambda}$  and  $\underline{Z}$  are matrix of dimension  $2T \times k(k-1)$ . With element-wise division  $\oslash$  and square  $(\cdot)^2$ , the updates become:

$$\begin{aligned} \Lambda^{(t+1)} &= V^\top \underline{Z}^\top \oslash \mathbf{1}_p^\top \underline{V}^2 \\ V^{(t+1)} &= \Lambda^\top \underline{Z} \oslash \mathbf{1}_n^\top \underline{\Lambda}^2 \end{aligned} \quad (26)$$

To facilitate the implementation, we summarize both the construction of the matrix  $\mathbb{M}$  and the optimization steps in Algorithm 1.

---

**Algorithm 1: Multi-class AUC**


---

**Definition:** Number of classes  $k$ , number of classification instance  $n$ .  
**Input:** Classification probability  $\mathbb{P} \in [0, 1]^{n \times k}$ , factorization weight  $Q_j = \mathbb{R}_+^K$ , true class label  $Y \in [1, \dots, k]^n$ , threshold quantile  $\mathbb{T} = \{\tau_1, \dots, \tau_T\}$  with  $0 \leq \tau_1 < \dots < \tau_T \leq 1$ .  
**for**  $i = 1, 2, \dots, k$  **do**  $\triangleright$  construct  $\mathbb{M}^{tp}, \mathbb{M}^{fp}$  according to Eq (10)  
     Extract probability vector from the  $i$ -th column of  $\mathbb{P}$ :  $\mathbb{P}_i \equiv \mathbb{P}(Y = i) = \mathbb{P}_{*i}$   
     Compute the pair-wise threshold according to quantile  $\mathbb{T}$  of  $\mathbb{P}_i$ .  
     Define class 0 samples  $Y(0) = \{Y_l : Y_l \in I_i\}_{l=1}^n$   
     **for**  $j = 1, 2, \dots, k, j \neq i$  **do**  
         Define class 1 samples  $Y(1) = \{Y_l : Y_l \in I_j\}_{l=1}^n$   
         **for**  $t = 1, 2, \dots, T$  **do**  
             Compute the  $\mathbb{M}^{tp}$  element according to  $\tau_t$  quantile of  $\mathbb{P}_i$   
             Compute the  $\mathbb{M}^{fp}$  element according to  $\tau_t$  quantile of  $\mathbb{P}_i$   
         Set  $W_j = n_0(j) \times n_1(j)$  as in Eq (4).  
**Factorization** on  $\mathbb{M} = [\mathbb{M}^{tp}, \mathbb{M}^{fp}]^\top$  with weight  $W_j = W_j \times Q_j$ :  
**for**  $i = 1, \dots, 2T$  and  $j = 1, \dots, k(k-1)$  **do**  $\triangleright$  initialization  
     Set  $\mu_{ij}^{(0)}$  as a perturbed version of  $\mathbb{M}_{ij}$  and  $\eta_{ij}^{(0)} = g(\mu_{ij}^{(0)})$ ;  
     Set  $L^{(0)}$  as the first  $q$  columns of the  $j$  index  $[\eta_{ij}^{(0)}]_{i=1, \dots, 2T; j=1, \dots, k(k-1)}^\top$ ;  
     **for**  $t = 0, 1, \dots$  (until convergence) **do**  $\triangleright$  vectorized updates  
         Compute  $\underline{\Lambda}, \underline{V}, \underline{Z}$  as in Eq (25);  
         Update  $V$  and  $\Lambda$  according to Eq (23);  
     Obtain  $\Lambda_0^{tp}, \Lambda_0^{fp}$  through the projection on  $(\hat{\Lambda}, \hat{V}^\top)$  DNF result using Eq (13)  
**Result:** Binomial factorized  $\Lambda_0^{tp}$  and  $\Lambda_0^{fp}$  and metric  $\mathcal{D}$  according to Eq (22)

---

## 5. Experiments

To support our findings, we conducted simulation studies to demonstrate that 1. our method can successfully distinguish the classifiers of different performance quality, 2. our method is invariant to class skewness, 3. our ROC plot can reasonably take the weights specification into consideration, and 4. our ROC plot can provide reasonable confidence interval to quantify the evaluation uncertainty.

### 5.1 Discriminative Experiment

To define a well-understood multi-class classification problem, we simulated multinomial regression data conditional on known covariates  $X \in \mathbb{R}^{n \times p}$  and known coefficients  $B \in \mathbb{R}^{p \times k}$ . Specifically, the covariate  $X$  and coefficients  $B$  are indexed with sample size  $n = 50,000$  and dimension  $p = 10$ . Their values are fixed after a row-wise/column-wise random sample from  $X_i \sim N(0, I_p), B_j \sim N(1, I_p), \forall i = 1, \dots, n, \forall j = 1, \dots, k-1$ . The number of class  $k$  is chosen to be 5 as a starting example. The classification assignment probability is obtained

conditional on multinomial regression coefficient  $B \in \mathbb{R}^{p \times (k-1)}$ :

$$\vec{\mathbb{P}}_i = (\mathbb{P}(Y_i = 1), \mathbb{P}(Y_i = 2), \dots, \mathbb{P}(Y_i = k)) = \text{Softmax}((1, X_i B))$$

The ground truth label  $Y_i$  can then be obtained according to a multinomial random sample with probability  $\vec{\mathbb{P}}_i$ :

$$Y_i \sim \text{Multinomial}(\vec{\mathbb{P}}_i) \quad (27)$$

Or to make the setup even simpler, we can also simply assign the label according to the highest element of simulated vector  $\vec{\mathbb{P}}_i$ :

$$Y_i = \underset{i=1, \dots, k}{\operatorname{argmax}} \vec{\mathbb{P}}_i \quad (28)$$

Taking  $\vec{\mathbb{P}}_i$  as the output of a classification model, the label setup in (28) thus corresponds to a perfect classifier since we are essentially fitting the label according to the probability assignment. The label setup in (27) will correspond to a more realistic classifier due to the random multinomial assignment from  $\vec{\mathbb{P}}_i$ .

With fixed random seed, the number of observations associated with each class from Eq (27) is correspondingly 9,534, 6,042, 11,586, 12,957, and 9,881. The number of observations of each class from Eq (28) is correspondingly 9471, 5,111, 12,159, 13,543, 9,716. Below, we plotted the histogram of the generated class label:

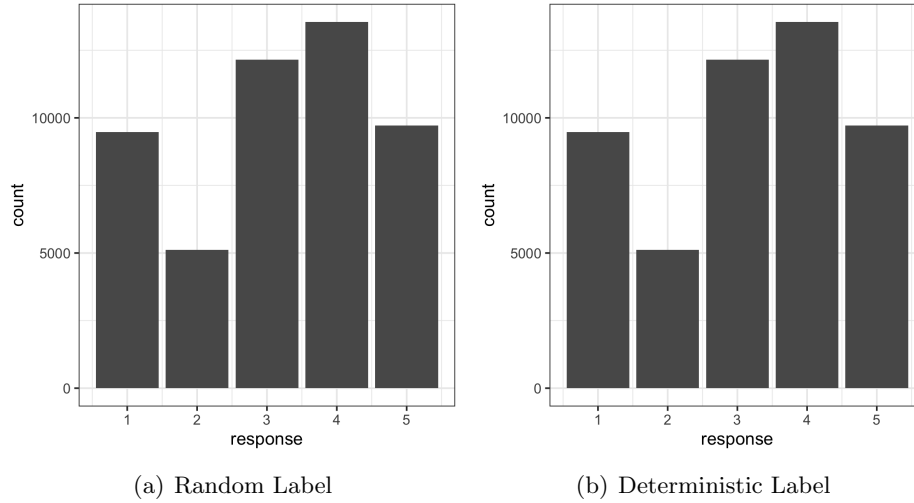


Figure 2: **Histogram of simulated class labels**

Then to create classifiers of different performance quality, we train classifiers with partial availability of the generating covariance  $X$ . That is, we gradually decrease the availability of the dimension of covariate  $X$  to obtain classifiers with inferior prediction capability. Specifically, we fit a multinomial regression with only the first  $d$  ( $d = 1, 2, \dots, p$ ) columns of

$X$ , denoted as  $X^d$ .

$$X = \begin{bmatrix} \underbrace{X_{*1}, \dots, X_{*d}}_{X^d}, \dots, X_{*p} \end{bmatrix}$$

To benchmark different levels' classifiers against the random classifier, we also provide one classifier with another independent simulation of  $\tilde{X}_i \sim N(0, I_p), i = 1, 2, \dots, n$ . Predicting the outcome conditional on  $\tilde{X}$  is considered as a random classifier because we know  $\tilde{X} \perp\!\!\!\perp Y$  from the generating process.

Lastly, for each of the fitted multinomial regression conditional on  $X^d$ , we can obtain a fitted probability  $\hat{\mathbb{P}}_i^d = \mathbb{P}(Y_i | X_i^d B^d) \in [0, 1]^k$  with  $B^d \in \mathbb{R}^{d \times k}$ . We then input both the true class label  $Y$  and the fitted probability  $\hat{\mathbb{P}}^d$  into our evaluation Algorithm 1 with  $Q_j = \frac{1}{n_0 \times n_1}$  to weight each binary pair equally. From the output the algorithm, we can easily plot  $\text{logit}^{-1}(\Lambda_0^{tp})$  against  $\text{logit}^{-1}(\Lambda_0^{tp})$  to construct the ROC curves. The constructed curves are provided below:

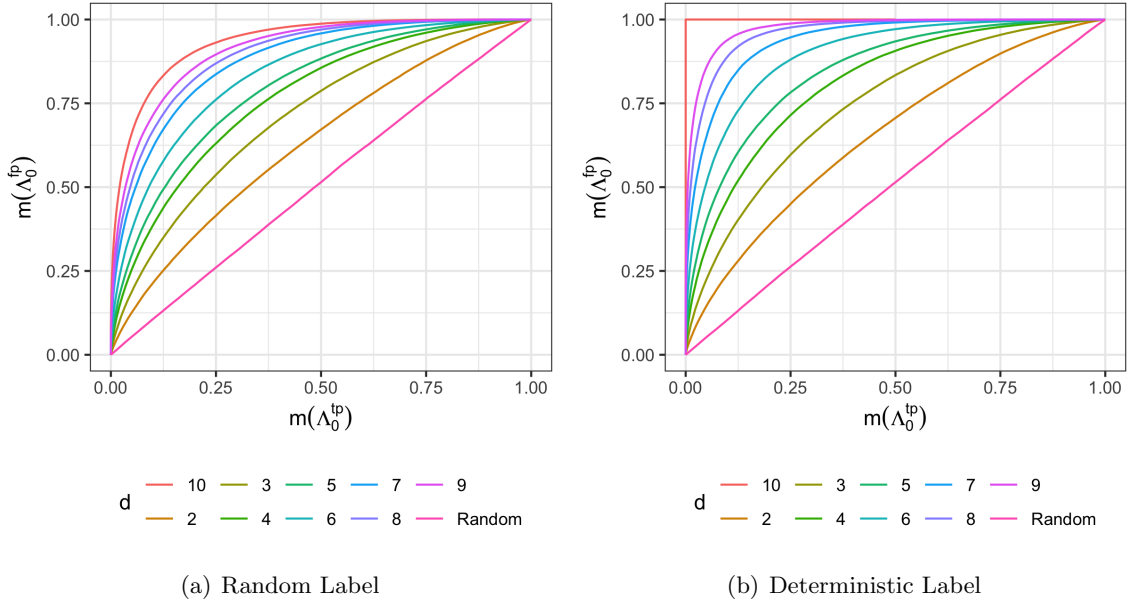


Figure 3: **Discriminative experiment**

As we can see from Figure 3, the optimal classifier with the full availability of  $X^p$  is plotted as the top-left curve with the highest AUC value. The classifier's performance deteriorates as we gradually decrease the availability of the covariate information defined by parameter  $d$ . The ROC curve of deterministic label (Eq (28)) is higher than the one of sampling labels (Eq (27)) with the best performance ROC = 1 given the full availability of predictor. The result provides the first evidence that our DMF factorization provides

effective discriminative ROC curve for multi-class classification. To further investigate the invariance to class-skewness property, we provide more empirical examples in the next subsection.

## 5.2 Class-skewness Experiment

From the simulation setup, we could obtain an imbalanced classification problem defined by  $(Y_n, X_n)$ . To make the classification problem more realistic, we take the random label assignment of Eq (27) as the ground truth label. Then to create a simulation example with a balanced class assignment, we simply sample an equal number  $n_{sub} = 6,042$  of observations within each class. Based upon this balanced simulation  $Y_{n_{sub}}, X_{n_{sub}}$ , we further create different level of skewness by sampling disproportionately the observations from  $k$  classes according to a simplex vector  $\vec{w} \equiv (w_1, w_2, \dots, w_k)$ .

The vector  $\vec{w}$  is repeatedly simulated from Dirichlet distribution  $\vec{w} \sim \text{Dir}(\vec{\alpha})$  with parameter  $\vec{\alpha} \equiv (\alpha_1, \alpha_2, \dots, \alpha_k)$ . For simplicity, we choose  $\alpha_1 = \alpha_2 = \dots = \alpha_k = \alpha$ . To investigate the impact of parameter  $\alpha$  on simulated skewness, we quantify the data skewness using the ratios among the Dirichlet-sampled weights  $\frac{w_i}{w_j}, \forall i \neq j$ . To summarize the skewness of the  $k$  different classes, we denote random variable  $Z = \max_{i \neq j} \frac{w_i}{w_j}$  as the surrogate of the overall class-skewness. The random variable  $Z$  can be understood as the ratio between the cardinality of majority class and the cardinality of minority class. As we can conclude from the setup, the higher the  $Z$  value is, the higher the class skewness is within the experiment dataset. We also empirically estimate and visualize (in Figure 4) the tail probability  $P(Z > c|\alpha)$  with 10,000 independent Dirichlet random sampling:

$$\hat{P}(Z > c|\alpha) = \frac{\sum_{i=1}^{10,000} 1_{\{Z_i > c\}}}{10,000}$$

To investigate the impact of parameter  $\alpha$  on level of class skewness under different number of classes ( $k$ ), we also calculated the probability estimate with number of classes  $k = 10$ :

As it is indicated by Figure 4, the lower the  $\alpha$  is, the higher is the ratio between majority and minority class (and thus the higher is the class-skewness). The skewness is also notably higher when the number of classes increase ( $k = 10$ ). To illustrate the skewness impact across different quality of the classifiers, we also experimented on three classifiers with parameter  $d = 2, 5, 9$  to represent classifiers of different level of performance. That is, for each of the three classifiers with different covariate availability ( $d = 2, 5, 9$ ), we sample disproportionately according to Dirichlet-generated weight vectors  $\vec{w} \sim \text{Dir}(\alpha)$ . The skewness parameter  $\alpha$  is varied among 2, 5, 9, based upon which, the sampling is repeated 30 times to create  $3 \times 30 = 90$  imbalanced testing sets for classifier performance evaluation.

To have a glimpse about the weight specification effect, we provided the visualization plots for both "weighted" and "unweighted" ROC-DMF. In the "unweighted" ROC-DMF plot, the factorization cost is universally equal by having  $Q_j = \frac{1}{W_j}, W_j = n_0(j) \times n_1(j), \forall j = 1, \dots, k(k-1)$  with  $n_0(j), n_1(j)$  defined in Eq (4). In the "weighted" ROC-DMF, we applied factorization weight according to the cardinality (or prior occurrence probability) of the true labels, i.e.  $\forall j = 1, \dots, k(k-1), Q_j = 1, W_j = n_0 \times n_1$ .

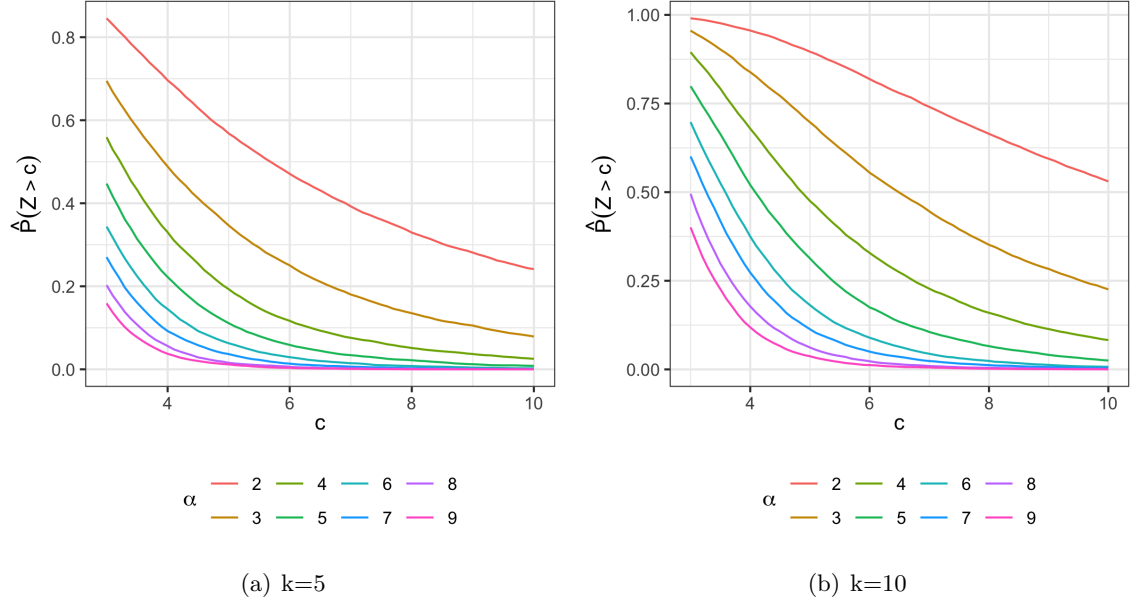
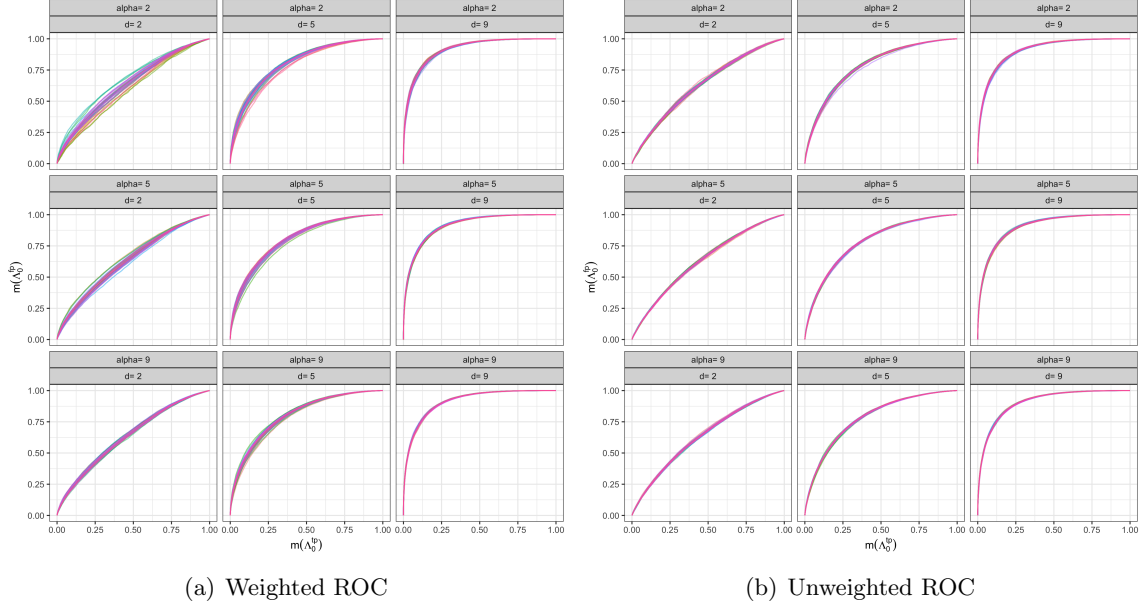

 Figure 4:  $\alpha$  impact on class-skewness


Figure 5: Class-skewness experiment

From Figure 5, we can conclude that for a given level of classifier defined by the availability covariate dimension  $d$ , the unweighted ROC curve barely varies. Because the weighted ROC result is obtained through prior probability weighted factorization, there are more variation associated with the low-quality classifiers. Comparing the curves row-wisely ac-

cording to different skewness level, we can conclude that our constructed ROC curve is invariant to class-skewness. To amplify the weight specification effect, we continue our experiments in the next section.

### 5.3 Weight Specification Experiment

To examine the impact of weights specification on the ROC plot, we created a naive soft classifier that universally predicts the majority-class with dominated probability. In this naive classifier, all observations universally have the highest probability of belonging to the majority class (e.g. class four from our original 50,000 examples). With equal misclassification cost, the created naive classifier should be considered a bad classifier since it universally predicts all the observations to be the majority class. However, this majority classifier might not be a bad classifier if

1. the cost of misclassifying the majority class into other classes is low.
2. the cost of misclassifying other classes into the majority class is high.
3. the reward of correctly classifying the majority class is high.
4. the reward of correctly classifying the other classes is low.

In fact, naively classifying all the classes into the majority class should be the optimal classifier if the relative weights and reward are extreme. The conclusion will be the opposite if the two cost requirements are exchanged. With the first index pair named as classification class and the second pair index named as reference class according to Figure 1, this cost requirement can be easily specified by the following column-wise operation:

1. decreasing the FPR weight such that the majority class is the classification class.
2. increasing the FPR weights such that the majority class is the reference class.
3. increasing the TPR weights such that the majority class is the classification class.
4. decreasing the TPR weights such that the majority class is the reference class.

To test the change of the ROC plot with respect to the change of the misclassification cost  $Q_j$ , we initialize all the pairwise FPR and pair-wise TPR weights to be one and gradually change the pair-wise TPR/FPR weights associated with the majority class through a multiplication/division on a parameter  $c$ . To illustrate the experiment design, we provide an example of TPR/FPR weight matrix with  $k = 3$  and the first class being the majority class below, that is we choose  $Q_j$  such that

$$W^{pos} \circ Q^{tp} = \begin{matrix} & \begin{matrix} (1, 2) & (1, 3) & (2, 1) & (2, 3) & (3, 1) & (3, 2) \end{matrix} \\ \begin{matrix} c & c & 1/c & 1 & 1/c & 1 \\ c & c & 1/c & 1 & 1/c & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & c & 1/c & 1 & 1/c & 1 \end{matrix} \end{matrix}, W^{neg} \circ Q^{fp} = \begin{matrix} & \begin{matrix} (1, 2) & (1, 3) & (2, 1) & (2, 3) & (3, 1) & (3, 2) \end{matrix} \\ \begin{matrix} 1/c & 1/c & c & 1 & c & 1 \\ 1/c & 1/c & c & 1 & c & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 1/c & 1/c & c & 1 & c & 1 \end{matrix} \end{matrix}$$

When  $c > 1$ , it corresponds to the scenario where the relative TPR weight of the majority class is high and the relative FPR weight of the majority class is low. This implies the majority classifier is optimal since the cost of mis-classification is low while the reward the correct classification is high. The scenario is exactly the opposite when  $c < 1$ . The parameter  $c$  is then changed from  $c = 0.1, c = 0.2, \dots, c = 0.9$ , and  $c = 1/0.1, c = 1/0.2, \dots, c = 1/0.9$ .

To demonstrate that this conclusion is invariant to different number of the classes, we choose  $n = 10,000$  and vary  $k$  from 5, 10, 15. To also validate the impact of class skewness, we simulate the 10,000 classification labels disproportionately according to direction distributed vector  $\vec{w} \sim \text{Dir}(\alpha)$  with  $\alpha = 2, 5, 9$ . The result is summarized below:

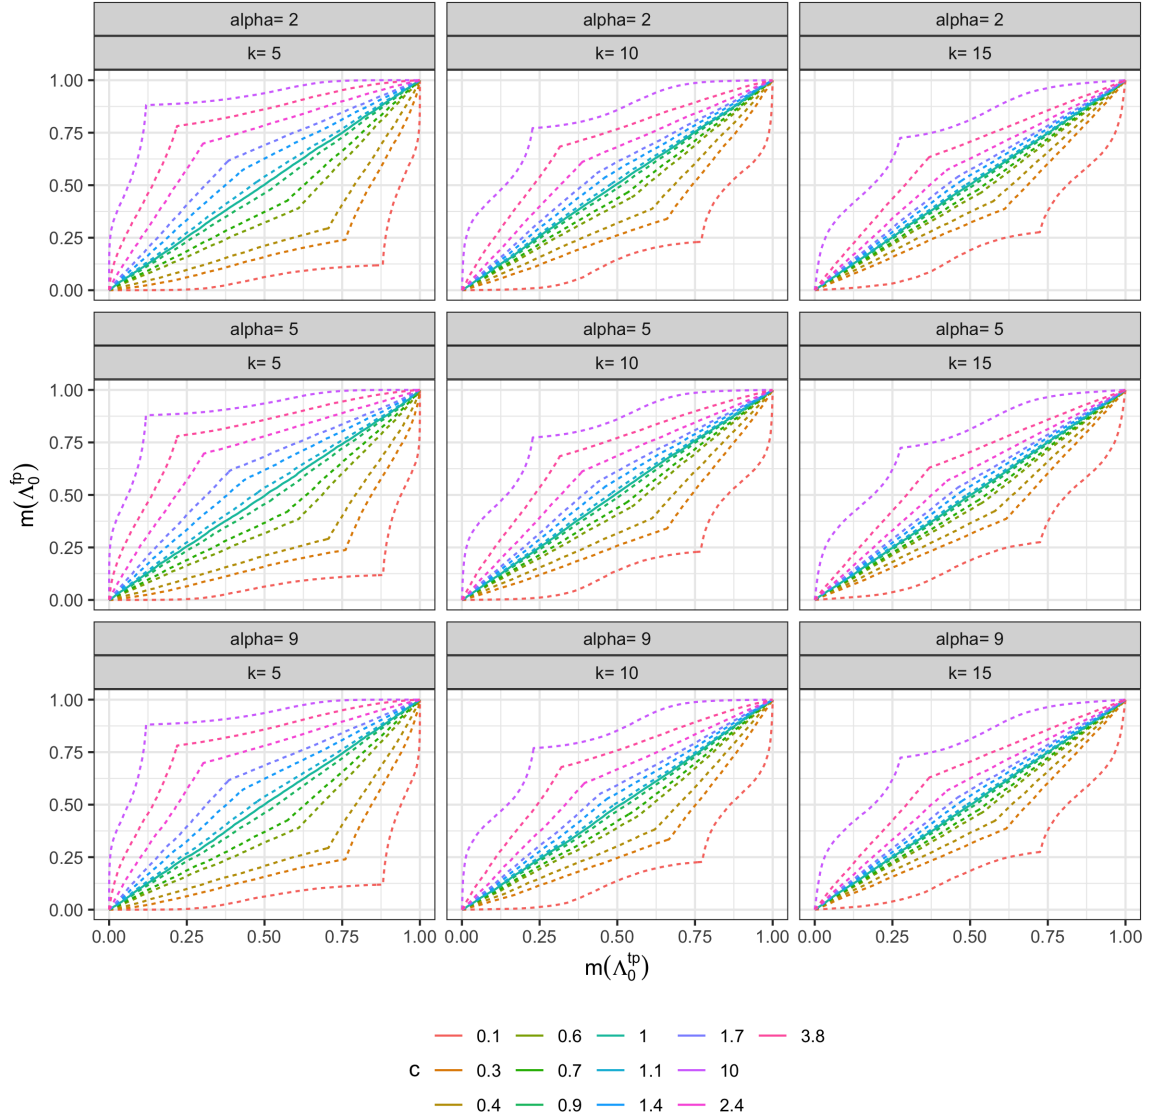


Figure 6: **Weight specification experiment**

As we can see from Figure 6, initially with  $c = 1$ , the majority classifier can be considered as a random classifier by having the ROC curve as a straight line along the diagonal. The weighted ROC curve then correctly reflected the weight specification by moving the ROC curve to the top-left if  $c > 1 (c \rightarrow \infty)$  and by moving the ROC curve to the bottom-right if  $c < 1 (c \rightarrow 0)$ .

## 5.4 Real World Experiment

To comprehensively examine our methods, we experimented with benchmark datasets that have been used by existing literature. To illustrate the invariance to class-skewness property, (Hand and Till, 2001) experimented on eight commonly used multi-class datasets from the UCI Repository of Machine Learning Databases<sup>1</sup>. For a comparison to the pair-wise AUC statistics (called  $\mathcal{M}$  measurement in (Hand and Till, 2001)), we applied our method on the same datasets and adopted the same sets of classification models with similar training procedures. Specifically, for each of the eight datasets, we fit logistic regression, k-nearest neighbor, and decision tree used in Splus(Venables et al., 1999).

We adopted the same training procedure by splitting the dataset into equal sizes of training and testing. It is also important to realize that due to the absence of a random seed number, the ranking of the three model performances can sometimes be different from the Hands' paper due to train/test split. We thus instead focus on the ranking discrepancies among different classification models between our statistics and the  $\mathcal{M}$  statistics conditional on our train/test split. To plot the confidence interval of our models, we bootstrap with binomial weight  $W_j = n_0(j) \times n_1(j)$  100 times the computed AUC and the plotted ROC after obtaining the estimation of unweighted  $\eta^{tp}$  and  $\eta^{fp}$  with  $Q_j \circ W_j = 1$ . The bootstrapped  $\mathbb{M}^{tp}$  and  $\mathbb{M}^{fp}$  can thus provide a confidence interval on the estimated  $\mathcal{D}$  statistics. We summarize the bootstrapped  $\mathcal{D}$  statistics for the six datasets with the following box plot. The pair-wise AUC score is plotted as a blue dot for comparison:

---

1. <https://archive.ics.uci.edu/ml/index.php>

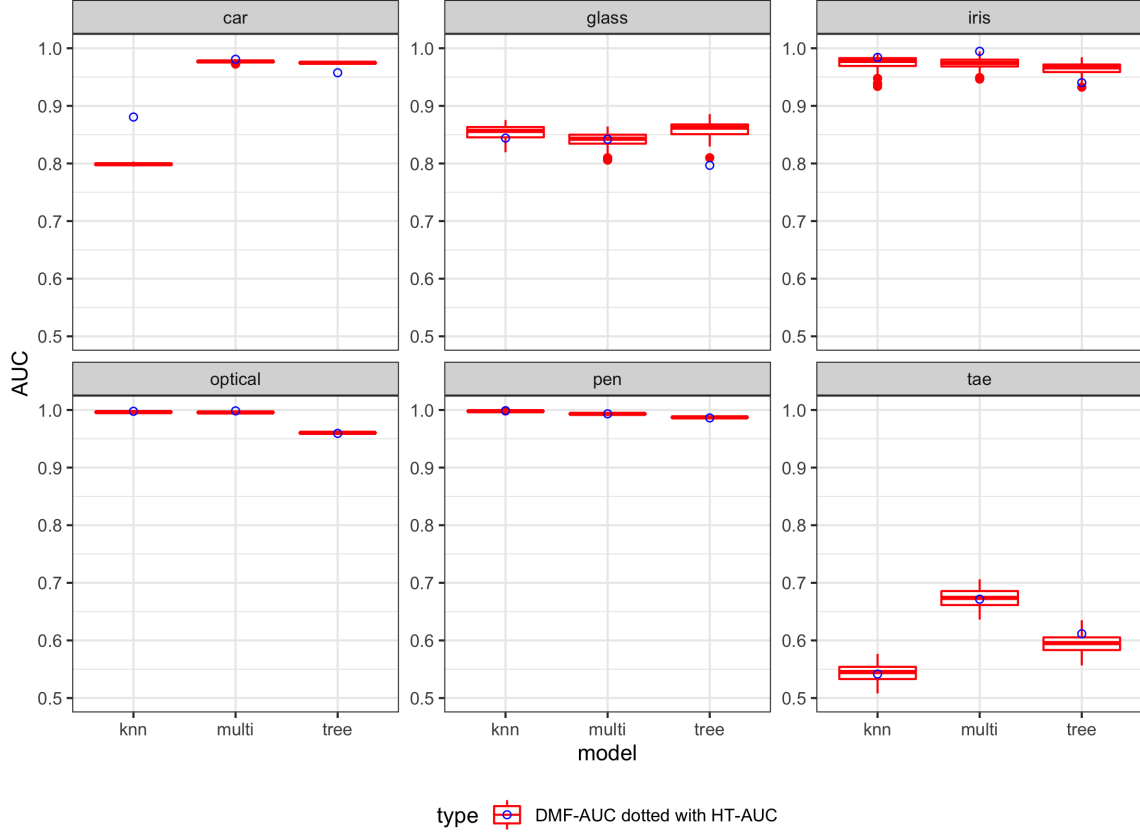


Figure 7: AUC comparison

As it is indicated in Figure 7, even though the exact AUC numbers obtained from  $\mathcal{M}$  and  $\mathcal{D}$  are different. The ranking of the model performance tends to agree. Moreover, for some cases when the  $\mathcal{M}$  statistics give a performance score with a negligible difference (e.g the IRIS dataset and the GLASS dataset), our  $\mathcal{D}$  method quantifies the uncertainty behind the AUC statistics with a confidence interval. This observation additionally substantiates the superiority of our method by having our evaluation metric being able to provide more discriminative ranking when the  $\mathcal{M}$  statistics fails to do so (Ling et al., 2003; Halimu et al., 2019). In fact, due to the bootstrap sampling of the confidence interval, one can readily estimate the probability of ranking statistics such as:  $\mathbb{P}(\text{AUC}(\text{tree}) > \text{AUC}(\text{multinomial}) > \text{AUC}(\text{knn}))$ . There are six possible permutation of the ranking statistics and we provide their corresponding estimate in the following table:

# MULTICLASS ROC

Dataset	knn>tree>multi	knn>multi>tree	multi>tree>knn	multi>knn>tree	tree>knn>multi	tree>multi>knn
car	-	-	0.84	-	-	0.16
optical	-	0.94	-	0.06	-	-
pen	-	1.00	-	-	-	-
tae	-	-	0.99	0.01	-	-
iris	0.16	0.36	0.15	0.22	0.04	0.07
glass	0.25	0.13	-	0.04	0.37	0.21

Table 1: Probability estimation of classification model ranking

As we can see from Table 1, the car, optical, pen and tae dataset tend to have more certainty in the ultimate classification ranking while it seems that more possible ranking orders are available for the Glass and IRIS dataset. One can also check on the shared threshold effect (ROC) by plotting the  $m(\Lambda_0^{tp})$  against  $m(\Lambda_0^{fp})$ :

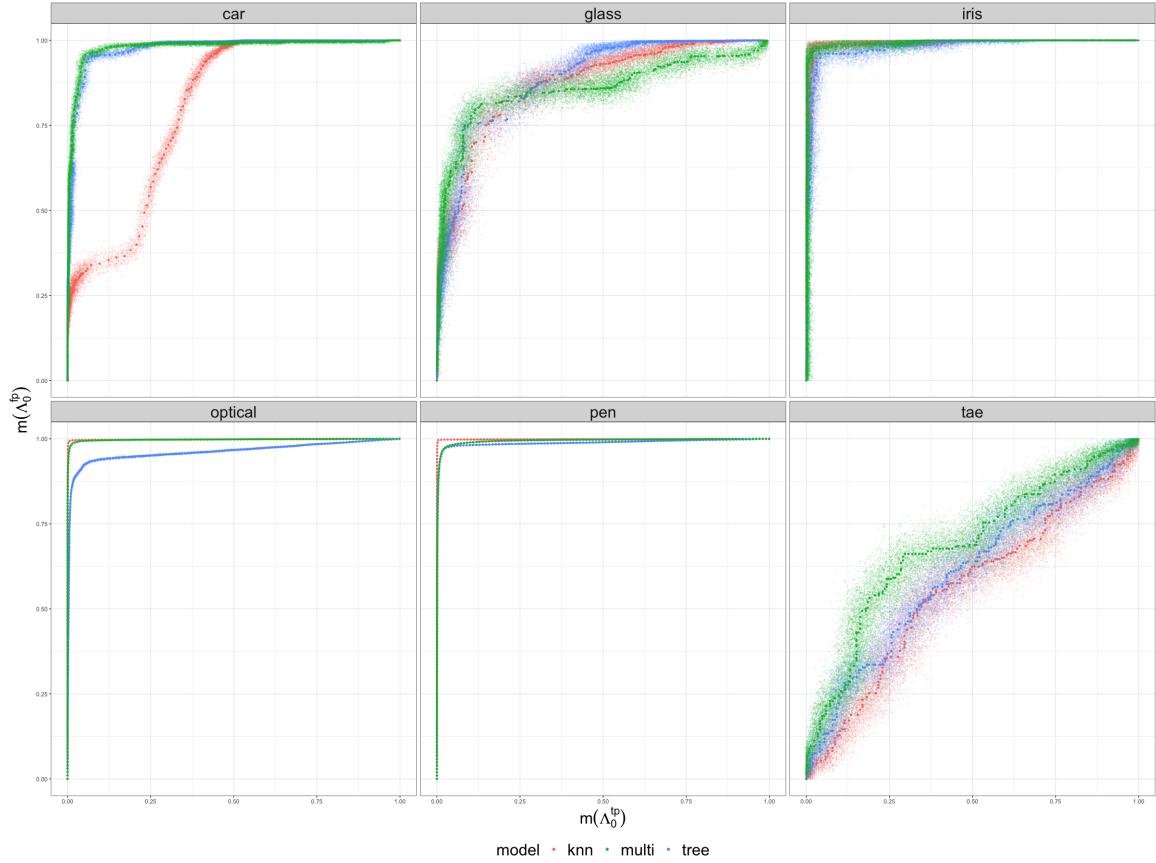


Figure 8: Multiclass ROC with confidence interval

From the AUC comparison of the GLASS dataset, we found that the multinomial regression and the KNN have similar AUC scores (0.8419 and 0.8440). Instead of naively concluding that KNN is preferred, a more careful conclusion can be drawn from the ROC plot in Figure 8. Specifically, we can find that the multinomial regression model tends to

have the highest TPR across all  $k(k-1)$  pairs with a high threshold. This can sometimes be preferred by applications (e.g. Alexa Question & Answering) that prefer a more confident label assignment by thresholding the prediction probability.

## 6. Conclusion

In this paper, we studied an evaluation method for multi-class classification models. The method visually summarize the classifier performance through a rank one factorization. The factorized components coincides with the sROC interpretation by having the column correlation being modeled as random effect. The evaluation method is not only invariant to class skewness but also supports a weight specification for mis-classification loss. A bootstrapped confidence interval is also available to quantify the uncertainty behind the evaluation.

## References

- Niall M Adams and David J Hand. Comparing classifiers when the misallocation costs are uncertain. *Pattern Recognition*, 32(7):1139–1147, 1999.
- LR Arends, TH Hamza, JC Van Houwelingen, MH Heijenbrok-Kal, MGM Hunink, and Theo Stijnen. Bivariate random effects meta-analysis of roc curves. *Medical Decision Making*, 28(5):621–638, 2008.
- S Balaswamy and R Vishnu Vardhan. Confidence interval estimation of an roc curve: an application of generalized half normal and weibull distributions. *Journal of Probability and Statistics*, 2015, 2015.
- Davide Chicco and Giuseppe Jurman. The advantages of the matthews correlation coefficient (mcc) over f1 score and accuracy in binary classification evaluation. *BMC genomics*, 21(1):1–13, 2020.
- Davide Chicco, Matthijs J Warrens, and Giuseppe Jurman. The matthews correlation coefficient (mcc) is more informative than cohen’s kappa and brier score in binary classification assessment. *IEEE*, 2021.
- Jacob Cohen. A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1):37–46, 1960.
- Chris Drummond and Robert C Holte. Cost curves: An improved method for visualizing classifier performance. *Machine learning*, 65(1):95–130, 2006.
- César Ferri, José Hernández-Orallo, and Miguel Angel Salido. Volume under the roc surface for multi-class problems. In *European conference on machine learning*, pages 108–120. Springer, 2003.
- Jan Gorodkin. Comparing two k-category assignments by a k-category correlation coefficient. *Computational biology and chemistry*, 28(5-6):367–374, 2004.

- Chongomweru Halimu, Asem Kasem, and SH Shah Newaz. Empirical comparison of area under roc curve (auc) and mathew correlation coefficient (mcc) for evaluating machine learning algorithms on imbalanced datasets for binary classification. In *Proceedings of the 3rd international conference on machine learning and soft computing*, pages 1–6, 2019.
- David J Hand and Robert J Till. A simple generalisation of the area under the roc curve for multiple class classification problems. *Machine learning*, 45(2):171–186, 2001.
- Charles X Ling, Jin Huang, Harry Zhang, et al. Auc: a statistically consistent and more discriminating measure than accuracy. In *Ijcai*, volume 3, pages 519–524, 2003.
- Brian W Matthews. Comparison of the predicted and observed secondary structure of t4 phage lysozyme. *Biochimica et Biophysica Acta (BBA)-Protein Structure*, 405(2):442–451, 1975.
- Charles E Metz. Basic principles of roc analysis. In *Seminars in nuclear medicine*, volume 8, pages 283–298. Elsevier, 1978.
- Lincoln E Moses, David Shapiro, and Benjamin Littenberg. Combining independent studies of a diagnostic test into a summary roc curve: data-analytic approaches and some additional considerations. *Statistics in medicine*, 12(14):1293–1316, 1993.
- Douglas Mossman. Three-way rocs. *Medical Decision Making*, 19(1):78–89, 1999.
- David MW Powers. Evaluation: from precision, recall and f-measure to roc, informedness, markedness and correlation. *arXiv preprint arXiv:2010.16061*, 2010.
- Foster Provost and Pedro Domingos. Well-trained pets: Improving probability estimation trees. *Raport instytutowy IS-00-04, Stern School of Business, New York University*, 2000.
- Foster Provost and Tom Fawcett. Robust classification systems for imprecise environments. In *AAAI/IAAI*, pages 706–713, 1998.
- Johannes B Reitsma, Afina S Glas, Anne WS Rutjes, Rob JPM Scholten, Patrick M Bossuyt, and Aeilko H Zwinderman. Bivariate analysis of sensitivity and specificity produces informative summary measures in diagnostic reviews. *Journal of clinical epidemiology*, 58(10):982–990, 2005.
- K Ross and D Page. Auc $\mu$ : a performance metric for multiclass machine learning models. In *Proceedings of the International Conference on Machine Learning*, pages 3439–3447, 2019.
- Carolyn M Rutter and Constantine A Gatsonis. A hierarchical regression approach to meta-analysis of diagnostic test accuracy evaluations. *Statistics in medicine*, 20(19):2865–2884, 2001.
- Anna N Angelos Tosteson and Colin B Begg. A general regression methodology for roc curve estimation. *Medical Decision Making*, 8(3):204–215, 1988.
- William N Venables, Brian D Ripley, et al. Modern applied statistics with s-plus, 1999.

Liang Wang and Luis Carvalho. Deviance matrix factorization. *Electronic Journal of Statistics*, 17(2):3762–3810, 2023.

Liang Wang and Luis Carvalho. Computational approaches for exponential-family factor analysis. *arXiv e-prints*, pages arXiv–2403, 2024.

Qiuming Zhu. On the performance of matthews correlation coefficient (mcc) for imbalanced dataset. *Pattern Recognition Letters*, 136:71–80, 2020.