
Estimating weighted areas under the ROC curve

Andreas Maurer

Istituto Italiano di Tecnologia
am@andreas-maurer.eu

Massimiliano Pontil

Istituto Italiano di Tecnologia & University College London
massimiliano.pontil@iit.it

Abstract

Exponential bounds on the estimation error are given for the plug-in estimator of weighted areas under the ROC curve. The bounds hold for single score functions and uniformly over classes of functions, whose complexity can be controlled by Gaussian or Rademacher averages. The results justify learning algorithms which select score functions to maximize the empirical partial area under the curve (pAUC). They also illustrate the use of some recent advances in the theory of nonlinear empirical processes.

1 Introduction

Using the area under the ROC curve (the AUC) to evaluate score functions has a long history in medical screening and bioinformatics ([9],[10]). In the last decades several algorithms have been developed to learn score functions which maximize the AUC ([7], [25]). In some cases, however, different regions of the false positive range may not be equally relevant to assess the quality of the score and should therefore be weighted differently, say with some nonconstant weight function W . In melanoma detection, for example, false negatives have disastrous consequences for a patient, while a certain number of false positives is tolerable. A good scoring function should then have very large values on the right hand side of the ROC plot. Such considerations have created interest in a *partial* area under the ROC curve (the pAUC), which only measures the area between two specified false positive rates, so that W is a step function ([27], [16], [19]). Several algorithms have been designed with the goal of maximizing the pAUC over classes of scoring functions ([21], [22], [20], [23]).

Any measurement or optimization of the AUC or pAUC must rely on a finite number of observations, which raises the questions of estimation and generalization. These problems are well understood for the AUC ([2], [6], [24]), where W is constant, but for more general weight functions estimation may be difficult to impossible. In Proposition 1 below we will show that the bias of its plug-in estimator may be bounded away from zero even in the limit of an infinite sample.

Our first contribution shows that these problems are absent if the weight function W is Lipschitz continuous. Theorem 2 shows that then the weighted AUC can be estimated at rate $n^{-1/2}$ by its plug-in estimator, where n is the sample size. Theorem 3 shows that the Lipschitz-weighted AUC can be uniformly estimated over a class of score functions at rate $n^{-1/2}$ if the Gaussian complexity of the class increases as $n^{1/2}$, which is standard for most function classes in machine learning. This result also implies a statistical performance guarantee for algorithms maximizing the empirical pAUC as in [22] or [20].

Even if W is Lipschitz the weighted AUC remains a nonlinear statistical functional of challenging complexity. Our second contribution is Theorem 7, which provides a general method to control the estimation errors of any statistical functional, which satisfies certain Lipschitz conditions with respect to the total-variation, Wasserstein or Kolmogorov metrics. The verification of these conditions for the weighted AUC then provides the proofs of Theorem 2 and Theorem 3.

1.1 Related work

The pAUC estimator is a special case of an L-estimator (see [26]). Asymptotic results for L-estimators have been known for a long time. Helmers [11] gave Berry-Esseen type rates of normal approximation for smoothed L-statistics. Finite sample bounds for special distributions are shown in [3]. The asymptotics of different estimators for the pAUC are considered in [8]. More recent asymptotic results for the pAUC are given in [28] and [1]. [14] announces a uniform bound of order $n^{-1/4}$ for left-partial areas. We do not know of any finite sample bounds comparable to Theorem 2 or uniform bounds of order $n^{-1/2}$ comparable to Theorem 3 below.

1.2 Notation

With $\mathbf{1}_S$ we denote the indicator function of a set S and with $H := \mathbf{1}_{(0,\infty)}$ the Heaviside function. $\mathcal{P}(\mathcal{X})$ is the set of non-negative measures and $\mathcal{P}_1(\mathcal{X}) \subset \mathcal{P}(\mathcal{X})$ is the set of probability measures on a measurable space \mathcal{X} respectively. For $\lambda \in [0, 1]$ and $\mu, \nu \in \mathcal{P}(\mathcal{X})$ we write $\lambda(\mu, \nu)$ for the convex combination $\lambda(\mu, \nu) = \lambda\mu + (1 - \lambda)\nu$. If $x \in \mathcal{X}$ then $\delta_x \in \mathcal{P}_1(\mathcal{X})$ is the unit mass at x . If $\mu \in \mathcal{P}_1(\mathcal{X})$ and $\mathbf{X} = (X_1, \dots, X_n) \sim \mu^n$ is an iid sample of size n the empirical measure is $\hat{\mu}(\mathbf{X}) = \frac{1}{n} \sum_{i=1}^n \delta_{X_i} \in \mathcal{P}_1(\mathcal{X})$. If \mathbf{X} is unambiguous we write $\hat{\mu} = \hat{\mu}(\mathbf{X})$.

$\mathcal{Y}^{\mathcal{X}}$ is the set of functions $f : \mathcal{X} \rightarrow \mathcal{Y}$, for $f \in \mathcal{Y}^{\mathcal{X}}$ and $\mu \in \mathcal{P}(\mathcal{X})$ the push-forward $f_{\#}\mu \in \mathcal{P}(\mathcal{Y})$ is defined $f_{\#}\mu(A) = \mu(f^{-1}(A))$, $A \subseteq \mathcal{Y}$. For $f \in \mathbb{R}^{\mathcal{X}}$ $\|f\|_{\infty} = \sup_{x \in \mathcal{X}} |f(x)|$ and if (\mathcal{X}, d) is a metric space then $\|f\|_{Lip} = \sup_{x \neq y \in \mathcal{X}} |f(x) - f(y)| / d(x, y)$.

For $\mu \in \mathcal{P}_1(\mathcal{X} \times \{0, 1\})$ (random labeled data) μ_0 and $\mu_1 \in \mathcal{P}(\mathcal{X})$ are defined $\mu_0(A) = \mu(A \times \{0\})$ and $\mu_1(A) = \mu(A \times \{1\})$ for $A \subseteq \mathcal{X}$, and $\mu(0) = \mu(\mathcal{X} \times \{0\})$ and $\mu(1) = \mu(\mathcal{X} \times \{1\})$ denote the relative frequencies of the two labels.

A summary of notation in tabular form is given in the supplement.

2 Weighted areas under the ROC-curve

Underlying the concept of the ROC-curve is the joint random occurrence of scores X in some open, bounded interval $\mathcal{I} \subseteq \mathbb{R}$ and binary labels $Y \in \{0, 1\}$. We assume $(X, Y) \sim \mu$ for some law $\mu \in \mathcal{P}_1(\mathcal{I} \times \{0, 1\})$, with corresponding un-normalized measures $\mu_0, \mu_1 \in \mathcal{P}(\mathcal{I})$ and scalar frequencies $\mu(0)$ and $\mu(1)$.

Any threshold $t \in \mathcal{I}$ on a score x induces a second labeling $H(x - t)$ alternative to y . The second labeling is considered correct iff it coincides with y . Correspondingly the true positive rate and false positive rate are the functions $tpr, fpr : \mathcal{I} \rightarrow [0, 1]$ defined respectively as

$$tpr(t) = \frac{\mu_1(t, \infty)}{\mu(1)} \text{ and } fpr(t) = \frac{\mu_0(t, \infty)}{\mu(0)}.$$

If we assume that μ_0 has a positive density, then fpr has a unique inverse, and the ROC-curve can be defined as the function $roc : [0, 1] \rightarrow [0, 1]$

$$roc(u) = tpr(fpr^{-1}(u)).$$

The ROC-curve gives the true positive rate corresponding to a specified false positive rate. Knowledge of the ROC-curve allows us to adjust the classification threshold according to respective cost estimates for false positives and false negatives This accounts for the great importance of ROC-curves in practice.

Let $W : [0, 1] \rightarrow \mathbb{R}_+$ be some specified weight function and consider the quantity

$$\int_{[0,1]} roc(u) W(1 - u) du.$$

With the change of variables $u = fpr(t)$ this is seen to be equal to

$$f_{W,\ell}(\mu) = \frac{1}{\mu(1)\mu(0)} \int_{\mathcal{I}^2} \ell(t' - t) W\left(\frac{\mu_0(-\infty, t]}{\mu(0)}\right) d\mu_1(t') d\mu_0(t), \quad (1)$$

as long as ℓ is the Heaviside function, $\ell = H$ (see supplement). Other choices of $\ell : \mathbb{R} \rightarrow [0, 1]$ replacing or approximating H will play a role when we want to learn scoring functions to optimize properties of the ROC-curve, but ℓ will generally be assumed nondecreasing. The equation (1) makes sense even when f_{pr} is not uniquely invertible, and it will serve as our definition of the statistical functional $f_{W,\ell} : \mathcal{P}_1(\mathcal{I} \times \{0, 1\}) \rightarrow \mathbb{R}_+$ in the sequel. Note that $f_{W,\ell}$ is monotonic in W and ℓ , and that its range lies in $[0, \|W\|_\infty]$.

An important special case is the "area under the curve" (AUC) obtained if $W \equiv 1$.

$$f_{1,H}(\mu) = \int_{[0,1]} \text{roc}(u) du = \frac{1}{\mu(1)\mu(0)} \int_{\mathcal{I}^2} H(t' - t) d\mu_1(t') d\mu_0(t).$$

$f_{1,H}(\mu)$ is the probability that $t' > t$ if t' and t are drawn independently from the positive and the negative conditional distributions. The AUC has been widely used to measure the quality of score functions and its statistical properties have been thoroughly analyzed from different perspectives ([7], [2], [6]).

If the respective costs of false negatives and false positives are very different, different regions of the false positive range are not equally relevant to measure the quality of the score. This leads to the consideration of a weighted or partial AUC where W is nonconstant and assigns different weights to different false positive rates. In the simplest case W it is a step function, such as $\mathbf{1}_{[a,b]}$ with $0 \leq a < b \leq 1$, but W may also be chosen to approximate the value of the ROC-curve itself at a given point.

2.1 Estimation

The measure μ , and with it the ROC-curve itself, are mathematical idealizations, accessible only through observations. We assume that we have access to an iid sample $(\mathbf{X}, \mathbf{Y}) = ((X_1, Y_1), \dots, (X_n, Y_n)) \sim \mu^n$ of labeled data. This defines the empirical measure $\hat{\mu} = \hat{\mu}(\mathbf{X}, \mathbf{Y}) \in \mathcal{P}_1(\mathcal{I} \times \{0, 1\})$ and corresponding empirical variants $\hat{\mu}_0, \hat{\mu}_1$ as well as the empirical rates $\hat{\mu}(0)$ and $\hat{\mu}(1)$. We will study the estimation of $f_{W,\ell}(\mu)$ by the plug-in estimator

$$\begin{aligned} f_{W,\ell}(\hat{\mu}) &= \frac{1}{\hat{\mu}(1)\hat{\mu}(0)} \int_{\mathcal{I}^2} \ell(t' - t) W\left(\frac{\hat{\mu}_0(-\infty, t]}{\hat{\mu}(0)}\right) d\hat{\mu}_1(t') d\hat{\mu}_0(t) \\ &= \frac{1}{|S_1||S_0|} \sum_{i \in S_1} \sum_{j \in S_0} \ell(X_i - X_j) W\left(\frac{|\{k \in S_0 : X_k \leq X_j\}|}{|S_0|}\right), \end{aligned}$$

where $S_1 = \{i : Y_i = 1\}$ and $S_0 = \{i : Y_i = 0\}$.

The dependence of the weight function on the order statistic of $\hat{\mu}_0/\hat{\mu}(0)$ causes problems which are absent in the estimation of the AUC as in [2]. If W has a discontinuity then the value of the functional depends discontinuously on the underlying law in the weak topology, and random fluctuations can make estimation impossible, even if the factor $1/(\hat{\mu}(1)\hat{\mu}(0))$ is bounded. The supplement gives proof of the following proposition, which shows that for discontinuous W the bias of the plug-in estimator may be bounded away from zero even in the limit of infinite sample sizes.

Proposition 1 *For $W = \mathbf{1}_{[1/2,1]}$ there exists a bounded, open interval \mathcal{I} and $\mu \in \mathcal{P}_1(\mathcal{I} \times \{0, 1\})$ such that for every even n and $\mathbf{X} \sim \mu^n$*

$$\lim_{n \rightarrow \infty} \mathbb{E}_{\mathbf{X} \sim \mu^n} [f_{W,H}(\hat{\mu}(\mathbf{X}))] = 1/4 < 1/2 = f_{W,H}(\mu).$$

To resolve this problem one might assume some regularity of μ , but since we wish to estimate a property of an *unknown* distribution from observations, we prefer to assume that the weight function W is Lipschitz.

We still need to exclude the case that one of the observed label frequencies $\hat{\mu}(0)$ and $\hat{\mu}(1)$ is too small relative to the sample size. For $\delta > 0$ we define the event

$$A_\delta = \left\{ \min\{\hat{\mu}(0), \hat{\mu}(1)\} > \sqrt{\frac{2 \ln(4/\delta)}{n}} \right\}. \quad (2)$$

We then have the following result.

Theorem 2 Suppose $W : \mathbb{R} \rightarrow [0, \infty)$ satisfies $\|W\|_\infty, \|W\|_{Lip} < \infty$, that $\ell : \mathcal{I} \rightarrow [0, 1]$ is nondecreasing and $\mu \in \mathcal{P}_1(\mathcal{I} \times \{0, 1\})$. Then for any $\delta > 0$ we have with probability at least $1 - \delta$ in $(\mathbf{X}, \mathbf{Y}) \sim \mu^n$ that A_δ implies

$$|f_{W,\ell}(\hat{\mu}(\mathbf{X}, \mathbf{Y})) - f_{W,\ell}(\mu)| \leq \frac{\|W\|_{Lip} + 9\|W\|_\infty}{\min\{\hat{\mu}(0), \hat{\mu}(1)\}^2} \sqrt{\frac{2 \ln(4/\delta)}{n}}.$$

Remarks:

1. For any event B the statement "with probability at least $1 - \delta$ it holds that A_δ implies B " means $\Pr\{A_\delta \cap B^c\} < \delta$, where B^c is the complement of B , so if we observe A_δ the bound holds with high probability.
2. Since $f_{W,H}$ is monotonic in W the bound implies one-sided bounds for discontinuous weight functions, such as the step functions defining the pAUC. If $\hat{W} \leq W_{Lip} \leq W$ where \hat{W} and W are discontinuous and W_{Lip} is Lipschitz, then with high probability $f_{W,H}(\mu) \geq f_{\hat{W},H}(\hat{\mu}) - O(1/\sqrt{n})$. Upper estimates and sandwich estimates may be constructed similarly.
3. As in the previous remark an approximation to the value of *roc* at some $u \in [0, 1]$ can be estimated with plug-in estimators $f_{\hat{W}_n,H}(\hat{\mu})$, if the weight functions \hat{W}_n are Lipschitz and, if regarded as probability densities, converge weakly to the unit mass δ_{1-u} .

2.2 Uniform bounds

Now let \mathcal{X} be some space of instances and let $\mu \in \mathcal{P}_1(\mathcal{X} \times \{0, 1\})$ be a law for labeled instances. Suppose that $\mathcal{H} \subseteq \mathcal{I}^{\mathcal{X}}$ is a class of candidate scoring functions $h : \mathcal{X} \rightarrow \mathcal{I}$. For every $h \in \mathcal{H}$ define $\bar{h} : \mathcal{X} \times \{0, 1\} \rightarrow \mathcal{I} \times \{0, 1\}$ by $\bar{h}(x, y) = (h(x), y)$. Given a sample $(\mathbf{X}, \mathbf{Y}) \sim \mu^n$, we would like to find $h \in \mathcal{H}$ so as to (approximately) maximize the value of $f_{W,H}(\bar{h}_{\#}\mu)$, where $\bar{h}_{\#}\mu$ is push-forward of μ under \bar{h} .

The strategy is to pick an h which (approximately) maximizes a regularized empirical surrogate $f_{W',\ell}(\bar{h}_{\#}\hat{\mu}(\mathbf{X}, \mathbf{Y}))$, where W' and ℓ are Lipschitz lower bounds of W and H respectively. The situation mirrors that of support vector machines, where one minimizes the hinge-loss as Lipschitz upper bound of the 0-1-loss.

Since the chosen score function h depends on (\mathbf{X}, \mathbf{Y}) , it becomes a random variable, and a justification of the method is given by a high probability bound on

$$\sup_{h \in \mathcal{H}} f_{W',\ell}(\bar{h}_{\#}\hat{\mu}) - f_{W,H}(\bar{h}_{\#}\mu) \leq \sup_{h \in \mathcal{H}} f_{W',\ell}(\bar{h}_{\#}\hat{\mu}) - f_{W',\ell}(\bar{h}_{\#}\mu),$$

where the inequality follows from $f_{W',\ell} \leq f_{W,H}$. We can bound the right hand side.

Theorem 3 Suppose $W : \mathbb{R} \rightarrow [0, \infty)$ satisfies $\|W\|_\infty, \|W\|_{Lip} < \infty$, that $\ell : \mathcal{I} \rightarrow [0, 1]$ is nondecreasing and $\|\ell\|_{Lip} < \infty$, that $\mathcal{H} \subseteq \mathcal{I}^{\mathcal{X}}$ and that $\mu \in \mathcal{P}_1(\mathcal{X} \times \{0, 1\})$. Let A_δ be as in (2). Then for any $\delta > 0$ we have with probability at least $1 - \delta$ in $(\mathbf{X}, \mathbf{Y}) \sim \mu^n$ that A_δ implies

$$\begin{aligned} & \sup_{h \in \mathcal{H}} |f_{W,\ell}(\bar{h}_{\#}\mu) - f_{W,\ell}(\bar{h}_{\#}\hat{\mu})| \\ & \leq \frac{8\sqrt{2\pi} \|\ell\|_{Lip} (\|W\|_\infty + \|W\|_{Lip})}{\hat{\mu}(0)^2 \hat{\mu}(1)} \frac{G(\mathcal{H})}{n} + \frac{\|W\|_{Lip} + 10\|W\|_\infty}{\min\{\hat{\mu}(0), \hat{\mu}(1)\}^2} \sqrt{\frac{4 \ln(16n/\delta)}{n}}. \end{aligned}$$

where $G(\mathcal{H})$ is the expected Gaussian complexity $G(\mathcal{H}) = \mathbb{E}_{\mathbf{X}} \mathbb{E}_\gamma \sup_{h \in \mathcal{H}} \sum_{i=1}^n \gamma_i h(X_i)$, with independent standard normal variables $\gamma_1, \dots, \gamma_n$.

Remarks.

1. Again we must observe A_δ for the bound to hold with high probability.
2. The use of Rademacher and Gaussian complexities has become a standard in learning theory, and there is a large body of literature giving bounds ([4], [15], [13], [12]), which can be substituted

above. These bounds apply to many different function classes, such as multi-layer neural networks or bounded sets of linear functionals in a reproducing-kernel-Hilbert space. Typically the bounds on the Gaussian complexity are of $O\left(\sqrt{n \ln(n)}\right)$ making the above bound $O\left(\sqrt{\ln(n)/n}\right)$.

3. Existing algorithms of [22] or [20] optimize the pAUC for discontinuous weight functions \hat{W} and practitioners may wish guarantees for discontinuous weight functions W . The next corollary, which directly follows from the monotonicity of $f_{W,\ell}$ in W and ℓ , shows that we can give such guarantees of $O\left(\sqrt{\ln(n)/n}\right)$, whenever a Lipschitz weight function W_{Lip} can be jammed between \hat{W} and W . This can be regarded as a statistical guarantee for the algorithms of [22] and [20].

Corollary 4 *Let $\hat{W}, W_{Lip}, W : [0, 1] \rightarrow [0, \infty)$, $\hat{W} \leq W_{Lip} \leq W$ and $\|W_{Lip}\|_{Lip} < \infty$. Then with probability at least $1 - \delta$ in $(\mathbf{X}, \mathbf{Y}) \sim \mu^n$ we have that A_δ implies*

$$\forall h \in \mathcal{H}, f_{W,H}(\bar{h}_{\#}\mu) \geq f_{\hat{W},\ell}(\bar{h}_{\#}\hat{\mu}) - B(n, W_{Lip}, \mathcal{H}, \hat{\mu}, \delta),$$

where $B(n, W_{Lip}, \mathcal{H}, \hat{\mu}, \delta)$ is the bound in Theorem 3.

3 Proofs

In this section we outline the proofs. Several technical details are given in the supplementary material.

3.1 Conditioning on the empirical label frequencies

The appearance of the true label frequencies $\mu(0)$ and $\mu(1)$ in the various denominators in the definition (1) of $f_{W,\ell}$ is a nuisance. But if the weight function is Lipschitz, we can approximate the estimation difference for $f_{W,\ell}$ by the estimation difference of another functional $g_{W,\ell,c}$, which is independent of the label frequencies and defined for $c > 0$ as

$$g_{W,\ell,c}(\mu) := \int_{\mathcal{I}} \int_{\mathcal{I}} \ell(t' - t) d\mu_1(t') W\left(\frac{\mu_0(-\infty, t]}{c}\right) d\mu_0(t).$$

In situations, where W, ℓ and c are unambiguously fixed, we will simply write g for $g_{W,\ell,c}$.

Lemma 5 *Let $\delta \in (0, 1)$ and A_δ as in (2). Then with probability at least $1 - \delta/2$ it holds that A_δ implies both $\mu(0) \geq \hat{\mu}(0)/2$ and*

$$|f_{W,\ell}(\hat{\mu}) - f_{W,\ell}(\mu)| \leq \frac{|g_{W,\ell,\mu(0)}(\hat{\mu}) - g_{W,\ell,\mu(0)}(\mu)|}{\hat{\mu}(0)\hat{\mu}(1)} + \frac{\|W\|_{Lip} + 8\|W\|_\infty}{\min\{\hat{\mu}(0), \hat{\mu}(1)\}^2} \sqrt{\frac{2 \ln(4/\delta)}{n}}.$$

The proof, a straightforward application of Hoeffdings inequality, is given in the supplement. If W is Lipschitz and $\hat{\mu}(0)$ and $\hat{\mu}(1)$ and n are reasonably large, the lemma allows us to bound the estimation error for $f_{W,\ell}$ in terms of the estimation error of the functional $g_{W,\ell,c}$, for which we have the following result.

Proposition 6 (i) *For $\delta > 0$ with probability at least $1 - \delta$*

$$|g_{W,\ell,c}(\mu) - g_{W,\ell,c}(\hat{\mu})| \leq \|W\|_\infty \sqrt{\frac{2 \ln(2/\delta)}{n}}. \quad (3)$$

(ii) *Under the conditions of Theorem 3 we have with probability at least $1 - \delta$ in $(\mathbf{X}, \mathbf{Y}) \sim \mu^n$ that*

$$\begin{aligned} & \sup_{h \in \mathcal{H}} |\mathbb{E} g_{W,\ell,c}(\bar{h}_{\#}\hat{\mu}) - g_{W,\ell,c}(\bar{h}_{\#}\hat{\mu}(\mathbf{X}, \mathbf{Y}))| \\ & \leq \frac{4\sqrt{2\pi} \|\ell\|_{Lip} \left(2\|W\|_\infty + c^{-1}\|W\|_{Lip}\right)}{n} G(\mathcal{H}) + 2\|W\|_\infty \sqrt{\frac{\ln(2/\delta)}{n}}. \end{aligned}$$

The proof of this proposition is not easy, but it easily implies Theorem 2 and Theorem 3.

Proof. (Theorem 2) A union bound of the inequalities in (i) and Lemma 5 together with some algebraic simplifications. ■

Proof. (Theorem 3) Set $\delta = n^{-1/2}$ in (3) and take the expectation inside the absolute value. Using $0 \leq g_{W,\ell,c} \leq \|W\|_\infty$ this gives for every μ (and every $\hat{h}_\# \mu$) the bias bound

$$|g_{W,\ell,c}(\mu) - \mathbb{E}g_{W,\ell,c}(\hat{\mu})| \leq \|W\|_\infty \left(\sqrt{\frac{2 \ln(2\sqrt{n})}{n}} + \sqrt{\frac{1}{n}} \right) \leq 2 \|W\|_\infty \sqrt{\frac{\ln(4n)}{n}}.$$

Combine this with (ii), set $c = \mu(0)$ and simplify to obtain

$$\begin{aligned} & \sup_{h \in \mathcal{H}} |g_{W,\ell,\mu(0)}(\hat{h}_\# \mu) - g_{W,\ell,\mu(0)}(\hat{h}_\# \hat{\mu}(\mathbf{X}, \mathbf{Y}))| \\ & \leq \frac{4\sqrt{2\pi} \|\ell\|_{Lip} \left(2 \|W\|_\infty + \mu(0)^{-1} \|W\|_{Lip} \right)}{n} G(\mathcal{H}) + 2 \|W\|_\infty \sqrt{\frac{2 \ln(8n/\delta)}{n}}. \end{aligned}$$

Now assume A_δ and combine with Lemma 5 in a union bound, using also $\mu(0) \geq \hat{\mu}(0)/2$, and simplify to obtain the conclusion. ■

3.2 Plug-in estimators for Lipschitz functionals

To establish Proposition 6 we will give a general method to prove high probability bounds on the estimation and uniform estimation error for Lipschitz functionals on $\mathcal{P}_1(\mathcal{U})$ with $\mathcal{U} \subseteq \mathbb{R}$.

For any $\mu, \nu \in \mathcal{P}(\mathbb{R})$ we define the metrics

$$\begin{aligned} d_{TV}(\mu, \nu) &= |\mu - \nu|(\mathbb{R}) \\ d_1(\mu, \nu) &= \int_{\mathbb{R}} |\mu((-\infty, t]) - \nu((-\infty, t])| dt \\ d_\infty(\mu, \nu) &= \sup_{t \in \mathbb{R}} |\mu((-\infty, t]) - \nu((-\infty, t])|. \end{aligned}$$

If $\mu, \nu \in \mathcal{P}_1(\mathbb{R})$ then d_{TV} is the total variation metric, d_1 is the 1-Wasserstein metric and d_∞ is the Kolmogorov metric. All these metrics have the convexity property

$$d(\lambda(\mu, \nu), \lambda(\mu', \nu')) \leq \lambda d(\mu, \mu') + (1 - \lambda) d(\nu, \nu') \quad (4)$$

for any $\lambda \in [0, 1]$ and $\mu, \mu', \nu, \nu' \in \mathcal{P}_1(\mathbb{R})$.

Theorem 7 Let $\mathcal{U} \subseteq \mathbb{R}$ and $f : \mathcal{P}_1(\mathcal{U}) \rightarrow \mathbb{R}$ and let $\delta > 0$. Consider the following conditions on f

(a) $\forall \mu, \nu \in \mathcal{P}_1(\mathcal{U}), f(\mu) - f(\nu) \leq L_\infty d_\infty(\mu, \nu)$.

(b) $\forall \mu, \nu \in \mathcal{P}_1(\mathcal{U}), f(\mu) - f(\nu) \leq L_1 d_1(\mu, \nu)$.

(c) $\forall \lambda \in [0, 1]$ and $\forall \mu, \mu', \nu, \nu' \in \mathcal{P}_1(\mathcal{U})$ we have

$$f(\lambda(\mu, \nu)) - f(\lambda(\mu, \nu')) - f(\lambda(\mu', \nu)) + f(\lambda(\mu', \nu')) \leq (1 - \lambda) \lambda L_2 d_1(\nu, \nu') d_{TV}(\mu, \mu').$$

Then

(i) Suppose f satisfies (a). Let $\mu \in \mathcal{P}_1(\mathcal{U})$ and $\hat{\mu} = \hat{\mu}(\mathbf{X})$ with $\mathbf{X} \sim \mu^n$. Then with probability at least $1 - \delta$

$$|f(\mu) - f(\hat{\mu})| \leq L_\infty \sqrt{\frac{\ln(2/\delta)}{2n}}.$$

(ii) Suppose f satisfies (a),(b) and (c). Let $\mathcal{H} \subseteq \mathcal{U}^\mathcal{X}$, $\mu \in \mathcal{P}_1(\mathcal{X})$ and $\hat{\mu} = \hat{\mu}(\mathbf{X})$ with $\mathbf{X} \sim \mu^n$. Then with probability at least $1 - \delta$

$$\sup_{h \in \mathcal{H}} f(h_\# \hat{\mu}) - \mathbb{E}f(h_\# \hat{\mu}) \leq \frac{\sqrt{8\pi}(L_1 + L_2)}{n} G(\mathcal{H}) + L_\infty \sqrt{\frac{\ln(1/\delta)}{n}}.$$

The proof uses two nontrivial auxiliary results. For (i) we need the following version of the Dvoretzky-Kiefer-Wolfowitz theorem as sharpened by Massart [17].

Theorem 8 *If μ is a probability measure on the real line and $\hat{\mu}$ is the empirical measure for $\mathbf{X} \sim \mu^n$ then for $t > 0$*

$$\Pr \{d_\infty(\mu, \hat{\mu}) > t\} \leq 2e^{-2nt^2}.$$

For part (ii) we use a result about nonlinear empirical processes [18], for which we need some additional notation. For $\mathbf{x} \in \mathcal{X}^n$ and $h \in \mathcal{U}^{\mathcal{X}}$ we write $h(\mathbf{x}) = (h(x_1), \dots, h(x_n)) \in \mathcal{U}^n$. For $k \in \{1, \dots, n\}$ and $y, y' \in \mathcal{U}$ we define the partial difference operator for functions $f : \mathcal{U}^n \rightarrow \mathbb{R}$ by

$$D_{y,y'}^k f(\mathbf{x}) = f(x_1, \dots, x_{k-1}, y, x_k, \dots, x_n) - f(x_1, \dots, x_{k-1}, y', x_k, \dots, x_n).$$

Theorem 9 *(see [18]) Let $\mathbf{X} = (X_1, \dots, X_n)$ be a vector of independent random variables with values \mathcal{X} and \mathbf{X}' iid to \mathbf{X} . Let $\mathcal{U} \subseteq \mathbb{R}$, $\mathcal{H} \subseteq \mathcal{U}^{\mathcal{X}}$ and $f : \mathcal{U}^n \rightarrow \mathbb{R}$. Then for any $\delta \in (0, 1)$, with probability at least $1 - \delta$,*

$$\sup_{h \in \mathcal{H}} f(h(\mathbf{X})) - \mathbb{E}[f(h(\mathbf{X}'))] \leq \sqrt{2\pi} (2M_L(f) + J_L(f)) G(\mathcal{H}) + M(f) \sqrt{n \ln(1/\delta)},$$

where the three seminorms M , J_L and M_L are defined as

$$\begin{aligned} M(f) &= \max_k \sup_{\mathbf{x} \in \mathcal{U}^n, y, y' \in \mathcal{U}} D_{y,y'}^k f(\mathbf{x}) \\ M_L(f) &= \max_k \sup_{\mathbf{x} \in \mathcal{U}^n, y, y' \in \mathcal{U}, y \neq y'} \frac{D_{y,y'}^k f(\mathbf{x})}{|y - y'|} \\ J_L &= n \max_{k,l:k \neq l} \sup_{\mathbf{x} \in \mathcal{U}^n, z, z', y, y' \in \mathcal{U}, y \neq y'} \frac{D_{y,y'}^k D_{z,z'}^l f(\mathbf{x})}{|y - y'|}. \end{aligned}$$

Proof. (of Theorem 7) (i) If (a) holds then by Theorem 8

$$\Pr \{|g(\mu) - g(\hat{\mu})| > t\} \leq \Pr \left\{ d_\infty(\mu, \hat{\mu}) > \frac{t}{L_\infty} \right\} \leq 2 \exp\left(-\frac{2nt^2}{L_\infty^2}\right).$$

Set the right hand side to δ and solve t for the conclusion.

(ii) Suppose (a), (b) and (c) hold. We will use Theorem 9 and bound the seminorms for $f(\hat{\mu})$. Fix $\mathbf{x} \in \mathcal{U}^n$ and $k \in \{1, \dots, n\}$ and set $\hat{\mu}_k = \frac{1}{n-1} \sum_{i:i \neq k} \delta_{x_i}$.

$$\begin{aligned} D_{y,y'}^k f(\hat{\mu}) &= f\left(\frac{n-1}{n} \hat{\mu}_k + \frac{1}{n} \delta_y\right) - f\left(\frac{n-1}{n} \hat{\mu}_k + \frac{1}{n} \delta_{y'}\right) \\ &\leq L_\infty d_\infty\left(\frac{n-1}{n} \hat{\mu}_k + \frac{1}{n} \delta_y, \frac{n-1}{n} \hat{\mu}_k + \frac{1}{n} \delta_{y'}\right) \leq \frac{L_\infty}{n} d_\infty(\delta_y, \delta_{y'}) \leq \frac{L_\infty}{n}, \end{aligned}$$

where the first inequality follows from (a), the second from convexity d_∞ (see (4)) and the last from $d_\infty(\delta_y, \delta_{y'}) \leq 1$. In the same way (b) gives $D_{y,y'}^k f(\hat{\mu}) \leq L_1 |y - y'|/n$, since $d_1(\delta_y, \delta_{y'}) = |y - y'|$. This gives $M(f) \leq L_\infty/n$ and $M_L(f) \leq L_1/n$.

To bound $J_L(f)$ let $l \neq k$ and write $\hat{\mu}_{kl} = \frac{1}{n-2} \sum_{i:i \notin \{k,l\}} \delta_{x_i}$, $\hat{\mu}_{k,z} = \frac{n-2}{n-1} \hat{\mu}_{kl} + \frac{1}{n-1} \delta_z$ and $\hat{\mu}_{k,z'} = \frac{n-2}{n-1} \hat{\mu}_{kl} + \frac{1}{n-1} \delta_{z'}$. Then

$$\begin{aligned} &D_{y,y'}^k D_{z,z'}^l f(\mathbf{x}) \\ &= f\left(\frac{n-1}{n} \hat{\mu}_{k,z} + \frac{1}{n} \delta_y\right) - f\left(\frac{n-1}{n} \hat{\mu}_{k,z} + \frac{1}{n} \delta_{y'}\right) \\ &\quad - f\left(\frac{n-1}{n} \hat{\mu}_{k,z'} + \frac{1}{n} \delta_y\right) + f\left(\frac{n-1}{n} \hat{\mu}_{k,z'} + \frac{1}{n} \delta_{y'}\right) \\ &\leq \frac{n-1}{n^2} L_{12} d_{TV}(\hat{\mu}_{k,z}, \hat{\mu}_{k,z'}) d_1(y, y') \leq \frac{L_2}{n^2} d_{TV}(\delta_z, \delta_{z'}) d_1(y, y') \leq \frac{2L_2}{n^2} |y - y'|. \end{aligned}$$

The first inequality follows from (c), the second from convexity of d_{TV} (see (4)) and the last from $d_{TV}(\delta_z, \delta_{z'}) \leq 2$. This gives $J_L(f) \leq 2L_2/n$. Substitution in the conclusion of Theorem 9 completes the proof. ■

3.3 Proof of Proposition 6

To apply Theorem 7 we need to control the Lipschitz properties of $g_{W,\ell,c}$. We do so at first with respect to the unnormalized measures μ_0 and μ_1 .

Proposition 10 Fix W, ℓ and c . The functional $g = g_{W,\ell,c}$ satisfies $\forall \mu, \mu', \nu, \nu' \in \mathcal{P}_1(\mathcal{I} \times \{0, 1\})$

(a) $g(\nu) - g(\nu') \leq \|W\|_\infty (d_\infty(\nu_0, \nu'_0) + d_\infty(\nu_1, \nu'_1))$

(b) if ℓ is Lipschitz then $g(\nu) - g(\nu') \leq \|W\|_\infty \|\ell\|_{Lip} (d_1(\nu_0, \nu'_0) + d_1(\nu_1, \nu'_1))$

(c) if W and ℓ are Lipschitz then

$$\begin{aligned} & g(\lambda(\mu, \nu)) - g(\lambda(\mu, \nu')) - g(\lambda(\mu', \nu)) + g(\lambda(\mu', \nu')) \\ & \leq (1 - \lambda) \lambda \|\ell\|_{Lip} \left(\frac{1}{c} \|W\|_{Lip} + \|W\|_\infty \right) (d_1(\nu_1, \nu'_1) + d_1(\nu_0, \nu'_0)) d_{TV}(\mu, \mu'). \end{aligned}$$

The proof is given in the supplement. We then eliminate the labels and replace the disconnected space $\mathcal{I} \times \{0, 1\}$ by a subset of the real line. Let $a = \sup \mathcal{I}$ and $b = \sup \mathcal{I} - \inf \mathcal{I}$. We map $\mathcal{I} \times \{0\}$ to $(-b, 0)$ in ascending order of \mathcal{I} , and we map $\mathcal{I} \times \{1\}$ to $(0, b)$ in descending order of \mathcal{I} .

More formally we define the bijection $\tau : (x, y) \in \mathcal{I} \times \{0, 1\} \mapsto (2y - 1)(a - x) \in (-b, 0) \cup (0, b)$. The push-forward $\mu \mapsto \tau_{\#}\mu$ is then a bijection between $\mathcal{P}_1(\mathcal{I} \times \{0, 1\})$ and $\mathcal{P}_1((-b, 0) \cup (0, b))$ with inverse $\tau_{\#}^{-1}$. Then for $A \subset \mathcal{I}$ we have $\mu_0(A) = \mu(A \times \{0\}) = \tau_{\#}\mu(A - a)$ and $\mu_1(A) = \tau_{\#}\mu(a - A)$. The supplement gives proof of the following lemma.

Lemma 11 For $\mu, \nu \in \mathcal{P}_1(\mathcal{I} \times \{0, 1\})$ we have

(i) $d_{TV}(\mu_0, \nu_0) + d_{TV}(\mu_1, \nu_1) = d_{TV}(\mu, \nu) = d_{TV}(\tau_{\#}\mu, \tau_{\#}\nu)$

(ii) $d_\infty(\mu_0, \nu_0) + d_\infty(\mu_1, \nu_1) \leq 2d_\infty(\tau_{\#}\mu, \tau_{\#}\nu)$

(iii) $d_1(\mu_0, \nu_0) + d_1(\mu_1, \nu_1) \leq 2d_1(\tau_{\#}\mu, \tau_{\#}\nu)$

Since $g_{W,\ell,c}(\mu) = g_{W,\ell,c} \circ \tau_{\#}^{-1}(\tau_{\#}\mu)$, we can shift perspective to the functional $g_{W,\ell,c} \circ \tau_{\#}^{-1}$ on $\mathcal{P}_1((-b, 0) \cup (0, b))$. Proposition 10 and Lemma 11 show that the functional $\mu \mapsto g_{W,\ell,c} \circ \tau_{\#}^{-1}(\mu)$ satisfies the Lipschitz condition (a) of Theorem 7 with $L_\infty = 2\|W\|_\infty d_\infty(\mu, \nu)$ and, if ℓ is Lipschitz, also (b) with $L_1 = \|W\|_\infty \|\ell\|_{Lip}$ and (c) with

$$L_2 = 2\|\ell\|_{Lip} \left(\frac{1}{c} \|W\|_{Lip} + \|W\|_\infty \right).$$

In Theorem 7 we then replace \mathcal{U} by $(-b, 0) \cup (0, b)$, f by $g_{W,\ell,c} \circ \tau_{\#}^{-1}$, \mathcal{X} by $(\mathcal{X} \times \{0, 1\})$, \mathcal{H} by the class of functions $\mathcal{H}' = \{(x, y) \mapsto \tau(h(x), y) : h \in \mathcal{H}\}$, and we substitute the constants L_∞, L_1 and L_2 by the values given above. Theorem 7 then gives us parts (i) and (ii) of Proposition 6, since by symmetry of the standard normal distribution

$$G(\mathcal{H}') = \mathbb{E}_{\mathbf{X}} \mathbb{E}_\gamma \sup_{h \in \mathcal{H}} \sum_{i=1}^n \gamma_i (2Y_i - 1)(a - h(X_i)) = G(\mathcal{H}).$$

4 Conclusion

In a nonparametric setting the estimation of partial areas under the ROC-curve by plug-in estimators was shown to be impossible. It is nevertheless possible to give error bounds for Lipschitz weight functions and uniform bounds which provide some justification for existing algorithms optimizing the pAUC. The method to control the uniform estimation errors for nonlinear statistical functionals in terms of Lipschitz conditions appears promising in a more general context.

Acknowledgments and Disclosure of Funding

This work was partially supported by a grant from SAP SE.

Broader Impact

A solid mathematical basis is beneficial to the development of practical statistical methods. We believe that the present work improves the understanding of ROC-curves and the optimization of score functions used in machine learning and medical diagnostics.

References

- [1] Gianfranco Adimari and Monica Chiogna. Jackknife empirical likelihood based confidence intervals for partial areas under roc curves. *Statistica Sinica*, pages 1457–1477, 2012.
- [2] Shivani Agarwal, Thore Graepel, Ralf Herbrich, Sarel Har-Peled, and Dan Roth. Generalization bounds for the area under the roc curve. *Journal of Machine Learning Research*, 6(Apr):393–425, 2005.
- [3] EA Baklanov and IS Borisov. Probability inequalities and limit theorems for generalized l-statistics. *Lithuanian Mathematical Journal*, 43(2):125–140, 2003.
- [4] P. Bartlett and S. Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3:463–482, 2002.
- [5] Stéphane Boucheron, Gábor Lugosi, Pascal Massart, et al. Concentration inequalities using the entropy method. *The Annals of Probability*, 31(3):1583–1614, 2003.
- [6] Stéphan Cléménçon, Gábor Lugosi, Nicolas Vayatis, et al. Ranking and empirical minimization of u-statistics. *The Annals of Statistics*, 36(2):844–874, 2008.
- [7] Corinna Cortes and Mehryar Mohri. Auc optimization vs. error rate minimization. In *Advances in neural information processing systems*, pages 313–320, 2004.
- [8] Lori E Dodd and Margaret S Pepe. Partial auc estimation and regression. *Biometrics*, 59(3):614–623, 2003.
- [9] David J Goodenough, Kurt Rossmann, and Lee B Lusted. Radiographic applications of receiver operating characteristic (roc) curves. *Radiology*, 110(1):89–95, 1974.
- [10] James A Hanley and Barbara J McNeil. The meaning and use of the area under a receiver operating characteristic (roc) curve. *Radiology*, 143(1):29–36, 1982.
- [11] Roelof Helmers et al. The order of the normal approximation for linear combinations of order statistics with smooth weight functions. *The Annals of Probability*, 5(6):940–953, 1977.
- [12] S. Kakade, S. Shalev-Shwartz, and A. Tewari. Regularization techniques for learning with matrices. *Journal of Machine Learning Research*, 13:1865–1890, 2012.
- [13] Sham M Kakade, Karthik Sridharan, and Ambuj Tewari. On the complexity of linear prediction: Risk bounds, margin bounds, and regularization. In *Advances in neural information processing systems*, pages 793–800, 2009.
- [14] 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.
- [15] V. Koltchinskii and D. Panchenko. Rademacher processes and bounding the risk of function learning. In J. Wellner E. Gine, D. Mason, editor, *High Dimensional Probability II*, pages 443–459. 2000.
- [16] Zhenqiu Liu and Terry Hyslop. Partial auc for differentiated gene detection. In *2010 IEEE International Conference on BioInformatics and BioEngineering*, pages 310–311. IEEE, 2010.

- [17] Pascal Massart. The tight constant in the dvoretzky-kiefer-wolfowitz inequality. *The annals of Probability*, pages 1269–1283, 1990.
- [18] Andreas Maurer and Massimiliano Pontil. Uniform concentration and symmetrization for weak interactions. In *Proceedings of the Thirty-Second Conference on Learning Theory*, pages 2372–2387, 2019.
- [19] Donna Katzman McClish. Analyzing a portion of the roc curve. *Medical Decision Making*, 9(3):190–195, 1989.
- [20] H. Narasimhan and S. Agarwal. Support vector algorithms for optimizing the partial area under the roc curve. *Neural computation*, 29(7):1919–1963, 2017.
- [21] Harikrishna Narasimhan and Shivani Agarwal. A structural svm based approach for optimizing partial auc. In *International Conference on Machine Learning*, pages 516–524, 2013.
- [22] Harikrishna Narasimhan and Shivani Agarwal. Svm pauc tight: a new support vector method for optimizing partial auc based on a tight convex upper bound. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 167–175. ACM, 2013.
- [23] Sakrapee Paisitkriangkrai, Chunhua Shen, and Anton Van Den Hengel. Efficient pedestrian detection by directly optimizing the partial area under the roc curve. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1057–1064, 2013.
- [24] Gengsheng Qin and Xiao-Hua Zhou. Empirical likelihood inference for the area under the roc curve. *Biometrics*, 62(2):613–622, 2006.
- [25] Alain Rakotomamonjy. Optimizing area under roc curve with svms. In *ROCAI*, pages 71–80, 2004.
- [26] Robert J Serfling. *Approximation theorems of mathematical statistics*, volume 162. John Wiley & Sons, 2009.
- [27] Stephen D Walter. The partial area under the summary roc curve. *Statistics in medicine*, 24(13):2025–2040, 2005.
- [28] Hanfang Yang, Kun Lu, and Yichuan Zhao. A nonparametric approach for partial areas under roc curves and ordinal dominance curves. *Statistica Sinica*, pages 357–371, 2017.

Appendix

For the convenience of the reader we restate results which appear in the main part of the paper.

A Proof of the formula (1) for the pAUC

Lemma 12 *We have*

$$\int_{[0,1]} \text{roc}(u) W(1-u) du = f_{W,H}(\mu).$$

Proof. With $u = f_{pr}(t) = \frac{\mu_0(t, \infty)}{\mu(0)}$ we get

$$\begin{aligned} & \int_{[0,1]} \text{roc}(u) W(1-u) du \\ &= \int_{[0,1]} t_{pr}(f_{pr}^{-1}(u)) W(1-u) du \\ &= \frac{1}{\mu(0)} \int_{\mathbb{R}} t_{pr}(t) W\left(\frac{\mu_0(-\infty, t]}{\mu(0)}\right) d\mu_0(t) \\ &= \frac{1}{\mu(0)\mu(1)} \int_{\mathbb{R}} \int_{\mathbb{R}} H(t'-t) d\mu_1(t') W\left(\frac{\mu((-\infty, t], 0)}{\mu(0)}\right) d\mu_0(t) \\ &= f_{W,H}(\mu), \end{aligned}$$

since

$$t_{pr}(t) = \frac{\mu_1(t, \infty)}{\mu(1)} = \frac{1}{\mu(1)} \int_{\mathbb{R}} H(t'-t) d\mu_1(t').$$

■

B Proof of the lower bound on the bias, Proposition 1

Proposition 1 *For $W = 1_{[1/2,1]}$ there exists a bounded, open interval I and $\mu \in P_1(\mathcal{I} \times \{0,1\})$ such that for every even n and $X \sim \mu^n$*

$$\lim_{n \rightarrow \infty} \mathbb{E}_{\mathbf{X} \sim \mu^n} [f_{W,H}(\hat{\mu}(\mathbf{X}))] = 1/4 < 1/2 = f_{W,H}(\mu).$$

Proof. (of Proposition 1) Let \mathcal{I} be any bounded open interval containing $[0, 1]$ and consider members of $\mathcal{P}_1(\mathcal{I} \times \{0, 1\})$ supported on the set $\{(0, 0), (1/2, 1), (1, 0)\}$. Every such probability measure is of the form

$$\mu_{s,r} := s(r\delta_{(0,0)} + (1-r)\delta_{(1,0)}) + (1-s)\delta_{(1/2,1)}$$

with $r, s \in (0, 1)$. Substitution in the definition (1) gives $f_{1_{[1/2,1]}, H}(\mu_{s,r}) = r\mathbf{1}_{[1/2,1]}(r)$.

Now let the underlying law μ be $\mu = \mu_{1/2, 1/2}$. Then $f_{1_{[1/2,1]}, H}(\mu) = 1/2$ but the empirical measure is $\hat{\mu} = \mu_{\hat{S}, \hat{R}}$, where \hat{S} and \hat{R} can be generated by trials of a fair coin and have binomial distributions concentrated around $1/2$. As the sample size n goes to infinity $\Pr\{\hat{R} < 1/2\} \rightarrow 1/2$ and for every $\epsilon > 0$ we, $\Pr\left\{\left|\hat{R} - 1/2\right| > \epsilon\right\} \rightarrow 0$. It follows that

$$\lim_{n \rightarrow \infty} \mathbb{E}\left[f_{W,H}(\mu_{\hat{S}, \hat{R}})\right] = \lim_{n \rightarrow \infty} \mathbb{E}\left[\hat{R} \mathbf{1}_{[1/2,1]}(\hat{R})\right] = 1/4 < 1/2 = f_{1_{[1/2,1]}, H}(\mu).$$

■

C Proof of Lemma 5

Lemma 5 *Let $\delta \in (0, 1)$ and A_δ as in (2). Then with probability at least $1 - \delta/2$ it holds that A_δ implies both $\mu(0) \geq \hat{\mu}(0)/2$ and*

$$|f_{W,\ell}(\hat{\mu}) - f_{W,\ell}(\mu)| \leq \frac{|g_{W,\ell,\mu(0)}(\hat{\mu}) - g_{W,\ell,\mu(0)}(\mu)|}{\hat{\mu}(0)\hat{\mu}(1)} + \frac{\|W\|_{Lip} + 8\|W\|_\infty}{\min\{\hat{\mu}(0), \hat{\mu}(1)\}^2} \sqrt{\frac{2\ln(4/\delta)}{n}}.$$

Proof. We have $f_{W,\ell}(\mu) = g_{W,\ell,\mu(0)}(\mu) / (\mu(0)\mu(1))$. Thus

$$\begin{aligned} & |f_{W,\ell}(\hat{\mu}) - f_{W,\ell}(\mu)| \\ &= \left| \left(\frac{g_{W,\ell,\mu(0)}(\hat{\mu})}{\hat{\mu}(0)\hat{\mu}(1)} - \frac{g_{W,\ell,\mu(0)}(\mu)}{\hat{\mu}(0)\hat{\mu}(1)} \right) + \left(f_{W,\ell}(\hat{\mu}) - \frac{g_{W,\ell,\mu(0)}(\hat{\mu})}{\hat{\mu}(0)\hat{\mu}(1)} \right) \right. \\ &\quad \left. + \left(\frac{g_{W,\ell,\mu(0)}(\mu)}{\hat{\mu}(0)\hat{\mu}(1)} - \frac{g_{W,\ell,\mu(0)}(\mu)}{\mu(0)\mu(1)} \right) \right| \\ &\leq \frac{|g_{W,\ell,\mu(0)}(\hat{\mu}) - g_{W,\ell,\mu(0)}(\mu)|}{\hat{\mu}(0)\hat{\mu}(1)} + \|W\|_{Lip} \left| \frac{1}{\hat{\mu}(0)} - \frac{1}{\mu(0)} \right| \\ &\quad + \|W\|_\infty \left| \frac{1}{\hat{\mu}(0)\hat{\mu}(1)} - \frac{1}{\mu(0)\mu(1)} \right| \\ &\leq \frac{|g_{W,\ell,\mu(0)}(\hat{\mu}) - g_{W,\ell,\mu(0)}(\mu)|}{\hat{\mu}(0)\hat{\mu}(1)} + \frac{\|W\|_{Lip}}{\hat{\mu}(0)\mu(0)} |\hat{\mu}(0) - \mu(0)| \\ &\quad + \left(\frac{\|W\|_\infty}{\hat{\mu}(0)\hat{\mu}(1)\mu(0)\mu(1)} |\hat{\mu}(0)\hat{\mu}(1) - \mu(0)\mu(1)| \right) \end{aligned}$$

Hoeffdings inequality [5] implies with probability at least $1 - \delta/2$ that

$$|\hat{\mu}(1) - \mu(1)| = |\hat{\mu}(0) - \mu(0)| \leq \sqrt{\frac{\ln(4/\delta)}{2n}},$$

so under A_δ we have $\mu(0) \geq \hat{\mu}(0)/2$ and $\mu(1) \geq \hat{\mu}(1)/2$, and substitution above gives

$$\begin{aligned} & |f_{W,\ell}(\hat{\mu}) - f_{W,\ell}(\mu)| \\ &\leq \frac{|g_{W,\ell,\mu(0)}(\hat{\mu}) - g_{W,\ell,\mu(0)}(\mu)|}{\hat{\mu}(0)\hat{\mu}(1)} \left(\frac{\|W\|_{Lip}}{\hat{\mu}(0)^2} + \frac{8\|W\|_\infty}{\min\{\hat{\mu}(0), \hat{\mu}(1)\}^2} \right) \sqrt{\frac{2\ln(4/\delta)}{n}} \end{aligned}$$

■

D Proof of the Lipschitz bound, Proposition 10

For the proof of this proposition we need some preliminary definitions and lemmata. We define

$$K(t) := \int_0^t W\left(\frac{u}{c}\right) du, F_\mu(t) := \mu_0(-\infty, t] \text{ and } \phi_\mu(t) := \int_{\mathbb{R}} \ell(t-t') d\mu_1(t'),$$

so $g = g_{W,\ell,c}$ can be written as a Riemann-Stieltjes Integral

$$g_{W,\ell,c}(\mu) = \int_{\mathbb{R}} \phi_\mu(t) dK(F_\mu(t)). \quad (5)$$

Lemma 6 *Let $\mu, \nu, \rho \in \mathcal{P}_1(\mathcal{I} \times \{0, 1\})$ and ℓ and W be bounded. Then $|\phi_\rho(t)(K(F_\mu(t)) - K(F_\nu(t)))| \rightarrow 0$ as $t \rightarrow \pm\infty$. Furthermore*

$$\int_{\mathbb{R}} \phi_\rho(t) (dK(F_\mu(t)) - dK(F_\nu(t))) = - \int_{\mathbb{R}} \phi'_\rho(t) (K(F_\mu(t)) - K(F_\nu(t))). \quad (6)$$

Proof.

$$\begin{aligned} |\phi_\rho(t)(K(F_\mu(t)) - K(F_\nu(t)))| &= \left| \phi_\rho(t) \int_{F_\nu(t)}^{F_\mu(t)} W\left(\frac{u}{c}\right) du \right| \\ &\leq \|\ell\|_\infty \|W\|_\infty |F_\mu(t) - F_\nu(t)| \rightarrow 0 \end{aligned}$$

as $t \rightarrow \pm\infty$. The identity follows from integration by parts. ■

Lemma 7 For $\mu, \mu' \in \mathcal{P}_1(\mathcal{I} \times \{0, 1\})$

(i) $\|\phi_\mu - \phi_{\mu'}\|_\infty \leq d_\infty(\mu_1, \mu'_1)$

If ℓ is Lipschitz, then all of the following hold

(ii) $\|\phi_\mu - \phi_{\mu'}\|_\infty \leq \|\ell\|_{Lip} d_1(\mu_1, \mu'_1)$

(iii) $\|\phi'_\mu - \phi'_{\mu'}\|_\infty \leq \|\ell\|_{Lip} d_{TV}(\mu_1, \mu'_1)$

(iv) $\|\phi'_\mu - \phi'_{\mu'}\|_1 \leq d_{TV}(\mu_1, \mu'_1)$

(v) $\|\phi'_\mu\|_\infty \leq \|\ell\|_{Lip}$

(vi) $\|\phi'_\mu\|_1 \leq 1$

Proof. (i) Recall that ℓ is assumed to be non-decreasing with $\ell_n : \mathbb{R} \rightarrow [0, 1]$. Approximate ℓ pointwise from below by non-decreasing, differentiable functions $\ell_n : \mathbb{R} \rightarrow [0, 1]$, $n \in \mathbb{N}$. By dominated convergence and integration by parts

$$\begin{aligned} |\phi_\mu(t) - \phi_{\mu'}(t)| &= \lim_{n \rightarrow \infty} \left| \int_{\mathbb{R}} \ell_n(t' - t) d(\mu_1 - \mu'_1)(t') \right| \\ &= \lim_{n \rightarrow \infty} \left| \int_{\mathbb{R}} \ell'_n(t' - t) (F(\mu_1, t') - F(\mu'_1, t')) dt' \right| \\ &\leq d_\infty(\mu, \mu') \lim_{n \rightarrow \infty} \int_{\mathbb{R}} |\ell'_n(t' - t)| dt' = d_\infty(\mu, \mu') \lim_{n \rightarrow \infty} \int_{\mathbb{R}} \ell'_n(t' - t) dt' \\ &\leq d_\infty(\mu, \mu'). \end{aligned}$$

by the assumed properties of ℓ .

If ℓ is Lipschitz, then ℓ is absolutely continuous and its derivative ℓ' exists almost everywhere and $\|\ell\|_{Lip} = \|\ell'\|_\infty$. (ii)

$$\begin{aligned} |\phi_\mu(t) - \phi_{\mu'}(t)| &\leq \left| \int_{\mathbb{R}} \ell(t' - t) d(\mu_1 - \mu'_1)(t') \right| \\ &= \left| \int_{\mathbb{R}} \ell'(t' - t) (F(\mu_1, t') - F(\mu'_1, t')) dt' \right| \\ &\leq \|\ell'\|_\infty d_1(\mu_1, \mu'_1) \end{aligned}$$

(iii)

$$|\phi'_\mu(t) - \phi'_{\mu'}(t)| = \left| \int_{\mathbb{R}} \ell'(t' - t) d(\mu_1 - \mu'_1)(t') \right| \leq \|\ell'\|_\infty d_{TV}(\mu_1, \mu'_1)$$

(iv)

$$\begin{aligned} \int |\phi'_\mu(t) - \phi'_{\mu'}(t)| dt &= \int_{\mathbb{R}} \left| \int_{\mathbb{R}} \ell'(t' - t) d(\mu_1 - \mu'_1)(t') \right| dt \\ &\leq \int_{\mathbb{R}} \int_{\mathbb{R}} |\ell'(t' - t)| d(|\mu_1 - \mu'_1|)(t') dt \\ &= \int_{\mathbb{R}} \left(\int_{\mathbb{R}} \ell'(t - t') dt \right) d(|\mu_1 - \mu'_1|)(t') \\ &= \int_{\mathbb{R}} d(|\mu_1 - \mu'_1|)(t') = d_{TV}(\mu_1, \mu'_1). \end{aligned}$$

The proofs of (v) and (vi) are similar to those of (iii) and (iv). ■

Now we can prove Proposition 10.

Proposition 10 Fix W, ℓ and c . The functional $g = g_{W, \ell, c}$ satisfies $\forall \mu, \mu', \nu, \nu' \in P_1(\mathcal{I} \times \{0, 1\})$

(a) $g(\nu) - g(\nu') \leq \|W\|_\infty (d_\infty(\nu_0, \nu'_0) + d_\infty(\nu_1, \nu'_1))$

(b) if ℓ is Lipschitz then

$$g(\nu) - g(\nu') \leq \|W\|_\infty \|\ell\|_{Lip} (d_1(\nu_0, \nu'_0) + d_1(\nu_1, \nu'_1))$$

(c) if W and ℓ are Lipschitz then

$$\begin{aligned} & g(\lambda(\mu, \nu)) - g(\lambda(\mu, \nu')) - g(\lambda(\mu', \nu)) + g(\lambda(\mu', \nu')) \\ & \leq (1 - \lambda) \lambda \|\ell\|_{Lip} \left(\frac{1}{c} \|W\|_{Lip} + \|W\|_\infty \right) (d_1(\nu_1, \nu'_1) + d_1(\nu_0, \nu'_0)) d_{TV}(\mu, \mu'). \end{aligned}$$

Proof. From (5) we have

$$g(\mu) - g(\nu) = \int_{\mathbb{R}} (\phi_\mu(t) - \phi_\nu(t)) dK(F_\mu(t)) + \int_{\mathbb{R}} \phi_{\nu, n}(t) (dK(F_\mu(t)) - dK(F_\nu(t))).$$

Proof of (a). Approximate ℓ as in the proof of Lemma 7 (i) and let

$$\phi_{\nu, n}(t) = \int_{\mathbb{R}} \ell_n(t - t') d\mu_1(t').$$

Using dominated convergence, the integration by parts formula (6) and Lemma 7 (i)

$$\begin{aligned} & g(\mu) - g(\nu) \\ & = \int_{\mathbb{R}} (\phi_\mu(t) - \phi_\nu(t)) W \left(\frac{F_\mu(t)}{c} \right) d\mu_0(t) - \lim_{n \rightarrow \infty} \int_{\mathbb{R}} \phi'_{\nu, n}(t) (K(F_\mu(t)) - K(F_\nu(t))) dt \\ & \leq \|W\|_\infty \left(\|\phi_\mu(\cdot) - \phi_\nu(\cdot)\|_\infty + \lim_{n \rightarrow \infty} \int_{\mathbb{R}} \phi'_{\nu, n}(t) |F_\mu(t) - F_\nu(t)| dt \right) \\ & \leq \|W\|_\infty \left(d_\infty(\mu_1, \nu_1) + d_\infty(\mu_0, \nu_0) \lim_{n \rightarrow \infty} \int_{\mathbb{R}} \phi'_{\nu, n}(t) dt \right) \\ & \leq \|W\|_\infty (d_\infty(\mu_1, \nu_1) + d_\infty(\mu_0, \nu_0)). \end{aligned}$$

Proof of (b). Now ℓ is Lipschitz and we can write similar to the above

$$\begin{aligned} g(\mu) - g(\nu) & \leq \|W\|_\infty \left(\|\phi_\mu(\cdot) - \phi_\nu(\cdot)\|_\infty + \int_{\mathbb{R}} \phi'_\nu(t) |F_\mu(t) - F_\nu(t)| dt \right) \\ & \leq \|W\|_\infty \|\ell\|_{Lip} d_1(\mu_1, \nu_1) + d_1(\mu_0, \nu_0), \end{aligned}$$

where we used Lemma 7 (ii) and (v).

Proof of (c). We write the second difference as

$$\begin{aligned} & g(\lambda(\mu, \nu)) - g(\lambda(\mu, \nu')) - g(\lambda(\mu', \nu)) + g(\lambda(\mu', \nu')) \\ & = \int_{\mathbb{R}} (\phi_{\lambda(\mu, \nu)}(t) - \phi_{\lambda(\mu, \nu')}(t) - \phi_{\lambda(\mu', \nu)}(t) + \phi_{\lambda(\mu', \nu')}(t)) dK(F_{\lambda(\mu, \nu)}(t)) \\ & + \int_{\mathbb{R}} (\phi_{\lambda(\mu', \nu)}(t) - \phi_{\lambda(\mu', \nu')}(t)) (dK(F_{\lambda(\mu, \nu)}(t)) - dK(F_{\lambda(\mu', \nu)}(t))) \\ & + \int_{\mathbb{R}} (\phi_{\lambda(\mu, \nu)}(t) - \phi_{\lambda(\mu', \nu)}(t)) (dK(F_{\lambda(\mu, \nu)}(t)) - dK(F_{\lambda(\mu, \nu')}(t))) \\ & + \int_{\mathbb{R}} \phi_{\lambda(\mu', \nu')}(t) (dK(F_{\lambda(\mu, \nu)}(t)) - dK(F_{\lambda(\mu, \nu')}(t)) - dK(F_{\lambda(\mu', \nu)}(t)) + dK(F_{\lambda(\mu', \nu')}(t))) \\ & = A + B + C + D. \end{aligned}$$

We bound the four terms in turn. The term A simply vanishes, because ϕ_μ is affine in μ . To bound B first note that for any t

$$\begin{aligned} K(F_{\lambda(\mu',\nu)}(t)) - K(F_{\lambda(\mu,\nu)}(t)) &= \int_{F_{\lambda(\mu,\nu)}(t)}^{F_{\lambda(\mu',\nu)}(t)} W\left(\frac{u}{c}\right) du = \int_{(1-\lambda)F_\mu(t)}^{(1-\lambda)F_{\mu'}(t)} W\left(\frac{u}{c} - \lambda F_\nu(t)\right) du \\ &\leq (1-\lambda) \|W\|_\infty |F_{\mu'}(t) - F_\mu(t)|. \end{aligned}$$

Using integration by parts and Lemma 7 (iii)

$$\begin{aligned} B &= \int_{\mathbb{R}} \left(\phi'_{\lambda(\mu',\nu)}(t) - \phi'_{\lambda(\mu,\nu)}(t) \right) \left(K(F_{\lambda(\mu',\nu)}(t)) - K(F_{\lambda(\mu,\nu)}(t)) \right) dt \\ &\leq (1-\lambda) \|W\|_\infty \int_{\mathbb{R}} \left| \phi'_{\lambda(\mu',\nu)}(t) - \phi'_{\lambda(\mu,\nu)}(t) \right| |F_{\mu'}(t) - F_\mu(t)| dt \\ &\leq (1-\lambda) \lambda \|\ell'\|_\infty \|W\|_\infty d_1(\mu_0, \mu'_0) d_{TV}(\nu_1, \nu'_1) \end{aligned}$$

The term C is bounded using Lemma 7 (ii) as

$$\begin{aligned} C &= \int_{\mathbb{R}} \left(\phi_{\lambda(\mu,\nu')}(t) - \phi_{\lambda(\mu',\nu')}(t) \right) \left(dK(F_{\lambda(\mu,\nu)}(t)) - dK(F_{\lambda(\mu',\nu')}(t)) \right) \\ &\leq \left\| \phi_{\lambda(\mu,\nu')}(\cdot) - \phi_{\lambda(\mu',\nu')}(\cdot) \right\|_\infty \left(\int_{\mathbb{R}} \left| W\left(\frac{F_{\lambda(\mu,\nu)}(t)}{c}\right) - W\left(\frac{F_{\lambda(\mu',\nu')}(t)}{c}\right) \right| d\mu_0(t) \right. \\ &\quad \left. + \int_{\mathbb{R}} W\left(\frac{F_{\lambda(\mu,\nu')}(t)}{c}\right) |d\nu_0(t) - d\nu'_0(t)| \right) \\ &\leq (1-\lambda) \lambda \|\ell'\|_\infty \left(\left\| \frac{W}{c} \right\|_{Lip} + \|W\|_\infty \right) d_1(\mu_1, \mu'_1) d_{TV}(\nu_0, \nu'_0). \end{aligned}$$

Finally we again use integration by parts to bound D .

$$\begin{aligned} D &= \int_{\mathbb{R}} \phi'_{\lambda(\mu',\nu')}(t) \left(K(F_{\lambda(\mu,\nu)}(t)) - K(F_{\lambda(\mu,\nu')}(t)) - K(F_{\lambda(\mu',\nu)}(t)) + K(F_{\lambda(\mu',\nu')}(t)) \right) dt \\ &= \int_{\mathbb{R}} \phi'_{\lambda(\mu',\nu')}(t) \left(\int_{F_{\lambda(\mu,\nu')}(t)}^{F_{\lambda(\mu,\nu)}(t)} W\left(\frac{u}{c}\right) du - \int_{F_{\lambda(\mu',\nu)}(t)}^{F_{\lambda(\mu',\nu')}(t)} W\left(\frac{u}{c}\right) du \right) dt \\ &= \int_{\mathbb{R}} \phi'_{\lambda(\mu',\nu')}(t) \left(\int_{F_{\lambda\nu'}(t)}^{F_{\lambda\nu}(t)} \left(W\left(\frac{u - (1-\lambda)F_\mu(t)}{c}\right) - W\left(\frac{u - (1-\lambda)F_{\mu'}(t)}{c}\right) \right) du \right) dt \\ &\leq (1-\lambda) \lambda \left\| \frac{W}{c} \right\|_{Lip} \int_{\mathbb{R}} \left| \phi'_{\lambda(\mu',\nu')}(t) \right| |F_\nu(t) - F_{\nu'}(t)| |F_\mu(t) - F_{\mu'}(t)| dt \\ &\leq (1-\lambda) \lambda \|\ell'\|_\infty \left\| \frac{W}{c} \right\|_{Lip} \int_{\mathbb{R}} |F_\nu(t) - F_{\nu'}(t)| |F_\mu(t) - F_{\mu'}(t)| dt \\ &\leq (1-\lambda) \lambda \|\ell'\|_\infty \left\| \frac{W}{c} \right\|_{Lip} d_1(\mu_0, \mu'_0) d_{TV}(\nu_0, \nu'_0), \end{aligned}$$

where we used Lemma 7 (v) in the second inequality. Adding the bounds and using $d_{TV}(\nu_i, \nu'_i) \leq d_{TV}(\nu, \nu')$, we get

$$A + B + C + D \leq (1-\lambda) \lambda \|\ell'\|_\infty \left(\left\| \frac{W}{c} \right\|_{Lip} + \|W\|_\infty \right) (d_1(\mu_0, \mu'_0) + d_1(\mu_1, \mu'_1)) d_{TV}(\nu, \nu')$$

■

E Proof of Lemma 11

Lemma 11] For $\mu, \nu \in P_1(\mathcal{I} \times \{0, 1\})$ we have

- (i) $d_{TV}(\mu_0, \nu_0) + d_{TV}(\mu_1, \nu_1) = d_{TV}(\mu, \nu) = d_{TV}(\tau_{\#}\mu, \tau_{\#}\nu)$
- (ii) $d_\infty(\mu_0, \nu_0) + d_\infty(\mu_1, \nu_1) \leq 2d_\infty(\tau_{\#}\mu, \tau_{\#}\nu)$
- (iii) $d_1(\mu_0, \nu_0) + d_1(\mu_1, \nu_1) \leq 2d_1(\tau_{\#}\mu, \tau_{\#}\nu)$

Proof. (i) is obvious from the definitions of μ_i, ν_i and $\tau_{\#}\mu, \tau_{\#}\nu$.

(ii)

$$\begin{aligned}
& d_{\infty}(\mu_0, \nu_0) + d_{\infty}(\mu_1, \nu_1) \\
&= \sup_{t \in \mathcal{I}} |\mu_0(-\infty, t] - \nu_0(-\infty, t]| + \sup_{t \in \mathcal{I}} |\mu_1(-\infty, t] - \nu_1(-\infty, t]| \\
&= \sup_{t \in \mathcal{I}} |\tau_{\#}\mu(-\infty, t-a] - \tau_{\#}\nu(-\infty, t-a]| + \sup_{t \in \mathcal{I}} |\tau_{\#}\mu[a-t, \infty) - \tau_{\#}\nu[a-t, \infty)| \\
&\leq \sup_{t \in \mathbb{R}} |\tau_{\#}\mu(-\infty, t] - \tau_{\#}\nu(-\infty, t]| + \sup_{t \in \mathbb{R}} |\tau_{\#}\mu[t, \infty) - \tau_{\#}\nu[t, \infty)| \\
&\leq 2d_{\infty}(\tau_{\#}\mu, \tau_{\#}\nu).
\end{aligned}$$

(iii)

$$\begin{aligned}
& 2d_1(\tau_{\#}\mu, \tau_{\#}\nu) \\
&= \int_{-\infty}^{\infty} |\tau_{\#}\mu(-\infty, t] - \tau_{\#}\nu(-\infty, t]| dt + \int_{-\infty}^{\infty} |\tau_{\#}\mu[t, \infty) - \tau_{\#}\nu[t, \infty)| dt \\
&\geq \int_{-\infty}^0 |\tau_{\#}\mu(-\infty, t] - \tau_{\#}\nu(-\infty, t]| dt + \int_0^{\infty} |\tau_{\#}\mu[t, \infty) - \tau_{\#}\nu[t, \infty)| dt \\
&= \int_{\mathcal{I}} |\tau_{\#}\mu(-\infty, t-a] - \tau_{\#}\nu(-\infty, t-a]| dt + \int_{\mathcal{I}} |\tau_{\#}\mu[a-t, \infty) - \tau_{\#}\nu[a-t, \infty)| dt \\
&= \int_{\mathcal{I}} |\mu_0(-\infty, t] - \nu_0(-\infty, t]| dt + \int_{\mathcal{I}} |\mu_1(-\infty, t] - \nu_1(-\infty, t]| dt \\
&= d_1(\mu_0, \nu_0) + d_1(\mu_1, \nu_1)
\end{aligned}$$

■

F Table of notation

Notation	Definition	Section
A_δ	Conditioning event	2.1
$d_{TV}(\mu, \nu)$	Total variation distance for $\mu, \nu \in \mathcal{P}_1(\mathcal{X})$	3.2
$d_1(\mu, \nu)$	1-Wasserstein distance for $\mu, \nu \in \mathcal{P}_1(\mathbb{R})$	3.2
$d_\infty(\mu, \nu)$	Kolmogorov distance for $\mu, \nu \in \mathcal{P}_1(\mathbb{R})$	3.2
$D_{y,y'}^k$	Partial difference operator	3.2
δ_x	Unit mass at x	1.2
fpr	False positive rate	2
$f_{\#}\mu$	Pushforward of measure μ under f	1.2
$f_{W,\ell}$	Weighted area und ROC curve	2
$g_{W,\ell,c}$	Functional independent of label frequancies	3.1
$G(\mathcal{H})$	Gaussian complexity	2.2
H	Heaviside function $1_{(0,\infty)}$	1.2
\mathcal{H}	Class of scoring functions	2.2
\bar{h}	$\bar{h}(x, y) = (h(x), y)$ for $h \in \bar{h}$	2.2
ℓ	Loss function	2
μ_0	$\mu_0(A) = \mu(A \times \{0\})$ for $\mu \in \mathcal{P}_1(\mathcal{X} \times \{0, 1\})$	1.2
μ_1	$\mu_1(A) = \mu(A \times \{1\})$ for $\mu \in \mathcal{P}_1(\mathcal{X} \times \{0, 1\})$	1.2
$\hat{\mu}(\mathbf{X})$	Empirical measure for sample \mathbf{X}	1.2
$\mathcal{P}(\mathcal{X})$	Nonnegative measures on \mathcal{X}	1.2
$\mathcal{P}_1(\mathcal{X})$	Probability measures on \mathcal{X}	1.2
roc	ROC-curve	2
tpr	True positive rate	2
τ	Bijection $I \times \{0, 1\} \rightarrow (-b, 0) \cup (0, b)$	3.3
W	Weight function	2
$\ \cdot\ _\infty$	Supremum norm	1.2
$\ \cdot\ _{Lip}$	Lipschitz seminorm	1.2