
Estimating the Arc Length of the Optimal ROC Curve and Lower Bounding the Maximal AUC

Song Liu*

School of Mathematics
University of Bristol
Bristol, BS8 1UG, UK
song.liu@bristol.ac.uk

Abstract

In this paper, we show the arc length of the optimal ROC curve is an f -divergence. By leveraging this result, we express the arc length using a variational objective and estimate it accurately using positive and negative samples. We show this estimator has a non-parametric convergence rate $O_p(n^{-\beta/4})$ ($\beta \in (0, 1]$ depends on the smoothness). Using the same technique, we show the surface area between the optimal ROC curve and the diagonal can be expressed via a similar variational objective. These new insights lead to a novel classification procedure that maximizes an approximate lower bound of the maximal AUC. Experiments on CIFAR-10 datasets show the proposed two-step procedure achieves good AUC performance in imbalanced binary classification tasks.

1 Introduction

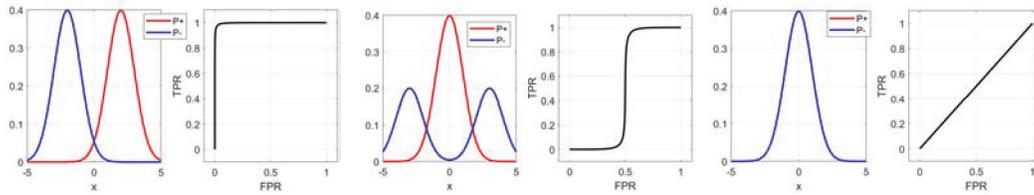
The study of Receiver operating characteristic (ROC) curves has a long history in medicine [26], psychology [15] and radiology [16, 13]. In machine learning, ROC curves have been primarily used to analyze the performance of different classification algorithms [8, 9]. Indeed, the Area Under the Curve (AUC) encodes a classifier's ranking accuracy, making it a preferable performance metric for imbalanced class classification [8, 3]. In recent years, ROC curves have also been used in comparing two distributions and achieved promising results. Examples include analyzing the mode collapsing issue of Generative Adversarial nets (GAN) [25], and diagnosing the performance of an amortized Markov Chain Monte Carlo [18].

In applications that require computing statistical discrepancy between distributions (e.g. GAN [14] or Variational Inference (VI) [1]), f -divergences are widely used discrepancy measures. The family of f -divergences includes Kullback-Leibler divergence [23] and Total Variation distance. It has been shown that f -divergences, generally, can be expressed via variational objectives and efficiently approximated from empirical samples [29, 30].

Since the ROC curves are used as performance metrics in many two sample applications, are they in any way related to f -divergences? For example, can AUC be an f -divergence between positive and negative data distributions given some classification score function? An earlier investigation proves that the answer is no when the score function is the likelihood ratio [31]. Nonetheless, this result inspired us to look for f -divergence from other geometries of the ROC curve.

In this paper, we show that, when using the likelihood ratio score, a novel f -divergence arises from the *arc length* of the corresponding ROC curve. By leveraging this result, we can express the arc length using a variational objective and approximate it using only samples from two distributions.

*Webpage: <https://allmodelsarewrong.net>



(a) different distributions, long ROC (b) different distributions, long ROC (c) same distribution, short ROC

Figure 1: ROC curves generated for one dimensional datasets using the identity classification score function $t(x) = x$. Notice that the arc length of ROC curves seem to be a good indication on how different the positive (red) and negative (blue) data distributions are.

We show this arc length estimator is also a consistent estimator to the arctangent of likelihood ratio and has a non-parametric convergence rate $O_p(n^{-\beta/4})$, where $\beta \in (0, 1]$ depends on the smoothness of the true arctangent likelihood ratio. Moreover, by parameterizing the ROC curves of positive and negative *mixture* distributions, the surface area between the optimal ROC curve and the diagonal can be expressed via a similar variational objective. With the help of our arctangent ratio estimator, we can approximately maximize a lower bound to this surface area. We point out the similarity between this lower bound maximization and the classic AUC maximization [3]. We show our approximated optimal score achieves comparable performance to a state of the art AUC maximizer in an imbalanced classification problem on CIFAR-10 dataset.

2 Background

2.1 An Illustrative Example

ROC curves are frequently used to visualize binary classification performance, and we often rely on the Area Under the Curve (AUC) as a numerical metric for selecting a good classifier. In this section, we highlight an often overlooked ROC geometry: arc length. We illustrate its potential as a good discrepancy measure between positive and negative data distributions.

Let us consider the one-dimensional distributions and the ROC curves in Figure 1. In this example, ROC curves are generated using the identity score function $t(x) = x$.

In both (a) and (b), p_+ and p_- are quite different, and thus the discrepancies between positive and negative distributions should be large in both cases. In (c), the densities p_+ and p_- are the same and consequently their discrepancy should be smaller than that of both (a) and (b). Further, notice that both (a) and (b) have long ROC curves (≈ 2) while (c) has a shorter one ($\approx \sqrt{2}$). This example suggests that the more similar p_+ and p_- are, the shorter the ROC curve is.

However, we can see that $t(x) = x$ is a special choice: If $t(x) = 0$, the arc length will not reflect any discrepancy between data distributions at all. This observation inspires the following questions: Why is the arc length in this example good at telling the differences between two data distributions? Are there other score functions whose ROC arc lengths are also good discrepancy measures? What are the practical applications of studying the arc length of ROC? In the following sections, we study the arc length of a ROC curve under a probabilistic framework and provide answers to these questions.

2.2 ROC Curve in a Probabilistic Setting

Suppose we have positive and negative datasets $X_+ := \{\mathbf{x}_i^+\}_{i=1}^{n_+}$ and $X_- := \{\mathbf{x}_i^-\}_{i=1}^{n_-}$ drawn from two distributions \mathbb{P}_+ and \mathbb{P}_- respectively. These distributions have respective probability density functions $p_+(\mathbf{x})$ and $p_-(\mathbf{x})$ that are both defined on the domain $\mathcal{X} \subseteq \mathbb{R}^d$. A classification score function (score function for short) takes a sample \mathbf{x} as input and outputs a real-valued score. Suppose we have a score function $t(\mathbf{x}) \in \mathbb{R}$. We classify \mathbf{x} as positive if $t(\mathbf{x}) > \tau$, where τ is a threshold.

Let F_+ and F_- denote the Cumulative Distribution Functions (CDFs) of $t(\mathbf{x}^+)$ and $t(\mathbf{x}^-)$ respectively. Then we can define the following quantities:

- False Positive Rate (FPR) at threshold τ , $\tilde{F}_-(\tau) := 1 - F_-(\tau)$
- True Positive Rate (TPR) at threshold τ , $\tilde{F}_+(\tau) := 1 - F_+(\tau)$
- The ROC curve of a score function t : the graph of function $\tilde{F}_+[\tilde{F}_-^{-1}(s)]$, where $s \in [0, 1]$.

The above definition of ROC curve requires F_- to be strictly increasing. In this paper, we assume F_+ , F_- to be both strictly increasing. Obviously, both F_+ and F_- depend on the choice of score function t .

2.3 Arc Length of ROC Curve

Due to the strict monotonicity of F_+ and F_- , \tilde{F}_+ and \tilde{F}_- form a bijective parameterization of the ROC curve in the sense that each point on this ROC curve can be written as $(\tilde{F}_-(\tau_0), \tilde{F}_+(\tau_0))$ for a unique $\tau_0 \in \mathbb{R}$. Using the line integral formula, the arc length of an ROC curve for a fixed score function t can be expressed using the derivatives of \tilde{F}_- and \tilde{F}_+ :

$$\widehat{\text{ROC}}(t) := \int_{-\infty}^{\infty} \sqrt{[\partial_\tau \tilde{F}_+(\tau)]^2 + [\partial_\tau \tilde{F}_-(\tau)]^2} d\tau = \int_{-\infty}^{\infty} \sqrt{f_+(\tau)^2 + f_-(\tau)^2} d\tau, \quad (1)$$

where $f_+(\tau)$ and $f_-(\tau)$ are the density functions of $t(\mathbf{x}^+)$ and $t(\mathbf{x}^-)$ respectively. Although (1) is an elementary result, it has been seldom discussed in the ROC literature. Authors in [5, 6] have proposed a performance metric computed over the ‘‘ROC hypersurface’’ and (1) is used to justify such a metric in a binary classification setting.

Using (1), we can confirm a simple geometric fact whose proof can be found in Appendix A:

Proposition 1. $\widehat{\text{ROC}}(t) \in [\sqrt{2}, 2]$, for all t .

This result reflects the geometric observation that any monotone curve (such as ROC curve) starts and ends at two opposite corners of the ROC space $[0, 1]^2$ has an arc length between $\sqrt{2}$ and 2.

3 f -divergences Arising from ROC Arc Length

3.1 f -divergence of Score Distributions

Among many discrepancies measures, f -divergence has been widely used in many applications.

Definition 1. Let p and q be densities of two continuous distributions. An f -divergence is defined as:

$$D_g(p|q) := \mathbb{E}_q \left[g \left(\frac{p(\mathbf{z})}{q(\mathbf{z})} \right) \right], \text{ where } g \text{ is convex and lower-semicontinuous satisfying } g(1) = 0.$$

Now let us slightly rewrite (1). Assuming f_- is strictly positive (in which case F_- is strictly increasing), we can write

$$\widehat{\text{ROC}}(t) - \sqrt{2} = \mathbb{E}_{f_-} \sqrt{\left[\frac{f_+(\tau)}{f_-(\tau)} \right]^2 + 1} - \sqrt{2} = \mathbb{E}_{f_-} \left[g \left(\frac{f_+(\tau)}{f_-(\tau)} \right) \right], \quad (2)$$

where $g(s) = \sqrt{s^2 + 1} - \sqrt{2}$. Equation (2) yields the first important result of this paper: $\widehat{\text{ROC}}(t) - \sqrt{2}$ is an f -divergence between score densities f_+ and f_- since $g(s)$ is a convex function and $g(1) = 0$.

This result confirms, for any given t , $\widehat{\text{ROC}}(t)$ is a good discrepancy for measuring positive and negative scores (i.e., score distributions). It also explains why the ROC arc length in Figure 1 is a good discrepancy measure: Given $t(x) = x$, the score distributions are same as the data distributions in each plot. The arc length of ROCs in Figure 1 plots are f -divergences of the score distributions hence are also f -divergences of data distributions.

Although $\widehat{\text{ROC}}(t) - \sqrt{2}$ is an f -divergence of score distributions, it is not an f -divergence of the positive and negative data distributions for general t . The choice $t(x) = x$ in the toy example does not have simple analogues for higher dimensional datasets. Are there choices of t such that the arc lengths of their ROC curves are good discrepancy measures on data distribution? In what follows, we show when t is a bijective transform of the likelihood ratio $\frac{p_+(\mathbf{x})}{p_-(\mathbf{x})}$, $\widehat{\text{ROC}}(t)$ encodes the differences between $p_+(\mathbf{x})$ and $p_-(\mathbf{x})$ in the form of an f -divergence between $p_+(\mathbf{x})$ and $p_-(\mathbf{x})$.

3.2 f -divergences of Data Distributions

Using the law of the unconscious statistician, we can express $\widehat{\text{ROC}}(t)$ in terms of an expectation with respect to the negative data density $p_-(\mathbf{x})$: $\mathbb{E}_{p_-} \sqrt{\left[\frac{f_+(t(\mathbf{x}))}{f_-(t(\mathbf{x}))}\right]^2 + 1}$. Consider a special family of score functions: $t^*(\mathbf{x}) = \gamma\left(\frac{p_+(\mathbf{x})}{p_-(\mathbf{x})}\right)$ where γ is any strictly increasing function. Due to the Neyman-Pearson lemma [27], $\text{ROC}(t^*)$ has the highest TPR at any FPR level. Geometrically speaking, they dominate all other ROC curves in an ROC plot and have the maximal AUC. Hence, we refer to t^* as the optimal score and $\text{ROC}(t^*)$ as the optimal ROC curve. For convenience, we denote the $\text{ROC}(t^*)$ as ROC^* which reads “rock star”. It can be shown that

$$\frac{f_+(t^*(\mathbf{x}_0))}{f_-(t^*(\mathbf{x}_0))} = \frac{\int_{\mathbf{x}:\gamma\left(\frac{p_+(\mathbf{x})}{p_-(\mathbf{x})}\right)=t^*(\mathbf{x}_0)} p_+(\mathbf{x})d\mathbf{x}}{\int_{\mathbf{x}:\gamma\left(\frac{p_+(\mathbf{x})}{p_-(\mathbf{x})}\right)=t^*(\mathbf{x}_0)} p_-(\mathbf{x})d\mathbf{x}} = \frac{p_+(\mathbf{x}_0)}{p_-(\mathbf{x}_0)}, \quad \forall \gamma \quad (3)$$

where the second equality holds due to $\frac{\int_D a(\mathbf{x})d\mathbf{x}}{\int_D b(\mathbf{x})d\mathbf{x}} = \gamma^{-1}(C)$, when $\gamma\left(\frac{a(\mathbf{x})}{b(\mathbf{x})}\right) \equiv C$, $\forall \mathbf{x} \in D$. When $\gamma(\mathbf{x}) = \mathbf{x}$, (3) expresses a known result [7], and is often given in plain English as “the density ratio of the likelihood ratio score is the likelihood ratio itself”.

Finally, $\widehat{\text{ROC}}^*$ takes an elegant form free from t^* or γ : $\mathbb{E}_{p_-(\mathbf{x})} \sqrt{\left[\frac{p_+(\mathbf{x})}{p_-(\mathbf{x})}\right]^2 + 1}$. Equivalently,

$$\widehat{\text{ROC}}^* - \sqrt{2} = \mathbb{E}_{p_-(\mathbf{x})} \left[g\left(\frac{p_+(\mathbf{x})}{p_-(\mathbf{x})}\right) \right] \quad (4)$$

We can see that the same f -divergence arises from computing the arc length of ROC^* . However, unlike the f -divergence given in (2), (4) is an f -divergence between *data* distributions, not score distributions. It shows that as long as we use the optimal scores, the arc length of the optimal ROC can indeed reflect the differences between *data* distributions. From now on, we will refer to $\widehat{\text{ROC}}^* - \sqrt{2}$ as the ROC divergence. To the best of our knowledge, (4) has not been presented in literature before.

By definition, the ROC divergence is symmetric. Moreover, as a result of Proposition 1, the ROC divergence is upper bounded by $2 - \sqrt{2}$ and lower bounded by 0. Some geometric properties of ROC^* can be found in Section B.

4 Estimating the Arc Length of ROC^*

4.1 A Variational Objective

To numerically approximate the arc length using samples alone, we leverage that $\widehat{\text{ROC}}^* - \sqrt{2}$ is an f -divergence. Utilizing Fenchel’s duality [20], authors in [29] show that an f -divergence $D_g(p_+|p_-)$ has a variational representation:

$$\begin{aligned} D_g(p_+|p_-) &= \int_{\mathcal{X}} p_-(\mathbf{x}) g\left[\frac{p_+(\mathbf{x})}{p_-(\mathbf{x})}\right] d\mathbf{x} = \int_{\mathcal{X}} p_-(\mathbf{x}) \sup_u \left\{ u(\mathbf{x}) \cdot \left[\frac{p_+(\mathbf{x})}{p_-(\mathbf{x})}\right] - g'[u(\mathbf{x})] \right\} d\mathbf{x} \\ &= \sup_u \int_{\mathcal{X}} p_+(\mathbf{x}) u(\mathbf{x}) - \int_{\mathcal{X}} p_-(\mathbf{x}) g'[u(\mathbf{x})] d\mathbf{x}, \end{aligned}$$

where g' is the convex conjugate of g and the supremum is taken over all measurable functions. In the case of the ROC divergence, $g(z) = \sqrt{z^2 + 1} - \sqrt{2}$ and $z \in [0, \infty]$ thus g has a convex conjugate $g'(z') = -\sqrt{1 - z'^2} - \sqrt{2}$, $z' \in [0, 1]$. Rewriting $\widehat{\text{ROC}}^* - \sqrt{2}$ using the above variational representation and dropping the $-\sqrt{2}$, we obtain:

$$\widehat{\text{ROC}}^* = \sup_{u \in [0,1]} \mathbb{E}_{p_+}[u(\mathbf{x})] + \mathbb{E}_{p_-}[\sqrt{1 - u^2(\mathbf{x})}].$$

We reparameterize $u(\mathbf{x}) = \sin[v(\mathbf{x})]$, where $v \in [0, \pi/2]$:

$$\widehat{\text{ROC}}^* = \sup_{v \in [0, \pi/2]} \mathbb{E}_{p_+} \sin[v(\mathbf{x})] + \mathbb{E}_{p_-} \cos[v(\mathbf{x})]. \quad (5)$$

Differentiating the objective in (5) for v and setting the derivative to zero, we can see the supremum is attained at $v^* = \text{atan} \frac{p_+}{p_-}$. In other words, the optimal v^* is the arctangent likelihood ratio function.

It is also interesting to see how v^* is visualized in the ROC plot. We can see the tangent of $\widehat{\text{ROC}}^*$ at an FPR level $s_0 \in [0, 1]$ is

$$\partial_s \tilde{F}_+(\tilde{F}_-^{-1}(s_0)) = \frac{f_+(\tilde{F}_-^{-1}(s_0))}{f_-(\tilde{F}_-^{-1}(s_0))} = \frac{p_+(\mathbf{x}_0)}{p_-(\mathbf{x}_0)}, \quad (6)$$

where \mathbf{x}_0 is any point in \mathcal{X} that satisfies the equality $\gamma \left(\frac{p_+(\mathbf{x}_0)}{p_-(\mathbf{x}_0)} \right) = \tilde{F}_-^{-1}(s_0)$. (6) is a known result [9]. In other words, $v^* = \text{atan} \frac{p_+}{p_-}$ is the *slope angle* of $\widehat{\text{ROC}}^*$. See Figure 2 for a visualization of the tangent of $\widehat{\text{ROC}}^*$ expressed by the likelihood ratio.

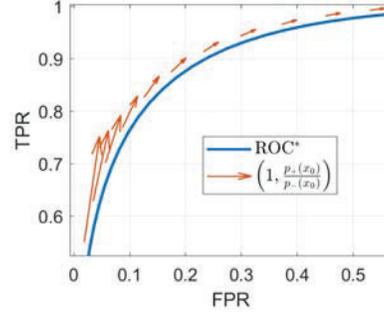


Figure 2: $\widehat{\text{ROC}}^*$ and its tangent marked by vector $\left(1, \frac{p_+(\mathbf{x}_0)}{p_-(\mathbf{x}_0)}\right)$ (scaled to fit). $p_+ = \mathcal{N}(1, 1), p_- = \mathcal{N}(-1, 1)$.

Moreover, using (5), we can obtain a relationship between $\widehat{\text{ROC}}^*$ and the total variation distance between \mathbb{P}_+ and \mathbb{P}_- (denoted as $\text{TV}(\mathbb{P}_+, \mathbb{P}_-)$).

Proposition 2. $\text{TV}(\mathbb{P}_+, \mathbb{P}_-)$ has a lower and upper bound expressed via $\widehat{\text{ROC}}^*$:

$$\max_{a \in [0, 1]} \frac{2}{\pi} \left[\frac{\widehat{\text{ROC}}^* - 2\sqrt{1 - a^2}}{a} + \arccos(a) - \arcsin(a) \right] \leq \text{TV}(\mathbb{P}_+, \mathbb{P}_-) \leq \widehat{\text{ROC}}^* - 1.$$

Proof of this proposition can be found in Section C. This proposition justifies that $\widehat{\text{ROC}}^*$ is a valid measure of the discrepancy between \mathbb{P}_+ and \mathbb{P}_- . We compare this bound with other known TV bounds in Section 6.1.

4.2 A Tractable Objective for Estimating $\text{atan} \frac{p_+}{p_-}$

To use (5) in practice, we need to find an appropriate function class \mathcal{F} . We can simply restrict v to a bounded (parametric/non-parametric) function class \mathcal{F} and solve the sample version of (5):

$$\max_{v \in [0, \pi/2], v \in \mathcal{F}} \frac{1}{n_+} \sum_{i=1}^{n_+} \sin(v(\mathbf{x}_i^+)) + \frac{1}{n_-} \sum_{i=1}^{n_-} \cos(v(\mathbf{x}_i^-)). \quad (7)$$

In practice, enforcing the boundedness $v \in [0, \pi/2]$ over \mathcal{X} is difficult. We can relax (7) by only enforcing the boundedness constraint of v on the sample dataset $X_+ \cup X_-$.

For example, by letting \mathcal{F} be a Reproducing Kernel Hilbert Space (RKHS) [33], we can translate (7) into the following optimization problem:

$$\begin{aligned} \hat{v} &:= \underset{v \in \mathcal{H}}{\text{argmin}} \ell(v) + \frac{\lambda}{2} \|v\|_{\mathcal{H}}^2, \quad \ell(v) := -\frac{1}{n_+} \sum_{i=1}^{n_+} \sin\langle v, \varphi(\mathbf{x}_i^+) \rangle - \frac{1}{n_-} \sum_{i=1}^{n_-} \cos\langle v, \varphi(\mathbf{x}_i^-) \rangle \\ \text{s.t.} & \langle v, \varphi(\mathbf{x}) \rangle \in \left[0, \frac{\pi}{2}\right], \quad \forall \mathbf{x} \in X_+ \cup X_-, \end{aligned} \quad (8)$$

where \mathcal{H} is a RKHS with a positive definite kernel $k(\mathbf{x}, \mathbf{x}') = \langle \varphi(\mathbf{x}), \varphi(\mathbf{x}') \rangle$, $\|\cdot\|_{\mathcal{H}}$ is the RKHS norm and $\frac{\lambda}{2} \|v\|_{\mathcal{H}}^2$ is the regularization term. The optimizer $\langle \hat{v}, \varphi(\mathbf{x}) \rangle$ is an estimation of $v^*(\mathbf{x})$, the arctangent of the likelihood ratio. (8) is a strictly convex optimization and thus, if a solution \hat{v} exists, it must be unique.

Instead of modelling $\text{atan} \frac{p_+}{p_-}$, we can opt for modelling the log likelihood ratio $\log \frac{p_+}{p_-}$. However, this modelling choice results in a non-convex optimization thus presents extra challenges in the theoretical analysis. Details can be found in Section K.

4.3 Finite Sample Guarantee

We show that the solution of (8), $\langle \hat{v}, \varphi(\mathbf{x}) \rangle$ converges to the true arctangent likelihood ratio (or its projection onto \mathcal{H}) as the number of samples goes to infinity. Below are a few regularity conditions.

Assumption 1. *There exists a unique $v^* \in \mathcal{H}$, such that $\mathbb{E}[\nabla_v \ell(v^*)] = 0$ and $\langle v^*, \varphi(\mathbf{x}) \rangle \in [0, \pi/2]$ holds for all $\mathbf{x} \in \mathcal{X}$.*

A sufficient condition of the above condition is specified in the following proposition.

Proposition 3. *If there exists a unique $v^* \in \mathcal{H}$, such that $\langle v^*, \varphi(\mathbf{x}) \rangle = \text{atan} \left[\frac{p_+(\mathbf{x})}{p_-(\mathbf{x})} \right]$ then Assumption 1 holds.*

The proof can be found in Appendix F. Proposition 3 states if model is correctly specified and identifiable then Assumption 1 holds. It is possible that there exists a $v^* \in \mathcal{H}$ satisfying $\mathbb{E}[\nabla_v \ell(v^*)] = 0$ which does not meet the boundedness constraint $[0, \pi/2]$. In this paper we only consider situations where Assumption 1 holds, which includes all situations where the model is correctly specified and some situations where the model is misspecified.

Assumption 2. *Let $n_{\min} = \min(n_+, n_-)$. There exists a subspace $\mathcal{H}^* := \{v \in \mathcal{H} \mid \|v - v^*\|_{\mathcal{H}}^2 \leq \delta_{n_{\min}}^2\}$, such that $\forall v \in \mathcal{H}^*, \forall \mathbf{x} \in X_+ \cup X_-, \langle v, \varphi(\mathbf{x}) \rangle \in (0, \frac{\pi}{2})$ holds with high probability. The sequence $\delta_{n_{\min}}$ is monotonically decreasing as n_{\min} grows to infinity.*

Assumption 2 states all v within a vicinity of v^* are in the interior of (8)'s feasible region with high probability. The following proposition gives a sufficient condition under which Assumption 2 holds.

Proposition 4. *Suppose $\|\varphi(\mathbf{x})\|_{\mathcal{H}} \leq 1$. If our model is correctly specified as described in Proposition 3 and $\forall \mathbf{x} \in \mathcal{X}$, $\text{atan} \left[\frac{p_+(\mathbf{x})}{p_-(\mathbf{x})} \right] \in [R_1, R_2]$, for some R_1 and R_2 such that $\frac{\pi}{2} > R_2 > R_1 > 0$, then there exists $N > 0$ such that Assumption 2 holds when $n_{\min} > N$.*

The proof can be found in Appendix D. Since we use RKHS as the estimator function class, our final assumption is that v^* should be reasonably smooth. In previous works such an assumption depends on the decay of the integral operator's eigenvalues [36, 10]. In this paper, we measure the smoothness using the *range space* technique which has been recently adopted in [11, 35]. We define

$$\Sigma_v := \mathbb{E}[\nabla_v^2 \ell(v)] = \mathbb{E}_{p_+}[\sin\langle v, \varphi(\mathbf{x}) \rangle \cdot \varphi(\mathbf{x}, \cdot) \otimes \varphi(\mathbf{x}, \cdot)] + \mathbb{E}_{p_-}[\cos\langle v, \varphi(\mathbf{x}) \rangle \cdot \varphi(\mathbf{x}, \cdot) \otimes \varphi(\mathbf{x}, \cdot)],$$

where \otimes denotes the outer product. Given $v_0 \in \mathcal{H}$, Σ_{v_0} is an integral operator on $u \in \mathcal{H}$ and

$$\Sigma_{v_0} u = \mathbb{E}_{p_+}[\sin\langle v_0, \varphi(\mathbf{x}) \rangle \cdot \varphi(\mathbf{x}, \cdot) \cdot u(\mathbf{x})] + \mathbb{E}_{p_-}[\cos\langle v_0, \varphi(\mathbf{x}) \rangle \cdot \varphi(\mathbf{x}, \cdot) \cdot u(\mathbf{x})].$$

By definition Σ_{v_0} is a positive, self-adjoint operator, in the sense that $\langle u, \Sigma_{v_0} u \rangle \geq 0$, $\langle u, \Sigma_{v_0} v \rangle = \langle \Sigma_{v_0} u, v \rangle, \forall v, u \in \mathcal{H}$. Moreover, some algebra shows that Σ_{v_0} is also a bounded and compact operator. See Section G for more details.

Next, we assume the true arctangent ratio function (or its projection) is in the range space of Σ_{v^*} .

Assumption 3. *Let $\mathcal{R}(\Sigma_{v^*})$ denote the range space of Σ_{v^*} . There exists $0 < \beta \leq 1$, $v^* \in \mathcal{R}(\Sigma_{v^*}^\beta)$, where C^β is the fraction power of a compact, positive and self-adjoint operator C .*

Note that the larger β is, the smoother the functions in the range space are. More discussions on the range space assumption can be found in Section 4.2, [35]. Now we are ready to state our theorem:

Theorem 1 (Convergence Rate of \hat{v}). *Suppose Assumptions 1, 2 and 3 hold and \hat{v} exists. If $\|\varphi(\mathbf{x})\|_{\mathcal{H}} \leq 1$ and*

$$\lambda = \frac{T}{n_{\min}^{1/4}}, \quad \frac{K}{n_{\min}^{\beta/4}} \leq \delta_{n_{\min}} \leq \frac{4}{\max(B_+, B_-)},$$

where $B_+ = \|(\Sigma_{v^*} + \lambda \mathbf{I})^{-1} \mathbb{E}_{p_+}[\varphi(\mathbf{x})]\|_{\mathcal{H}}$, $B_- = \|(\Sigma_{v^*} + \lambda \mathbf{I})^{-1} \mathbb{E}_{p_-}[\varphi(\mathbf{x})]\|_{\mathcal{H}}$ and $T \geq 1, K > 0$ are constants that do not depend on n_{\min} , then there exists a constant $N > 0$ such that when $n_{\min} > N$, $\|\hat{v} - v^*\|_{\mathcal{H}} = O_p(n_{\min}^{-\beta/4})$.

The proof can be found at Appendix E. Theorem 1 shows that, under mild assumptions, \hat{v} is indeed a good estimator for $\text{atan} \left[\frac{p_+(\mathbf{x})}{p_-(\mathbf{x})} \right]$. Some discussions comparing our results with convergence results

proved by Nguyen et al. [28] can be found in Section M. Since $\text{atan} \left[\frac{p_+(\mathbf{x})}{p_-(\mathbf{x})} \right]$ is an optimal score that gives rise to ROC^* , our estimator may have some interesting applications such as outlier detection [19] or Neyman-Pearson classification [37]. We will defer discussions on those applications in future works. In the next section, we employ our arctangent likelihood ratio estimator in an application of lower bounding the maximal AUC.

5 Approximately Lower Bounding the Maximal AUC

Finding a score function t that approximately maximizes AUC is an important task in binary classification. Let us denote the AUC of $\text{ROC}(t)$ as $\text{AUC}(t)$. It can be seen that

$$\text{AUC}(t) = \int_{[0,1]} \tilde{F}_+(\tilde{F}_-^{-1}(s)) ds = \mathbb{E}_{p_-} \mathbb{E}_{p_+} [\mathbb{1}(t(\mathbf{x}_+) \geq t(\mathbf{x}_-))].$$

Due to the Neyman-Pearson lemma, ROC^* has the maximum AUC among all ROC curves. Denote the AUC of ROC^* as AUC^* . Consider the following inequalities:

$$\text{AUC}^* = \underbrace{\sup_t \mathbb{E}_{p_-} \mathbb{E}_{p_+} [\mathbb{1}(t(\mathbf{x}^+) \geq t(\mathbf{x}^-))]}_{(i)} \geq \sup_{t \in \mathcal{F}'} \underbrace{\mathbb{E}_{p_-} \mathbb{E}_{p_+} [L(t(\mathbf{x}^+), t(\mathbf{x}^-))]}_{(ii)}, \quad (9)$$

where $L(a, b)$ is a continuous and concave lower bound of the indicator function $\mathbb{1}(a > b)$. Due to the Neyman-Pearson lemma, the supremum of (i) is only attained when $t(\mathbf{x}) = \gamma \left(\frac{p_+(\mathbf{x})}{p_-(\mathbf{x})} \right)$ where γ is a strictly increasing function. Replacing the expectations in (ii) with sample averages yields the optimization problem of AUC maximization [3, 12]:

$$\max_{t \in \mathcal{F}'} \frac{1}{n_- n_+} \sum_{i=1}^{n_-} \sum_{j=1}^{n_+} L(t(\mathbf{x}_j^+), t(\mathbf{x}_i^-)). \quad (10)$$

The objective above is also referred to as Wilcoxon-Mann-Whitney statistic [16]. Therefore, we can see that AUC maximization is a procedure that approximates an optimal score function by maximizing a lower bound of AUC^* .

Now, we show a different way of lower bounding AUC^* with the help of $\text{atan} \left(\frac{p_+(\mathbf{x})}{p_-(\mathbf{x})} \right)$. We have seen that how $(\tilde{F}_-(\tau), \tilde{F}_+(\tau))$ parameterizes an ROC curve in Section 2.3. In fact, the area between ROC^* and the diagonal line from $(0, 0)$ to $(1, 1)$ can be similarly parameterized by considering ROC curves of positive and negative *mixture* score distributions.

Let F_+^* and F_-^* denote CDFs of any optimal score. For $\alpha \in [0, .5]$, we can define CDFs of α -mixtures of F_+^* and F_-^* as follows

$$\begin{aligned} F_-^*(\tau, \alpha) &:= (1 - \alpha) F_-^*(\tau) + \alpha F_+^*(\tau), \\ F_+^*(\tau, \alpha) &:= \alpha F_-^*(\tau) + (1 - \alpha) F_+^*(\tau). \end{aligned}$$

Then, FPR $(\tilde{F}_-^*(\tau, \alpha))$ and TPR $(\tilde{F}_+^*(\tau, \alpha))$ for these α -mixtures can be defined accordingly. We visualize densities of $F_+^*(\alpha)$ and $F_-^*(\alpha)$ for different α in Figure 3.

Further, we can see that the 2-dimensional coordinate $\mathbf{r}(\tau, \alpha) := (\tilde{F}_-^*(\tau, \alpha), \tilde{F}_+^*(\tau, \alpha))$ parameterizes the area between ROC^* and the diagonal line in $[0, 1]^2$:

- When fixing α and varying τ , the coordinates give rise to a smooth curve in ROC space from $[0, 0]$ to $[1, 1]$.
 - When $\alpha = 0$, such a curve is ROC^* .
 - When $\alpha = .5$, such a curve is the diagonal line.
- When fixing $\tau = \tau_0$ and varying α , the coordinates produce a straight line segment connecting $(\tilde{F}_-^*(\tau_0, 0), \tilde{F}_+^*(\tau_0, 0))$ and $(\tilde{F}_-^*(\tau_0, .5), \tilde{F}_+^*(\tau_0, .5))$.

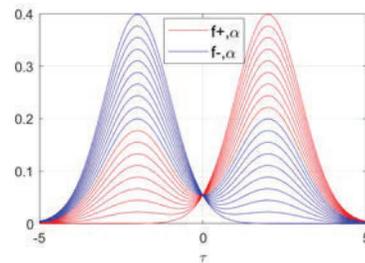


Figure 3: Densities $f_+^*(\tau, \alpha)$ and $f_-^*(\tau, \alpha)$ for $\alpha \in [0, .5]$.

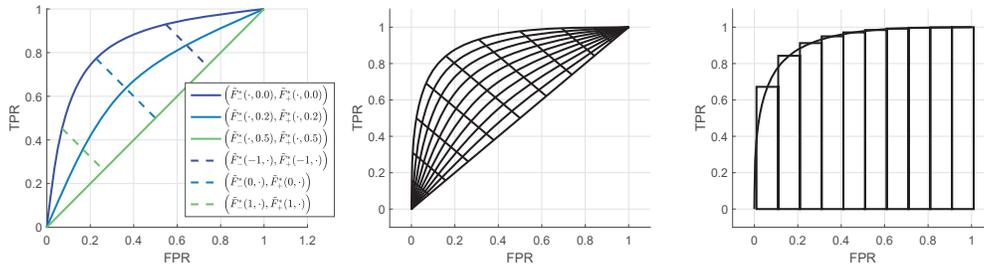


Figure 4: Left: $(\tilde{F}_-^*(\tau, \alpha), \tilde{F}_+^*(\tau, \alpha))$ parameterizes the surface between ROC^* and the diagonal line in $[0, 1]^2$. This plot is created by setting $p_+ = \mathcal{N}(1, 1)$, $p_- = \mathcal{N}(-1, 1)$ and $t^*(x) = \frac{1}{2} \log \frac{p_+(x)}{p_-(x)} = x$. Middle: Our parameterization “mesh” divides AUC into surface elements (small patches on the plot). Right: Wilcoxon-Mann-Whitney statistic divides AUC into histogram bars.

The left plot in Figure 4 visualizes this parameterization. Now the surface area sandwiched between ROC^* and the diagonal line can be computed using a surface integration:

$$\text{AUC}^* - .5 = \int_{\text{dom}(\tau)} \int_{[0, .5]} \|\partial_\tau \mathbf{r}(\tau, \alpha) \times \partial_\alpha \mathbf{r}(\tau, \alpha)\| \, d\alpha d\tau, \quad (11)$$

where \times denotes the cross product. After some algebra and applying the Fenchel duality technique in Section 4.1, we prove that AUC^* can be expressed as the supremum of a variational objective similar to (5):

Proposition 5. $\text{AUC}^* = \frac{\sqrt{2}A}{2} + \frac{1}{2}$,

$$A := \sup_{v \in [0, \pi/2]} \mathbb{E}_{p_+} \left[w \left(\text{atan} \frac{p_+(\mathbf{x})}{p_-(\mathbf{x})} \right) \sin[v(\mathbf{x})] \right] + \mathbb{E}_{p_-} \left[w \left(\text{atan} \frac{p_+(\mathbf{x})}{p_-(\mathbf{x})} \right) \cos[v(\mathbf{x})] \right], \quad (12)$$

where $w(\tau) := \sin(\tau + \frac{\pi}{4}) \cdot |F_+^*(\tau) - F_-^*(\tau)|$. The supremum of (12) is attained at $v^* = \text{atan} \frac{p_+}{p_-}$.

The proof can be found in Appendix H in the supplementary material. A lower bound of A can be obtained by restricting v to a function class.

Evaluating w requires us to evaluate $\text{atan} \frac{p_+(\mathbf{x})}{p_-(\mathbf{x})}$, F_+^* and F_-^* which are not readily available. However, Section 4.3 shows that the empirical estimator (8) is a consistent estimator of $\text{atan} \frac{p_+(\mathbf{x})}{p_-(\mathbf{x})}$ under mild conditions. Therefore, we propose the following two-step procedure to approximately lowerbound A :

Algorithm 1 Two-step Procedure for Approximately Lower Bounding A

1. Obtain $\hat{t}(\mathbf{x}) := \langle \hat{v}, \phi(\mathbf{x}) \rangle$ using (8). Approximate F_+^* and F_-^* using \hat{F}_+ and \hat{F}_- which are empirical CDFs of $\hat{t}(\mathbf{x}^+)$ and $\hat{t}(\mathbf{x}^-)$.
2. Optimize the empirical version of (12) by restricting v to a feasible function class and plugging in estimates obtained in the earlier step, i.e.,

$$\hat{v} := \underset{v \in [0, \pi/2], v \in \mathcal{F}}{\text{argmax}} \frac{1}{n_+} \sum_{i=1}^{n_+} \hat{w} [\hat{t}(\mathbf{x}_i^+)] \cdot \sin[v(\mathbf{x}_i^+)] + \frac{1}{n_-} \sum_{i=1}^{n_-} \hat{w} [\hat{t}(\mathbf{x}_i^-)] \cdot \cos[v(\mathbf{x}_i^-)], \quad (13)$$

$$\text{where } \hat{w}(\tau) := \sin(\tau + \frac{\pi}{4}) \left| \hat{F}_+(\tau) - \hat{F}_-(\tau) \right|.$$

Note that (13) is nothing but a weighted sample objective (7). Thus, it can be easily optimized by the algorithm that solves a weighted version of (7) given the approximated weights in the first step. In

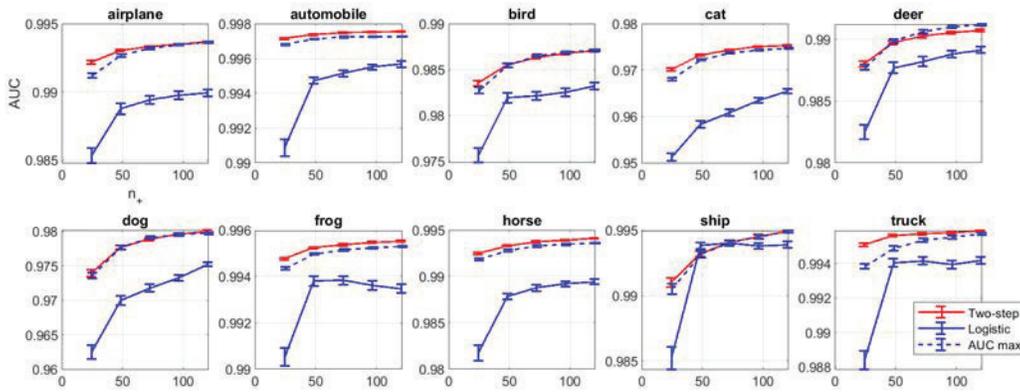


Figure 6: Testing AUC of one-versus-the-rest classification on CIFAR-10 dataset.

practice, we simply run the solver for (8) twice: The first time we run it without weights then run it again with weights $\hat{w}[\hat{t}(x_i)]$ calculated from the first run.

Since the above algorithm also approximates an optimal score ($\text{atan} \frac{p_+}{p_-}$) by maximizing an approximated lower bound of AUC^* , it is natural to wonder how the maximizer \hat{v} of (13) would perform in AUC maximization tasks. In the next section, we show that our two-step algorithm achieves a promising AUC performance compared to a state of the art AUC maximizer.

Computing AUC^* using (11) is different from using Wilcoxon-Mann-Whitney statistic (i.e., (13)): Our approach divides the space between ROC^* and the diagonal into small surface elements and then adds them up. Wilcoxon-Mann-Whitney statistic adds up all histogram bars, which are TPRs at different FPR levels. Our approach requires \bar{F}_+ and \bar{F}_- to be differentiable with respect to τ , which means the score distributions cannot be discrete. However, Wilcoxon-Mann-Whitney can compute the AUC of discrete score distributions without a problem. This difference is visualized in the middle and right plots of Figure 4.

6 Experiments

6.1 Numerical Comparison of Divergences and TV Bounds

In this experiment, we numerically compare the ROC divergence, the upper and lower bound in Proposition 2 with several other divergences and some known bounds of TV in Figure 5. In this numerical simulation, $p_+ = \mathcal{N}(0, 1)$ and $p_- = \mathcal{N}(\delta, 1)$. We plot ROC divergence, Jensen Shannon divergence, Wasserstein distance and TV between p_+ and p_- as δ grows from 0 to 5. We can see the (rescaled) ROC divergence closely resembles TV. When $\delta > 1.45$, the upper bound given in Proposition 2 is the tightest among known TV upper bounds [2, 4] (Pinsker's upper-bound, Bretagnolle & Huber's upper bound, Le Cam's upper bound). This suggests that combining our upper-bound with existing bounds may produce an even tighter bound for TV.

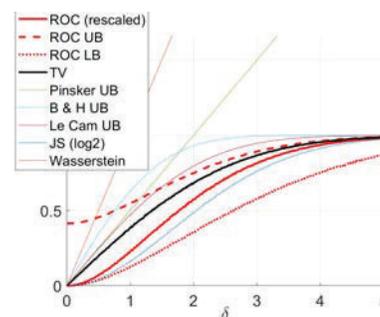


Figure 5: Comparison of various divergences and bounds of TV.

6.2 Imbalanced Classification on CIFAR-10

In this section, we test if the \hat{v} obtained in our two-step procedure (13) is indeed a good score function in terms of AUC in imbalanced classification tasks. We use a widely known image classification dataset CIFAR-10 [22]. The performance is compared with an AUC maximizer which maximizes the empirical lower bound in (10) and a vanilla logistic regression classifier. We set the surrogate loss $L(a, b) := -(1 - (a - b))^2$ in the AUC maximizer, as

suggested in [12]. All methods use linear models with no regularization terms since our models are simple and we have sufficient samples. Particularly, the \hat{t} and \hat{v} in our two-step algorithm is obtained using (8) by setting $\varphi(\mathbf{x}) = \mathbf{x}$ and $\lambda = 0$. The AUC maximization (AUC max) is implemented using SPAUC method [24].

Instead of using the raw features, we extract 50 dimensional bounded features by training a residual network [17] on the training dataset using the 10-class cross entropy loss. The structure of the network is included in the supplementary material. After obtaining features, we construct datasets for 10 different one-versus-the-rest classification tasks. For a single task, we pick a class and obtain X_+ by randomly sampling from this class n_+ times in the training set. Similarly, X_- is obtained by randomly sampling from the rest of the classes n_- times. In our experiments, we set $n_+ = 24, 48, 72, 96, 120$ and fix $n_- = 1000$ to create imbalanced positive and negative datasets. We run all three methods and obtain the corresponding score functions. For each class, we repeat the experiment 96 times using different random samples. We use the testing and training split provided by the dataset itself.

Our experiments can be seen as a transfer learning task which reuses predictive features trained for a multi-class classifier for one-versus-the-rest binary classification tasks.

The average AUCs computed on the testing dataset and their standard errors over 96 runs over different n_+ sample sizes are shown in Figure 6. Our method has approximately equal performance with the AUC maximizer despite not directly maximizing the AUC. This observation indicates that \hat{v} can be a good score function in AUC maximization tasks. Both of the methods significantly outperform vanilla logistic regression.

6.3 Discussions on Computational Complexity

Without loss of generality, assume $n_+ = n_- = n$. The naive calculation of objective (10) has a computational complexity $O(n^2)$ since we evaluate the loss function L at each pair of samples. However, authors in [21] have shown that the objective function in (10) can also be computed with $O(n \log(n))$ complexity for hinge loss (and decomposable loss functions). A recent work [39] simplifies the computation of (10) for the squared loss function L with an unbiased estimate. Suppose $t(\mathbf{x}) := \langle v, \mathbf{x} \rangle$, then the *negative* objective of (10) is an unbiased estimate of

$$1 + \text{Var}_{p_+}[\langle v, \mathbf{x} \rangle] + \text{Var}_{p_-}[\langle v, \mathbf{x} \rangle] + 2\langle v, \mathbb{E}_{p_-}[\mathbf{x}] - \mathbb{E}_{p_+}[\mathbf{x}] \rangle + \langle v, \mathbb{E}_{p_-}[\mathbf{x}] - \mathbb{E}_{p_+}[\mathbf{x}] \rangle^2. \quad (14)$$

After we approximating (14) with empirical terms, the computation can be done with a complexity $O(n)$. When implemented in an online fashion, it has a computational complexity of one datum.

In comparison, the objective (7) and (13) are summation of $\sin / \cos(v(\mathbf{x}))$ evaluated at each datum, so computing the objective/gradient has a computational complexity $O(n)$. Computing \hat{F}_+ and \hat{F}_- requires sorting our dataset, which has an average complexity $O(n \log n)$. However, once our datasets are sorted, $\hat{F}_+(\hat{t}(\mathbf{x}_0)) = \frac{i}{n_+}$, where i is the index of $\hat{t}(\mathbf{x}_0)$ in the sorted set $\{\hat{t}(\mathbf{x}_i)\}_{i=1}^{n_+}$.

7 Conclusions

In this paper, we show that a novel f -divergence arises from the arc length of the optimal ROC curve. The arc length can be accurately estimated from positive and negative samples using a variational expression. It is also an estimator for $\text{atan } p_+/p_-$ and has a convergence rate $O_p(n^{-\beta/4})$. Finally, we show that the area between the optimal ROC curve and the diagonal can be parameterized using a similar variational objective. It leads to a two-step procedure that approximately lower bounds the maximal AUC which achieves a promising result in AUC maximization tasks.

Acknowledgments and Disclosure of Funding

The author would like to thank Prof. Peter Flach and Dr. Hao Song for their helpful discussions. The author would like to thank four anonymous reviewers for their insightful comments. In particular, we thank Reviewer WSB and the Area Chair for pointing out the computational complexity inaccuracies in our initial version.

References

- [1] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.
- [2] C. Canonne. A short note on an inequality between kl and tv. *arXiv:2202.07198*, 2022.
- [3] C. Cortes and M. Mohri. Auc optimization vs. error rate minimization. In *Advances in Neural Information Processing Systems 16 (NeurIPS 2003)*, volume 16, 2003.
- [4] L. Devroye, A. Mehrabian, and T. Reddad. The total variation distance between high-dimensional gaussians. *arXiv:1810.08693*, 2018.
- [5] D. C. Edwards and C. E. Metz. A utility-based performance metric for roc analysis of n -class classification tasks. In *Medical Imaging 2007: Image Perception, Observer Performance, and Technology Assessment*, volume 6515, pages 21 – 30. SPIE, 2007.
- [6] D. C. Edwards and C. E. Metz. Optimality of a utility-based performance metric for roc analysis. In *Medical Imaging 2008: Image Perception, Observer Performance, and Technology Assessment*, volume 6917, pages 122 – 127. SPIE, 2008.
- [7] S. Eguchi and J. Copas. A class of logistic-type discriminant functions. *Biometrika*, 89(1):1–22, 2002.
- [8] T. Fawcett. An introduction to roc analysis. *Pattern Recognition Letters*, 27(8):861–874, 2006.
- [9] P. A. Flach. Roc analysis. In *Encyclopedia of machine learning and data mining*, pages 1–8. Springer, 2016.
- [10] K. Fukumizu. *Exponential manifold by reproducing kernel Hilbert spaces*, page 291–306. Cambridge University Press, 2009.
- [11] K. Fukumizu, L. Song, and A. Gretton. Kernel bayes’ rule: Bayesian inference with positive definite kernels. *The Journal of Machine Learning Research*, 14(1):3753–3783, 2013.
- [12] W. Gao, R. Jin, S. Zhu, and Z-H Zhou. One-pass auc optimization. In *Proceedings of the 30th International Conference on Machine Learning (ICML 2013)*, volume 28, pages 906–914, 2013.
- [13] D. J. Goodenough, K. Rossmann, and L. B Lusted. Radiographic applications of receiver operating characteristic (roc) curves. *Radiology*, 110(1):89–95, 1974.
- [14] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems 27 (NeurIPS 2014)*, pages 2672–2680, 2014.
- [15] D. M. Green and J. A. Swets. *Signal detection theory and psychophysics*, volume 1. Wiley New York, 1966.
- [16] J. A. Hanley and B. J. McNeil. The meaning and use of the area under a receiver operating characteristic (roc) curve. *Radiology*, 143(1):29–36, 1982.
- [17] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [18] J. Hermans, V. Begy, and G. Louppe. Likelihood-free MCMC with amortized approximate ratio estimators. In *Proceedings of the 37th International Conference on Machine Learning (ICML2020)*, volume 119, pages 4239–4248, 2020.
- [19] S. Hido, Y. Tsuboi, H. Kashima, M. Sugiyama, and T. Kanamori. Statistical outlier detection using direct density ratio estimation. *Knowledge and information systems*, 26(2):309–336, 2011.
- [20] J.B. Hiriart-Urruty and C. Lemaréchal. *Fundamentals of convex analysis*. Springer Science & Business Media, 2004.

- [21] T. Joachims. A support vector method for multivariate performance measures. In *Proceedings of the 22nd international conference on Machine learning (ICML)*, pages 377–384, 2005.
- [22] A. Krizhevsky. Learning multiple layers of features from tiny images. 2009. technical report.
- [23] S. Kullback and R. A. Leibler. On information and sufficiency. *Annals of Mathematical Statistics*, 22:79–86, 1951.
- [24] Y. Lei and Y. Ying. Stochastic proximal auc maximization. *Journal of Machine Learning Research*, 22(61):1–45, 2021.
- [25] Z. Lin, A. Khetan, G. Fanti, and S. Oh. Pacgan: The power of two samples in generative adversarial networks. In *Advances in Neural Information Processing Systems 31 (NeurIPS 2018)*, volume 31, 2018.
- [26] L. B. Lusted. Signal detectability and medical decision-making. *Science*, 171(3977):1217–1219, 1971.
- [27] J. Neyman and E. S. Pearson. On the problem of the most efficient tests of statistical hypotheses. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 231:289–337, 1933.
- [28] X. Nguyen, M. J. Wainwright, and M. I. Jordan. Estimating divergence functionals and the likelihood ratio by penalized convex risk minimization. In *Advances in Neural Information Processing Systems 20*, pages 1089–1096. 2008.
- [29] X. Nguyen, M. J. Wainwright, and M. I. Jordan. Estimating divergence functionals and the likelihood ratio by convex risk minimization. *IEEE Transactions on Information Theory*, 56(11):5847–5861, 2010.
- [30] S. Nowozin, B. Cseke, and R. Tomioka. f-gan: Training generative neural samplers using variational divergence minimization. In *Advances in Neural Information Processing Systems 29 (NeurIPS 2016)*, pages 271–279, 2016.
- [31] M. D. Reid and R. C. Williamson. Information, divergence and risk for binary experiments. *Journal of Machine Learning Research*, 12(22):731–817, 2011.
- [32] L. Rosasco, M. Belkin, and E. De Vito. On learning with integral operators. *Journal of Machine Learning Research*, 11(2), 2010.
- [33] B. Scholkopf and A. J. Smola. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2001.
- [34] B. K. Sriperumbudur, K. Fukumizu, A. Gretton, B. Schölkopf, and G. RG Lanckriet. On the empirical estimation of integral probability metrics. *Electronic Journal of Statistics*, 6: 1550–1599, 2012.
- [35] B. K. Sriperumbudur, K. Fukumizu, A. Gretton, A. Hyvärinen, and R. Kumar. Density estimation in infinite dimensional exponential families. *Journal of Machine Learning Research*, 18(57): 1–59, 2017.
- [36] I. Steinwart, D. R. Hush, and C. Scovel. Optimal rates for regularized least squares regression. In *Proceedings of Conference on Learning Theory (COLT) 2009*, pages 79–93, 2009.
- [37] X. Tong. A plug-in approach to neyman-pearson classification. *Journal of Machine Learning Research*, 14(56):3011–3040, 2013.
- [38] M. J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge University Press, 2019.
- [39] Y. Ying, L. Wen, and S. Lyu. Stochastic online auc maximization. In *Advances in Neural Information Processing Systems 29 (NeurIPS 2016)*, volume 29, 2016.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?
[Yes] In introduction, we discussed the background/usages ROC curves, f -divergence estimation and provided a full paragraph summary of the contribution of this paper. We give the same contribution rundown in the abstract as well.
 - (b) Did you describe the limitations of your work?
[Yes] In Section 5, we included a brief discussion comparing the proposed maximal AUC lower bounding approach vs. the classic Wilcoxon-Mann-Whitney statistic.
 - (c) Did you discuss any potential negative societal impacts of your work? [N/A]
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them?
[Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results?
[Yes] See Section 4.3
 - (b) Did you include complete proofs of all theoretical results?
[Yes] See Appendix
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)?
[Yes] Full instructions on how to reproduce our experiments is provided in the supplementary material.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)?
[Yes] See Section 6.2
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)?
[Yes] See Section 6.2
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)?
[N/A] The Computation time is not compared in this paper.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators?
[Yes] See Section 6.2
 - (b) Did you mention the license of the assets?
[No] The authors of CIFAR-10 do not specify a license for the dataset.
 - (c) Did you include any new assets either in the supplemental material or as a URL?
[No] We did not create any new asset in this research.
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating?
[No] We did not collect any data ourselves.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content?
[N/A] CIFAR-10 dataset we use in our experiment does not contain any identifiable human information.
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable?
[N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable?

[N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation?

[N/A]