

---

# Binary Classification with Karmic, Threshold-Quasi-Concave Metrics

---

Bowei Yan<sup>1</sup> Oluwasanmi Koyejo<sup>2</sup> Kai Zhong<sup>1</sup> Pradeep Ravikumar<sup>3</sup>

## Abstract

Complex performance measures, beyond the popular measure of accuracy, are increasingly being used in the context of binary classification. These complex performance measures are typically not even decomposable, that is, the loss evaluated on a batch of samples cannot typically be expressed as a sum or average of losses evaluated at individual samples, which in turn requires new theoretical and methodological developments beyond standard treatments of supervised learning. In this paper, we advance this understanding of binary classification for complex performance measures by identifying two key properties: a so-called Karmic property, and a more technical threshold-quasi-concavity property, which we show is milder than existing structural assumptions imposed on performance measures. Under these properties, we show that the Bayes optimal classifier is a threshold function of the conditional probability of positive class. We then leverage this result to come up with a computationally practical plug-in classifier, via a novel threshold estimator, and further, provide a novel statistical analysis of classification error with respect to complex performance measures.

## 1. Introduction

Binary classification, with the goal of predicting a binary response given input features, is perhaps *the* classical problem in machine learning, with wide ranging applications. A key ingredient in binary classification is a *performance measure*, that quantifies how well a given classifier fits the data. While the performance measure of *accuracy* has been the predominant focus of both theory and practice, it has severe limitations in many practical settings, such as imbal-

anced classes, and where different types of errors made by the classifier have different costs (Gu et al., 2009; Wallace et al., 2011). Accordingly, practitioners in applied machine learning settings such as information retrieval and medical diagnosis have developed complex performance metrics that capture important trade-offs between different types of errors; we have collated a few instances in Table 1. A key complication with many complex classification performance metrics, such as the F-measure (Manning et al., 2008) and Harmonic Mean (Kennedy et al., 2009), is that they cannot be decomposed into the sum or average of individual losses on each sample. Even the simple performance measure of precision — the fraction of correct positive predictions, among the set of positive predictions — is not a sum of individual losses on each sample. Thus, the standard theoretical and practical treatments of supervised learning, such as standard empirical risk minimization that minimizes the empirical expectation of a loss evaluated on a single random example, are not applicable.

This practical reality has motivated research into effective and efficient algorithms tailored to complex non-decomposable performance measures. One class of approaches extend standard empirical risk minimization to this non-decomposable setting, which often relies on strong assumptions on either the form of the classifiers, such as requiring linear classifiers (Narasimhan et al., 2015a), or restricted to specific performance measures such as F-measure (Parambath et al., 2015). An alternative approach is the plug-in estimator, where we first derive the form of the Bayes optimal classifier, estimate the statistical quantities associated with the Bayes optimal classifier, and finally “plug-in” the sample estimates of the population quantities to then obtain the overall estimate of the Bayes optimal classifier. In particular, for many complex performance metrics, the Bayes optimal classifier is simply a thresholding of the conditional probability of the positive class (Koyejo et al., 2014; Narasimhan et al., 2014), so that the plug-in estimator requires (a) an estimate of the conditional probability, and (b) the associated threshold. Plug-in methods have been of particular interest in non-parametric functional estimation as they typically require weaker assumptions on the function class and are often easy to implement.

In this paper, we seek to advance our understanding and practice of binary classification under complex non-

---

<sup>1</sup>University of Texas at Austin, Austin, Texas, USA <sup>2</sup>University of Illinois at Urbana-Champaign, Champaign, Illinois, USA <sup>3</sup>Carnegie Mellon University, Pittsburgh, Pennsylvania, USA. Correspondence to: Bowei Yan <bowei@utexas.edu>.

decomposable performance measures. We show that for a very broad class of performance measures, encompassing a large set of performance measures used in practice, the Bayes optimal classifier is simply a thresholding of the conditional probability of the response. Towards this general result, we identify two key properties that a performance measure could satisfy. The first is what we call a ‘‘Karmic’’ property that loosely has the performance measure be more sensitive to an increase in true positives and true negatives, and a decrease in false positives and false negatives. The second is a more technical property we call threshold-quasi-concavity, which in turn ensures the performance measure is well-behaved around an optimal threshold. As we show these properties are satisfied by performance metrics used in practice, and in particular, these conditions are milder than existing results that restrict either the structural form of the performance measures, or impose strong shape constraints such as particular monotonicities.

Our general result has two main consequences, which we investigate further: one algorithmic, and the other for the analysis of classification error for general performance measures. As the algorithmic consequence, we leverage the derived form of the Bayes optimal classifier, and some additional general assumptions on the performance measures, to provide a tractable algorithm to estimate the threshold, which coupled with an estimator of the conditional probability, provides a tractable ‘‘plug-in estimator’’ of the Bayes optimal classifier. Towards the statistical analysis consequence, we provide an analysis of the excess classification error, but with respect to general non-decomposable performance measures, of the general class of plugin-estimators for our class of Bayes optimal classifiers. Our analysis of classification error rates for such plug-in classifiers depend on three natural quantities: the rate of convergence for the conditional probability estimate, the rate of convergence for the threshold estimate, and a measurement of noise in the data. For the last part, we extend margin or low-noise assumptions for binary classification with the accuracy performance measure to complex performance measures. *Low noise* assumptions, proposed by Mammen et al. (1999) in the context of the accuracy performance measure, bounds the noise level in the neighborhood of the Bayes optimal threshold i.e.  $\frac{1}{2}$  for standard classification. Under such a low-noise assumption, Audibert et al. (2007) derive fast convergence rates for plug-in classification rules based on the smoothness of the conditional probability function. Similar margin assumptions have also been introduced for density level set estimation by Polonik (1995). We provide a natural extension of such a low-noise assumption, under which we provide explicit rates of convergence of classification error with respect to complex performance measures. We provide corollaries for both parametric and non-parametric instances of our general class of plugin-classifiers.

The rest of the paper is organized as below. In Section 2 we introduce the problem and relevant notations. The characterization and properties of Bayes optimal classifier are derived in Section 3. We discuss the algorithm for estimating the plug-in estimator in Section 4, and present the statistical convergence guarantee in Section 5. Applications of the derived rate for two special cases, Gaussian generative model and  $\beta$ -Hölder class conditional probability are presented in Section 6 where explicit convergence rates are provided. We conclude the paper in Section 7. Detailed proofs are deferred to the supplementary materials.

## 2. Problem Setup and Preliminaries

Binary classification entails predicting a binary label  $Y \in \{\pm 1\}$  associated with a feature vector  $X \in \mathcal{X} \subset \mathbb{R}^d$ . Such a function mapping  $f : \mathcal{X} \mapsto \{\pm 1\}$  from the feature space  $\mathcal{X}$  to the labels  $\{\pm 1\}$  is called a binary classifier. Let  $\Theta = \{f : \mathcal{X} \rightarrow \{\pm 1\}\}$  denote a set of binary classifiers. We assume  $(X, Y)$  has distribution  $\mathbb{P} \in \mathcal{P}$ , and let  $\eta(x) := \mathbb{P}(Y = 1 | X = x)$  denote the conditional probability of the label  $Y$  given feature vector  $x$ .

A key quantity is the confusion matrix, that consists of four population quantities: true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN).

**Definition 2.1** (Confusion Matrix). *For any classifier  $f : \mathbb{R}^d \mapsto \{\pm 1\}$ , its confusion matrix is defined as  $\mathbf{C}(f, \mathbb{P}) := [TP(f, \mathbb{P}), FP(f, \mathbb{P}), FN(f, \mathbb{P}), TN(f, \mathbb{P})] \in [0, 1]^4$ , where:*

$$\begin{aligned} TP(f, \mathbb{P}) &= \mathbb{P}(Y = +1, f(X) = +1), \\ FP(f, \mathbb{P}) &= \mathbb{P}(Y = -1, f(X) = +1), \\ FN(f, \mathbb{P}) &= \mathbb{P}(Y = +1, f(X) = -1), \\ TN(f, \mathbb{P}) &= \mathbb{P}(Y = -1, f(X) = -1). \end{aligned} \tag{1}$$

Another key ingredient is the utility or performance measure  $\mathcal{U} : \Theta \times \mathcal{P} \rightarrow \mathbb{R}$ , that measures the performance of a classifier. In this paper, we focus on complex binary classification performance measures that can be expressed as a function of the confusion matrix. Formally,  $\mathcal{U}(f, \mathbb{P}) = \mathcal{G}(\mathbf{C}(f, \mathbb{P}))$ . When it is clear from context, we will drop the dependency of the distribution  $\mathbb{P}$  in  $\mathbf{C}$  and  $\mathcal{U}$ . The confusion-matrix functions  $\mathcal{G}$  corresponding to popular performance measures are listed in Table 1. Given the performance measure  $\mathcal{U}$ , we are interested in the corresponding Bayes optimal classifier:

$$f^* = \arg \max_{f \in \Theta} \mathcal{U}(f, \mathbb{P}). \tag{2}$$

Given any candidate classifier  $f$ , we are then interested in the excess risk  $\mathcal{U}(f^*, \mathbb{P}) - \mathcal{U}(f, \mathbb{P})$ , which is the utility gap between a given classifier  $f$  and the corresponding Bayes optimal.

**Assumption 1** (Karmic Performance Measure). *The confusion-matrix function  $\mathcal{G}$  corresponding to the performance measure  $\mathcal{U}$  is Lipschitz continuous, and satisfies the*

Table 1. Examples of evaluation metrics. Notation:  $\text{TPR} = \frac{\text{TP}}{\text{TP}+\text{FN}}$ ;  $\text{TNR} = \frac{\text{TN}}{\text{TN}+\text{FP}}$ .

METRIC	DEFINITION	REFERENCE	$\mathcal{G}(\mathbf{C})$
ACCURACY	$\text{TP} + \text{TN}$		$(1, 0, 0, 1)\mathbf{C}$
ARITHMETIC MEAN (AM)	$(\text{TPR} + \text{TNR})/2$	MENON ET AL. (2013)	$(\frac{C_1}{C_1+C_3} + \frac{C_4}{C_2+C_4})/2$
YOUDEN'S INDEX	$\text{TPR} + \text{TNR} - 1$	YOUDEN (1950)	$\frac{C_1}{C_1+C_3} + \frac{C_4}{C_2+C_4} - 1$
$F_\beta$	$\frac{(1+\beta^2)\text{TP}}{(1+\beta^2)\text{TP}+\beta^2\text{FN}+\text{FP}}$	VAN RIJSBERGEN (1974)	$\frac{(1+\beta^2, 0, 0, 0)\mathbf{C}}{(1+\beta^2, 1, \beta^2)\mathbf{C}}$
LINEAR-FRACTIONAL	$\frac{a_1\text{TP}+a_2\text{FP}+a_3\text{FN}+a_4\text{TN}}{b_1\text{TP}+b_2\text{FP}+b_3\text{FN}+b_4\text{TN}}$	KOYEJO ET AL. (2014)	$\frac{\mathbf{a}^T\mathbf{C}}{\mathbf{b}^T\mathbf{C}}$
G-MEAN	$\sqrt{\text{TPR} \cdot \text{TNR}}$	DASKALAKI ET AL. (2006)	$\sqrt{\frac{C_1 C_4}{(C_1+C_3)(C_2+C_4)}}$
Q-MEAN	$1 - \sqrt{\frac{(1-\text{TPR})^2+(1-\text{TNR})^2}{2}}$	KUBAT ET AL. (1997)	$1 - \sqrt{\frac{(\frac{C_3}{C_1+C_3})^2+(\frac{C_2}{C_2+C_4})^2}{2}}$
H-MEAN	$\frac{2}{1/\text{TPR}+1/\text{TNR}}$	KENNEDY ET AL. (2009)	$2 / \left( \frac{C_1+C_3}{C_1} + \frac{C_2+C_4}{C_4} \right)$

condition that  $\nabla\mathcal{G}(\mathbf{C})^T(1, -1, -1, 1)^T \geq C_B$ , for some constant  $C_B > 0$ .

We term performance measures that satisfy the condition  $\nabla\mathcal{G}(\mathbf{C})^T(1, -1, -1, 1)^T \geq C_B$  as ‘‘Karmic measures’’, since it guarantees a lower bound on the sensitivity of the performance measure in the direction of increasing true positives and true negatives, and decreasing false positives and false negatives. While our Karmic assumption slightly weakens the existing monotonicity assumption used in literature, it is worth pointing out that the analysis in (Narasimhan et al., 2014) requires not only monotonicity but also additional assumptions (Assumption B in (Narasimhan et al., 2014)). Assumption B assumes the existence and uniqueness of an optimal threshold, which turns out to be non-trivial to check. Our analysis on the threshold-quasi-concavity closes this gap.

Assumption 1 is satisfied if  $\mathcal{G}$  is strictly monotonically increasing with respect to TP, TN and decreasing with respect to FP, FN. Such an assumption is natural in that one would typically prefer a classifier with higher TP for fixed TN and vice versa (Narasimhan et al., 2015b). This condition is satisfied by most metrics in common use e.g. for the  $F_1$  measure,  $\nabla\mathcal{G}(\mathbf{C})^T(1, -1, -1, 1)^T = \frac{4(\text{FP}+\text{FN})}{(2\text{TP}+\text{FP}+\text{FN})^2}$  is strictly positive as long as the data is not fully separable.

## 2.1. Related Work

**Representation of the Bayes Optimal Classifier.** The Bayes optimal classifier under the accuracy metric is classically known to be a thresholding of the conditional probability of the response, with the threshold of  $1/2$  (see e.g. Devroye et al. (2013)). This property of Bayes optimal classifier having the thresholded form is called the probability thresholding principle for binary classification by Lewis (1995). Prior work has also shown that the thresholding principle, with a metric dependent threshold, for more complex specific measures such as F-measure (Jansche, 2007;

Zhao et al., 2013), Arithmetic Mean (AM) (Menon et al., 2013), linear-fractional performance metrics (Koyejo et al., 2014), and monotonic concave metrics (Narasimhan et al., 2015a).

**Plug-in Classifiers for Complex Metrics.** For Bayes optimal classifiers that have thresholded form, a line of work has devised plug-in classifiers that then estimate the threshold, and the conditional probability of response. For the AM metric, Menon et al. (2013) show that the threshold is simply the proportion of the positive class. For linear fractional functions, Koyejo et al. (2014) provide an implicit characterization of the optimal threshold, but the solution of which in turn requires the knowledge of the optimal classifier, which is unknown in practice. As a practical estimator, Koyejo et al. (2014) propose an exhaustive search for the threshold over all data samples, and show that the resulting algorithm is consistent, but for which non-asymptotic convergence rates are not known. Narasimhan et al. (2014) also note the importance of estimating the optimal threshold, but do not provide practical algorithms. As we show in Section 3, the empirical risk as a function of the threshold is in general neither convex nor concave. Hence, care must be taken to construct an optimization algorithm that guarantees convergence to the true threshold.

**Estimators designed for specific Utility Functions.** Perhaps the most studied non-decomposable performance metric is the F-measure (Nan et al., 2012; Joachims, 2005; Zhao et al., 2013), with wide use in information retrieval and related areas, and for which researchers have developed tailored estimators (Nan et al., 2012; Joachims, 2005) as well as risk bounds for these estimators (Zhao et al., 2013). For instance, Busa-Fekete et al. (2015) propose a scalable online F-measure estimator for large-scale datasets, with a root finding algorithm for the threshold update which exploits special properties of the F-measure. Similarly, for the

Arithmetic Mean (AM) measure, Menon et al. (2013) design a consistent optimization scheme, based on a balanced classification-calibrated surrogate to AM. Unfortunately, these techniques are not easily extended to general complex performance metrics.

**Algorithms for General Classification Measures.** Joachims (2005) poses the classification problem as a structured prediction problem, and for linear classifiers, propose a structural SVM solver, but for which neither consistency nor explicit convergence rates are known. Kar et al. (2014) proposes an online gradient descent algorithm which requires function classes that satisfy a uniform convergence property, which is difficult to verify apriori. Along similar lines, Narasimhan et al. (2015a) propose a stochastic gradient method, that involves a linearization of classification metric. Their proposed approach depends strongly on the assumption of a linear (or kernelized) classifier, and it is not obvious that the procedure can be extended to more complex non-linear function classes.

### 3. The Bayes Optimal Classifier Revisited

In this section, we characterize the Bayes-optimal classifier for the broad class of Karmic performance measures, that satisfy Assumption 1. We then show that with one additional assumption, we call threshold-quasi-concavity, the optimal threshold can be guaranteed to be unique. This result will be crucial for the design and analysis of our computationally efficient threshold finding procedure in Section 4.

Denote  $\mu$  as the measure corresponding to the marginal distribution of  $X$ . The utility is Frechét differentiable, whose Frechét derivative of  $\mathcal{U}$  may be computed as:

$$\begin{aligned} [\nabla \mathcal{U}(f)]_x &= \nabla \mathcal{G}(\mathcal{C}(f))^T \cdot [\nabla \mathcal{C}(f)]_x \\ &= \frac{1}{2} (\nabla \mathcal{G}(\mathcal{C})^T(1, -1, -1, 1)^T \eta(x) \\ &\quad - \nabla \mathcal{G}(\mathcal{C})^T(0, -1, 0, 1)^T) d\mu(x) \end{aligned}$$

From the Karmic measure Assumption 1, we know that  $\nabla \mathcal{G}(\mathcal{C})^T(1, -1, -1, 1)^T > 0$ . We define the ‘‘Bayes critical set’’ of  $\mathcal{G}(f, \mathbb{P})$  for any  $f \in \mathcal{F}$  as the set of instances where the utility has zero derivative:

$$A_3(f) = \left\{ x : \eta(x) = \frac{\nabla \mathcal{G}(\mathcal{C})^T(0, -1, 0, 1)^T}{\nabla \mathcal{G}(\mathcal{C})^T(1, -1, -1, 1)^T} \right\}.$$

For notational simplicity, we have omitted the dependency of  $g_i$  on  $\mathcal{C}(f)$ . Similarly, we will use  $A_3^* := A_3(f^*)$  to denote the Bayes critical set.

In this paper we focus on distributions where the critical set of  $\mathcal{G}(f, P)$  satisfies  $\mathbb{P}(A_3(f)) = 0$ . For instance, this is true for any distribution that satisfies the following assumption.

**Assumption 2** ( $\eta$ -continuity). *Let  $\nu$  denote the probability measure that is associated with random variable  $Z = \eta(X) = P(Y = 1|X)$ , then  $\nu$  is absolutely continuous with respect to  $\mu$ . Furthermore, the density of  $\eta(X)$ , denoted by  $p_\eta(\cdot)$ , has full support on  $[0, 1]$ , and is bounded everywhere.*

Absolute-continuity guarantees the existence of the density of  $Z$ . Armed with the above assumption on the conditional probability of the response, we can then characterize the Bayes optimal classifier as follows.

**Theorem 3.1** (Bayes Optimal Classifier as a Thresholding Function). *Suppose that  $\mathcal{U}$  is a performance measure that satisfies Assumption 1, and that  $\eta(X)$  satisfies Assumption 2. Let  $f^*$  be the Bayes classifier with respect to  $\mathcal{U}$  and  $\mathcal{C}^*$  be its confusion matrix. Then, for all  $x \in (A_3^*)^c$ ,*

$$f^*(x) = \text{sign} \left( \eta(x) - \frac{\nabla \mathcal{G}(\mathcal{C}^*)^T(0, -1, 0, 1)^T}{\nabla \mathcal{G}(\mathcal{C}^*)^T(1, -1, -1, 1)^T} \right). \quad (3)$$

**Threshold of Bayes optimal classifier** For some performance measures, the optimal threshold reduces to an absolute constant; for instance it has the value of  $1/2$  for the accuracy measure  $\mathcal{U}(f, \mathbb{P}) = \text{TP} + \text{TN}$  (see e.g. (Devroye et al., 2013)). In the general case however, the optimal threshold  $\delta^*$  is a solution of the fixed point equation:

$$(\nabla \mathcal{G}(\mathcal{C}^*)^T(0, -1, 0, 1)^T) / (\nabla \mathcal{G}(\mathcal{C}^*)^T(1, -1, -1, 1)^T) = \delta^*,$$

which is fixed point equation due to the dependency of  $\mathcal{C}^*$  on the threshold  $\delta^*$ . Theorem 3.1 guarantees the existence of a solution to the above fixed point equation, but not its uniqueness. As we will show in Section 5, uniqueness can be achieved with some additional regularity assumptions.

We note that Theorem 3.1 only imposes a weak Karmic assumption on the performance measure, which as stated in Section 2, is more general than even a simple strictly monotonicity assumption. In particular, it generalizes prior work such as (Koyejo et al., 2014; Menon et al., 2013), that impose more stringent assumptions (linear or linear fractional form of the measures, or strong monotonicity conditions).

We next briefly discuss why the critical set is crucial. Consider for instance the example studied in Narasimhan et al. (2014): with domain  $\mathcal{X} = \{x_1, x_2, x_3\}$ , a corresponding probability mass function  $(0.25, 0.5, 0.25)$ , and the conditional probability  $\eta = (0.49, 0.5, 0.51)$ . Narasimhan et al. (2014) show that for this setting, and for the case of the H-mean measure, there exist at least two deterministic Bayes optima:  $(-1, 1, -1)$  and  $(1, -1, 1)$ , which can be seen to not have a thresholded form i.e. it cannot be expressed as a (signed) thresholding of the conditional probability. Our analysis reveals why this is the case.

From the threshold expression in (3) from Theorem 3.1, the optimal threshold can be computed explicitly as  $\frac{\nabla \mathcal{G}(\mathbf{C}^*)^T(0, -1, 0, 1)^T}{\nabla \mathcal{G}(\mathbf{C}^*)^T(1, -1, -1, 1)^T} = \frac{1}{2}$ . Thus, the Bayes critical set  $A_3^* = \{x : \eta(x) = \frac{1}{2}\} = \{x_2\}$  has measure  $P(X \in A_3) = P(X = x_2) = \frac{1}{2} > 0$ . It is clear that the Bayes optimal classifier may not take a thresholded form on the Bayes critical set.

### 3.1. Uniqueness of the Bayes Optimal Threshold.

We are interested in characterizing mild conditions on the performance measure under which the fixed point equation characterizing the Bayes optimal threshold has a unique solution, under which case  $P(A_3^*) = 0$  (guaranteed by the  $\eta$ -continuity Assumption 2).

The performance measure restricted to classifiers that are threshold functions of the conditional probability, can be rewritten as a function of the conditional probability  $\eta$  and the threshold  $\delta$ .

**Definition 3.1.** We define  $\mathcal{V}_\eta(\delta, \mathbb{P}) := \mathcal{U}(f_{\eta, \delta}, \mathbb{P})$  as the performance measure of any threshold classifier  $f_{\eta, \delta}(x) = \text{sign}(\eta(x) - \delta)$ . Its arguments are the threshold  $\delta$ , and distribution  $\mathbb{P}$ , while the subscript  $\eta$  notes its dependence on the conditional probability  $\eta$ .

We next introduce the definition of quasi-concavity, and the assumption of  $\mathcal{V}$  being strictly quasi-concave.

**Definition 3.2.** A function  $f : \mathcal{X} \rightarrow \mathbb{R}$  is said to be quasi-concave if  $\forall x, y \in \mathcal{X}$ , such that  $f(x) \leq f(y)$ , it follows that  $\langle \nabla f(x), y - x \rangle \geq 0$ . We further say that  $f$  is strictly quasi-concave if it is quasi-concave and its gradient only vanishes at the global optimum, i.e.,  $f(y) < \max_{x \in \mathcal{X}} f(x) \Rightarrow \|\nabla f(y)\| > 0$ .

Quasi-concave functions have super level sets are convex sets, and moreover by definition are unimodal i.e. have a unique maximal point.

**Assumption 3.** (Threshold-Quasi-Concavity) The threshold-classifier performance measure  $\mathcal{V}_\eta(\delta, \mathbb{P})$  is strictly quasi-concave for  $\delta \in [0, 1]$ .

Assumption 3 seems abstract, but it entails that the performance measure is well-behaved as a function of the threshold. Moreover, it can be easily shown to hold for performance measures in practical use. We provide a proposition that shows that the assumption is satisfied for two important classes of performance measures: linear-fractional functions and concave functions.

**Proposition 3.1.** If Assumptions 1, 2 hold, and either: (a)  $\mathcal{G}$  is twice continuously differentiable and concave, or (b)  $\mathcal{G}$  is a linear fractional function  $\mathcal{G}(\mathbf{C}) = \frac{\mathbf{a}^T \mathbf{C}}{\mathbf{b}^T \mathbf{C}}$  with  $|\mathbf{b}^T \mathbf{C}| > 0$ . Then  $\mathcal{V}_\eta(\delta, \mathbb{P})$  is strictly quasi-concave.

**Theorem 3.2.** Under Assumption 3, the fixed-point equa-

tion:

$$\delta = \frac{\nabla \mathcal{G}(\mathbf{C}(f_\delta), \mathbb{P})^T(0, -1, 0, 1)^T}{\nabla \mathcal{G}(\mathbf{C}(f_\delta), \mathbb{P})^T(1, -1, -1, 1)^T}, \quad (4)$$

where  $f_\delta(x) = \text{sign}(\eta(x) - \delta)$ , has a unique fixed point  $\delta^* \in (0, 1)$ . Hence the threshold in Theorem 3.1 is uniquely defined.

Theorems 3.1 and 3.2 have two key consequences: first, we can use the representation to design plugin-estimators of the Bayes optimal classifier; second, it facilitates the statistical analysis for rates of convergence. We will discuss each of these two consequences in the following sections.

## 4. Algorithmic Consequence: Estimation of the Threshold

Theorem 3.1 shows that for Karmic performance measures, the Bayes optimal classifiers has the thresholded form as in Eq. (3), and moreover under the threshold-quasi-concavity Assumption 3, this threshold is unique. An immediate algorithmic consequence of this is to focus on plug-in classifiers that separately estimate the conditional probability, and the threshold. We present this plugin-classifier template in Algorithm 1. The template needs: (a) an estimator for conditional probability density  $\eta(x)$ , and (b) an estimator for the threshold. For the convenience of analysis, we divide the set of samples into two independent subsets: the conditional probability estimator is estimated using one subset, and the threshold is estimated using the other. In the coming subsec-

---

### Algorithm 1 Two-step Plug-in Classifier for General Metrics

---

- 1: **Input:** Training sample  $\{X_i, Y_i\}_{i=1}^n$ , utility measure  $\mathcal{U}$ , conditional probability estimator  $\hat{\eta}$ , stepsize  $\alpha$ .
  - 2: Randomly split the training sample into two subsets  $\{X_i^{(1)}, Y_i^{(1)}\}_{i=1}^{n_1}$  and  $\{X_i^{(2)}, Y_i^{(2)}\}_{i=1}^{n_2}$ ;
  - 3: Estimate  $\hat{\eta}$  on  $\{X_i^{(1)}, Y_i^{(1)}\}_{i=1}^{n_1}$ ;
  - 4: Estimate  $\hat{\delta}$  with  $\{X_i^{(2)}, Y_i^{(2)}\}$ ;
  - 5: **Output:**  $\hat{f}(x) = \text{sign}(\hat{\eta} - \hat{\delta})$ .
- 

tions we discuss how to estimate the conditional probability and the threshold respectively.

### 4.1. Estimation of Conditional Probability Function

The estimation of the conditional probability of the response plays a crucial role in the success of Algorithm 1, but we emphasize that it is not the focus of our paper. In particular, this is a well studied problem, and numerous methods have been proposed for both parametric and non-parametric model assumptions on the conditional probability function.

In this section we briefly discuss some common estimators, and defer additional details to Section 6.

**Parametric methods.** In a classical paper, Ng and Jordan (2002) compares two models of classification: one can either estimate  $P(Y)$  and  $P(X|Y)$  first, then get the conditional probability by Bayes rule (generative model approach); or directly estimate  $P(Y|X)$  (discriminative model approach). The two approaches can also be related. In particular, if  $P_{\theta_Y}(X|Y)$  belongs to exponential family, we have

$$P_{\theta_Y}(X|Y) = h(x) \exp(\langle \theta_Y, \phi(X) \rangle - A(\theta_Y)),$$

where  $\phi(X)$  is the set of sufficient statistics,  $\theta_Y$  is the vector of the true canonical parameters, and  $A(\theta)$  is the log-partition function. Using Bayes rule, we then have:

$$P(Y = 1|X) = \frac{1}{1 + \exp(-\langle \theta_1 - \theta_0, \phi(X) \rangle + c^*)}$$

where  $c^* = A(\theta_0) - A(\theta_1)$ . The conditional distribution can be seen to follow a logistic regression model, with the generative model sufficient statistics as the features, and the difference of the generative model parameters serving as the parameters of the discriminative model. A natural class of estimators for either the generative or discriminative models is based on Maximum likelihood Estimation (MLE). In Section 6, we derive the rate of convergence for the special case where the generative distributions are Gaussians with same covariances for both classes.

**Non-parametric methods.** One can also estimate  $\eta(x) = P(Y = 1|X)$  non-parametrically, where a common model assumption is some form of smoothness on  $\eta(x)$ . One popular class of smooth functions is the following.

**Definition 4.1** ( $\beta$ -Hölder class). *Let  $\beta > 0$ , denote  $\lfloor \beta \rfloor$  the maximal integer that is strictly less than  $\beta$ . For  $x \in \mathcal{X}$  and any  $\lfloor \beta \rfloor$ -times continuously differentiable real-valued function  $\eta$  on  $\mathcal{X}$ , we denote by  $\eta_x$  its Taylor polynomial of degree  $\lfloor \beta \rfloor$  at point  $x$ ,*

$$\eta_x(x') = \sum_{s \leq \lfloor \beta \rfloor} \frac{(x' - x)^s}{s!} D^s \eta(x).$$

$\beta$ -Hölder class is defined as the functions that satisfy, for  $\forall x, x' \in \mathcal{X}$ ,

$$|\eta_x(x) - \eta_x(x')| \leq C_\beta \|x - x'\|^\beta.$$

In particular, when  $0 \leq \beta < 1$ , we have  $|\eta(x) - \eta(x')| \leq C_\beta \|x - x'\|^\beta$  where  $\beta > 0$ .

We can then estimate  $\eta(x)$  from this family of smooth functions via locally polynomial estimators (Audibert et al., 2007), or kernel (conditional) density estimators (Jiang, 2017) with a properly chosen bandwidth.

## 4.2. Estimation of the Threshold

When  $\mathcal{V}_\eta$  is quasi-concave, a key consequence is that its gradient with respect to the threshold suffices to provide ascent direction information. We leverage this consequence, and summarize a simple binary search algorithm based on the sign of  $\mathcal{V}'_\eta(\delta, \mathbb{P})$  in Algorithm 2.

---

### Algorithm 2 Binary search for the optimal threshold

---

```

1: Input: Training sample  $\{X_i, Y_i\}_{i=1}^n$ , utility measure  $\mathcal{U}$ , conditional probability estimator  $\hat{\eta}$ , tolerance  $\epsilon_0$ .
2:  $\delta_\ell = 0; \delta_r = 1;$ 
3: while  $|\delta_\ell - \delta_r| \geq \epsilon_0$  do
4:   Evaluate  $s = \text{sign}(\mathcal{V}'_{\hat{\eta}}(\delta, \mathbb{P}_n))$ ;
5:   if  $s \geq 0$  then
6:      $\delta_\ell = \frac{\delta_\ell + \delta_r}{2};$ 
7:   else
8:      $\delta_r = \frac{\delta_\ell + \delta_r}{2};$ 
9:   end if
10: end while
11: Output:  $\frac{\delta_\ell + \delta_r}{2}$ .
    
```

---

In the next section, we then analyze the rates of convergence for the excess generalization error of the plug-in classifier learned from Algorithm 1, and with threshold estimated via Algorithm 2.

## 5. Statistical Analysis Consequence: Rates of Convergence

We next analyze the convergence rate of the excess utility. As we will show, the rates of convergence depend on three quantities: the noise level of the data distribution, the convergence rate of the conditional probability function, and the convergence rate of the threshold. We start by introducing some assumptions.

We assume that the estimator of the conditional probability of response satisfies the following condition.

**Assumption 4.** *Let  $\mathcal{S}_n$  denote a sample set of size  $n$ , and  $\eta_{\mathcal{S}_n}$  denote the conditional probability estimator learnt from  $\mathcal{S}_n$ . Then, for some absolute constants  $c_1, c_2 > 0$ , the conditional probability estimator satisfies the following condition:*

$$\sup_{\mathcal{S}_n} P(|\eta_{\mathcal{S}_n}(x) - \eta(x)| \geq \epsilon) \leq c_1 \exp(-c_2 a_n \epsilon^2) \quad \text{a.e.}$$

The convergence rate also depends on the noise in the training labels, which is typically captured via the probability mass near the Bayes optimal threshold. Here we generalize the classical margin assumption (sometimes also called low noise assumption) of Audibert et al. (2007), developed for the accuracy metric, to the case where the optimal threshold is not a fixed constant  $\frac{1}{2}$ :

**Assumption 5.** For some function  $C_0(\delta^*) > 0$  that depends on the threshold  $\delta^*$ , there exists an  $\alpha \geq 0$  such that

$$P_X(0 < |\eta(X) - \delta^*| \leq t) \leq C_0(\delta^*) t^\alpha.$$

The assumption characterizes the behavior of the regression function in the vicinity of the optimal threshold  $\delta^*$ . The case  $\alpha = 0$  bounds the probability by a constant potentially larger than one, and is trivially satisfied. The other extreme case  $\alpha = \infty$  is most advantageous for classification, since in this case the regression function  $\eta$  is bounded away from the threshold.

In cases where the threshold is not an absolute constant (such as  $1/2$ ), it has to be estimated from data. We make the following assumption on its convergence rate.

**Assumption 6.** Given a conditional probability estimate  $\hat{\eta}$  learned from an independent data source, the estimator  $\hat{\delta}_n$  of the threshold, from a sample set of size  $n$ , satisfies the following condition, for some absolute constant  $c_3 > 0$ :

$$P\left(\left|\mathcal{U}(\text{sign}(\hat{\eta} - \delta^*)) - \mathcal{U}(\text{sign}(\hat{\eta} - \hat{\delta}))\right| \geq b_n^{-1}\right) \leq n^{-c_3}.$$

Note that Assumption 6 allows the rate  $b_n$  to in turn depend on  $\hat{\eta}$ , or more specifically,  $\mathbb{E}_X|\eta(X) - \hat{\eta}(X)|$ . Moreover, it does not necessarily require that  $\hat{\delta}$  converge to  $\delta^*$ , only that their corresponding utility function values be close.

Armed with these largely notational assumptions, we can now provide the rate for the overall data-splitting two-step plug-in classifier described in Algorithm 1:

**Theorem 5.1.** Suppose Assumption 1 and 2 hold, and further that Assumptions 4 and 6 hold for some  $\hat{\eta}$  and  $\hat{\delta}$ . Let  $\mathcal{U}^* = \mathcal{U}(\text{sign}(\eta - \delta^*), \mathbb{P})$  be the Bayes optimal utility. If we split the data as  $n_1 = n_2 = \frac{n}{2}$ , then with probability greater than  $1 - n^{-c_4}$ :

$$\mathcal{U}^* - \mathcal{U}(\text{sign}(\hat{\eta} - \hat{\delta}), \mathbb{P}) \leq c_5 \max\left\{a_n^{-\frac{1+\alpha}{2}}, b_n^{-1}\right\}.$$

where  $c_4, c_5 > 0$  are absolute constants.

### 5.1. Key Lemmas

We provide a detailed proof of the theorem in the Appendix, but provide brief vignettes via some key lemmas that also provide some additional insight into the statistical analysis. A key tool when analyzing traditional binary classification is to turn the excess risk into an expectation of the absolute deviation of conditional probability from the threshold  $\frac{1}{2}$ . We show in the following lemma that a similar result holds with general optimal threshold:

**Lemma 5.1.** Let  $C_n$  and  $C^*$  be the vectorized confusion matrices associated with  $f_n = \text{sign}(\eta_n - \delta^*)$

and  $f^* = \text{sign}(\eta - \delta^*)$  respectively, where  $\delta^*$  is the threshold for the Bayes optimal classifier. Denote  $C_G := \nabla \mathcal{G}(C^*)^T(1, -1, -1, 1)$ , and  $C_H := \max_f \|\nabla^2 \mathcal{G}(C(f))\|_{op}$ , where  $\|\cdot\|_{op}$  refers to the operator norm of a matrix. If for some constant  $c_6$ ,  $a_n \geq c_6 \left(\frac{C_H}{C_G \min\{\delta^*, 1-\delta^*\}}\right)^2$ , then

$$\mathcal{G}(C^*) - \mathcal{G}(C_n) \geq \frac{1}{2} C_G \mathbb{E}[|\eta - \delta^*| \mathbf{1}(f_n \neq f^*)],$$

$$\mathcal{G}(C^*) - \mathcal{G}(C_n) \leq \frac{3}{2} C_G \mathbb{E}[|\eta - \delta^*| \mathbf{1}(f_n \neq f^*)].$$

This lemma thus helps us control the excess utility via the error of the conditional probability estimator  $\hat{\eta} - \eta$ . Armed with this result, and additionally using Assumption 4 on the convergence rate of the conditional probability estimator, we can then show that the excess utility converges at the rate  $O(a_n^{-\frac{1+\alpha}{2}})$ :

**Lemma 5.2.** Suppose that Assumptions 4 and 5 are satisfied, and that the Bayes optimal classifier is  $f^* = \text{sign}(\eta - \delta^*)$ . Then there exists a constant  $c_7 > 0$  which depend on  $\mathcal{G}$  and  $C(f^*)$ , such that  $\mathcal{U}(\text{sign}(\eta - \delta^*)) - \mathcal{U}(\text{sign}(\eta_n - \delta^*)) \leq c_7 a_n^{-\frac{1+\alpha}{2}}$ .

Lemma 5.2 describes the classification error rate when the optimal threshold is known. Stitching this together with Assumption 6 on the convergence rate of the threshold estimator can then be shown to yield the statement of Theorem 5.1.

### 5.2. Risk Bound for the Plugin Classifier from Algorithm 2

Prior work on threshold estimation for plug-in classifiers have ranged over brute-force search (Koyejo et al., 2014) with no rates of convergence, level-set based methods (Parambath et al., 2015) for the specific class of linear fractional metrics, and Frank-Wolfe based methods (Narasimhan et al., 2015b) for the specific class of concave performance metrics. However these estimators, in addition to focus on specific performance metrics, are only able to achieve a convergence rate of  $O(\max\{\mathbb{E}\|\hat{\eta}(X) - \eta(X)\|_1, 1/\sqrt{n}\})$ . This entails that even if when the conditional probability estimator has a fast convergence rate, the final convergence rate for these estimators will still be bounded by  $O(1/\sqrt{n})$ . In this section we show that our simple threshold search procedure in Algorithm 2 achieves a fast  $O(1/n)$  or  $O(1/a_n)$  rate of convergence by leveraging our analysis from Section 3.

**Lemma 5.3.** Assume that Assumptions 1, 2, 5 hold, and that the confusion-matrix function  $\mathcal{G}$  corresponding to the performance measure  $\mathcal{U}$  satisfies the same conditions as in Proposition 3.1. Let  $\hat{\delta}$  denote the output of Algorithm 2 with sample size  $n$  and tolerance  $\tau = \frac{\log n}{n}$ , and  $\hat{\eta}$  denote

a conditional probability estimator satisfying Assumption 4 obtained on an independent sample set of size  $n$ . Denoting  $\tilde{n} = \min\{n, a_n\}$ , we then have that the rate  $b_n$  in Assumption 6 satisfies:  $b_n = \frac{\log \tilde{n}}{n}$ .

An immediate corollary then gives the excess risk for the plug-in classifier.

**Corollary 5.1.** *Suppose Assumption 3 holds. If  $\tau = \frac{\log n}{n}$ ,  $n_1 = n_2 = \frac{n}{2}$ , then there exist constants  $c_8, c_9 > 0$ , such that with probability at least  $1 - \min\{n, a_n\}^{-c_8}$ ,*

$$\mathcal{U}(f^*, \mathbb{P}) - \mathcal{U}(\hat{f}, \mathbb{P}) \leq c_9 \max \left\{ \frac{\log n}{n}, \frac{\log a_n}{a_n}, a_n^{-\frac{1+\alpha}{2}} \right\}.$$

## 6. Explicit Rates for Specific Conditional Probability Models

In this section, we analyze two special cases where we can achieve explicit rate of convergence for the conditional probability estimation. For the first example, we consider the Gaussian generative model. We will show that the rate of convergence for the excess utility obtained in Theorem 5.1 is  $O(\frac{\log n}{n})$  in this case. The second example is for non-parametric kernel estimators when the conditional probability function satisfies certain smoothness assumption.

### 6.1. Gaussian Generative Model

Consider two Gaussian distributions with the same variance, without loss of generality we assume the covariance matrix is identity  $I_d$  for both classes. We define an asymmetric mixture of two Gaussians indexed by the centers and mixing weights.

$$\begin{aligned} \mathbb{P}_{\mu, \kappa} : \quad & P(Y = 1) = \kappa, P(Y = 0) = 1 - \kappa, \\ & X|Y = 1 \sim \mathcal{N}\left(\frac{\mu}{2}, I_d\right), X|Y = 0 \sim \mathcal{N}\left(-\frac{\mu}{2}, I_d\right). \end{aligned} \quad (5)$$

As stated in Section 4, we can compute the conditional probability and show that it can be fitted with a logistic regression model. Next we present results related to the key quantities in Theorem 5.1:  $a_n$  and  $\alpha$ .

**Lemma 6.1.** *Model defined in Eq. (5) with maximum likelihood estimator satisfies Assumption 4 with  $a_n = n$ .*

The following lemma specifies the margin assumption parameter for the above model.

**Lemma 6.2.** *The Gaussian generative model defined as in Eq. (5), satisfies Assumption 5 with  $\alpha = 1$ .*

Combining this result with Theorem 5.1 gives us the following corollary.

**Corollary 6.1.** *Assume Assumptions 1-5 hold,  $\mathbb{P}$  is generated from Eq. (5). Let  $\hat{f}$  be the output of Algorithm 1 with*

$\hat{\eta}$  estimated by MLE of logistic regression. We have with probability tending to 1,  $\mathcal{U}(f^*, \mathbb{P}) - \mathcal{U}(\hat{f}, \mathbb{P}) = O\left(\frac{\log n}{n}\right)$ .

For Gaussian generative models, fast rates of  $O(\frac{1}{n})$  are only known for 0-1 loss (Li et al., 2015). The logarithm factor can be further removed under 0-1 loss, or other cases when the threshold is known, as one can apply Lemma 5.2 with  $\alpha = 1$  and get exactly the same rate as in Li et al. (2015). Corollary 6.1 generalizes this result for a much broader class of utility functions, when the threshold is unknown and estimated from data.

### 6.2. $\beta$ -Hölder Densities

When the conditional probability function belongs to the  $\beta$ -Hölder class as defined in Definition 4.1, we have the following lemma on the convergence rates of  $\eta_n$ .

**Lemma 6.3.** *For  $\beta$ -Hölder conditional probability functions, there exists estimators such that Assumption 4 holds with  $a_n = n^{\frac{2\beta}{2\beta+d}}$ .*

Examples of such estimators include locally polynomial estimators (Audibert et al., 2007), or kernel (conditional) density estimators (Jiang, 2017). Combined with Theorem 5.1 we have the following corollary.

**Corollary 6.2.** *Assume Assumptions 1-5 hold and  $\mathbb{P}$  be a distribution where  $P(Y = 1|X)$  belongs to  $\beta$ -Hölder class. With locally polynomial estimators (Audibert et al., 2007) or kernel (conditional) density estimators (Jiang, 2017), we have:  $\mathcal{U}(f^*, \mathbb{P}) - \mathcal{U}(\hat{f}) = O\left(n^{-\frac{(\min\{\alpha, 1\} + 1)\beta}{2\beta+d}}\right)$ .*

The convergence rate obtained in Corollary 6.2 is faster than  $O(\frac{1}{\sqrt{n}})$  if  $\beta > \max\{\frac{d}{2\alpha}, \frac{d}{2}\}$ . It is worth pointing out that the fast rate is obtained via a trade-off between the parameter  $\alpha$  and  $\beta$ : to have a very smooth conditional probability function  $\eta$ , i.e., a large value of  $\beta$ , it cannot deviate from the critical level very abruptly, hence  $\alpha$  has to be small.

## 7. Conclusion

We study Bayes optimal classification for general performance metrics. We derive the form of the Bayes optimal classifier, provide practical algorithms to estimate this Bayes optimal classifier, and provide novel analysis of classification error with respect to general performance metrics, and in particular show our estimators are not only consistent but have fast rates of convergence. We also provide corollaries of our general results for some special cases, such as when the inputs are drawn from a Gaussian mixture generative models, or when the conditional probability function lies in a Hölder space, explicitly proving fast rates under mild regularity conditions.



## Acknowledgements

P.R. acknowledges the support of NSF via IIS-1149803, IIS-1664720, DMS-1264033, and PNC via the PNC Center for Financial Services Innovation.

## References

- Jean-Yves Audibert, Alexandre B Tsybakov, et al. Fast learning rates for plug-in classifiers. *The Annals of statistics*, 35(2):608–633, 2007.
- Róbert Busa-Fekete, Balázs Szörényi, Krzysztof Dembczynski, and Eyke Hüllermeier. Online f-measure optimization. In *Advances in Neural Information Processing Systems*, pages 595–603, 2015.
- Sophia Daskalaki, Ioannis Kostas, and Nikolaos Avouris. Evaluation of classifiers for an uneven class distribution problem. *Applied artificial intelligence*, 20(5):381–417, 2006.
- Luc Devroye, László Györfi, and Gábor Lugosi. *A probabilistic theory of pattern recognition*, volume 31. Springer Science & Business Media, 2013.
- Qiong Gu, Li Zhu, and Zhihua Cai. Evaluation measures of the classification performance of imbalanced data sets. In *International Symposium on Intelligence Computation and Applications*, pages 461–471. Springer, 2009.
- Martin Jansche. A maximum expected utility framework for binary sequence labeling. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, pages 736–743, 2007.
- Heinrich Jiang. Uniform convergence rates for kernel density estimation. In *International Conference on Machine Learning*, pages 1694–1703, 2017.
- Thorsten Joachims. A support vector method for multivariate performance measures. In *Proceedings of the 22nd international conference on Machine learning*, pages 377–384. ACM, 2005.
- Purushottam Kar, Harikrishna Narasimhan, and Prateek Jain. Online and stochastic gradient methods for non-decomposable loss functions. In *Advances in Neural Information Processing Systems*, pages 694–702, 2014.
- Kenneth Kennedy, Brian Mac Namee, and Sarah Jane Delany. Learning without default: A study of one-class classification and the low-default portfolio problem. In *Artificial Intelligence and Cognitive Science*, pages 174–187. Springer, 2009.
- Oluwasanmi O Koyejo, Nagarajan Natarajan, Pradeep K Ravikumar, and Inderjit S Dhillon. Consistent binary classification with generalized performance metrics. In *Advances in Neural Information Processing Systems*, pages 2744–2752, 2014.
- Miroslav Kubat, Stan Matwin, et al. Addressing the curse of imbalanced training sets: one-sided selection. In *ICML*, volume 97, pages 179–186. Nashville, USA, 1997.
- David D Lewis. Evaluating and optimizing autonomous text classification systems. In *Proceedings of the 18th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 246–254. ACM, 1995.
- Tianyang Li, Adarsh Prasad, and Pradeep K Ravikumar. Fast classification rates for high-dimensional gaussian generative models. In *Advances in Neural Information Processing Systems*, pages 1054–1062, 2015.
- Enno Mammen, Alexandre B Tsybakov, et al. Smooth discrimination analysis. *The Annals of Statistics*, 27(6): 1808–1829, 1999.
- Christopher D Manning, Prabhakar Raghavan, Hinrich Schütze, et al. *Introduction to information retrieval*, volume 1. Cambridge university press Cambridge, 2008.
- Aditya Menon, Harikrishna Narasimhan, Shivani Agarwal, and Sanjay Chawla. On the statistical consistency of algorithms for binary classification under class imbalance. In *International Conference on Machine Learning*, pages 603–611, 2013.
- Ye Nan, Kian Ming Chai, Wee Sun Lee, and Hai Leong Chieu. Optimizing f-measure: A tale of two approaches. *arXiv preprint arXiv:1206.4625*, 2012.
- Harikrishna Narasimhan, Rohit Vaish, and Shivani Agarwal. On the statistical consistency of plug-in classifiers for non-decomposable performance measures. In *Advances in Neural Information Processing Systems*, pages 1493–1501, 2014.
- Harikrishna Narasimhan, Purushottam Kar, and Prateek Jain. Optimizing non-decomposable performance measures: A tale of two classes. In *32nd International Conference on Machine Learning (ICML)*, 2015a.
- Harikrishna Narasimhan, Harish Ramaswamy, Aadirupa Saha, and Shivani Agarwal. Consistent multiclass algorithms for complex performance measures. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pages 2398–2407, 2015b.

Andrew Y Ng and Michael I Jordan. On discriminative vs. generative classifiers: A comparison of logistic regression and naive bayes. In *Advances in neural information processing systems*, pages 841–848, 2002.

Shameem A Puthiya Parambath, Nicolas Usunier, and Yves Grandvalet. Theory of optimizing pseudolinear performance measures: Application to f-measure. *arXiv preprint arXiv:1505.00199*, 2015.

Wolfgang Polonik. Measuring mass concentrations and estimating density contour clusters-an excess mass approach. *The Annals of Statistics*, pages 855–881, 1995.

Cornelis Joost Van Rijsbergen. Foundation of evaluation. *Journal of Documentation*, 30(4):365–373, 1974.

Byron C Wallace, Kevin Small, Carla E Brodley, and Thomas A Trikalinos. Class imbalance, redux. In *Data Mining (ICDM), 2011 IEEE 11th International Conference on*, pages 754–763. IEEE, 2011.

William J Youden. Index for rating diagnostic tests. *Cancer*, 3(1):32–35, 1950.

Ming-Jie Zhao, Narayanan Edakunni, Adam Pocock, and Gavin Brown. Beyond fano’s inequality: bounds on the optimal f-score, ber, and cost-sensitive risk and their implications. *The Journal of Machine Learning Research*, 14(1):1033–1090, 2013.