
On the Statistical Consistency of Plug-in Classifiers for Non-decomposable Performance Measures

Harikrishna Narasimhan[†], Rohit Vaish[†], Shivani Agarwal

Department of Computer Science and Automation

Indian Institute of Science, Bangalore – 560012, India

{harikrishna, rohit.vaish, shivani}@csa.iisc.ernet.in

Abstract

We study consistency properties of algorithms for non-decomposable performance measures that cannot be expressed as a sum of losses on individual data points, such as the F-measure used in text retrieval and several other performance measures used in class imbalanced settings. While there has been much work on designing algorithms for such performance measures, there is limited understanding of the theoretical properties of these algorithms. Recently, Ye et al. (2012) showed consistency results for two algorithms that optimize the F-measure, but their results apply only to an idealized setting, where precise knowledge of the underlying probability distribution (in the form of the ‘true’ posterior class probability) is available to a learning algorithm. In this work, we consider plug-in algorithms that learn a classifier by applying an empirically determined threshold to a suitable ‘estimate’ of the class probability, and provide a general methodology to show consistency of these methods for any non-decomposable measure that can be expressed as a continuous function of true positive rate (TPR) and true negative rate (TNR), and for which the Bayes optimal classifier is the class probability function thresholded suitably. We use this template to derive consistency results for plug-in algorithms for the F-measure and for the geometric mean of TPR and precision; to our knowledge, these are the first such results for these measures. In addition, for continuous distributions, we show consistency of plug-in algorithms for any performance measure that is a continuous and monotonically increasing function of TPR and TNR. Experimental results confirm our theoretical findings.

1 Introduction

In many real-world applications, the performance measure used to evaluate a learning model is non-decomposable and cannot be expressed as a summation or expectation of losses on individual data points; this includes, for example, the F-measure used in information retrieval [1], and several combinations of the true positive rate (TPR) and true negative rate (TNR) used in class imbalanced classification settings [2–5] (see Table 1). While there has been much work in the last two decades in designing learning algorithms for such performance measures [6–14], our understanding of the statistical consistency of these methods is rather limited. Recently, Ye et al. (2012) showed consistency results for two algorithms for the F-measure [15] that use the ‘true’ posterior class probability to make predictions on instances. These results implicitly assume that the given learning algorithm has precise knowledge of the underlying probability distribution (in the form of the true posterior class probability); this assumption does not however hold in most real-world settings.

In this paper, we consider a family of methods that construct a plug-in classifier by applying an empirically determined threshold to a suitable ‘estimate’ of the class probability (obtained using a model learned from a sample drawn from the underlying distribution). We provide a general method-

[†]Both authors contributed equally to this paper.

Table 1: Performance measures considered in our study. Here $\beta \in (0, \infty)$ and $p = \mathbf{P}(y = 1)$. Each performance measure here can be expressed as $\mathcal{P}_D^\Psi[h] = \Psi(\text{TPR}_D[h], \text{TNR}_D[h], p)$. The last column contains the assumption on the distribution D under which the plug-in algorithm considered in this work is statistically consistent w.r.t. a performance measure (details in Sections 3 and 5).

Measure	Definition	Ref.	$\Psi(u, v, p)$	Assumption on D
AM (1-BER)	$(\text{TPR} + \text{TNR})/2$	[17–19]	$\frac{u+v}{2}$	Assumption A
F $_\beta$ -measure	$(1 + \beta^2)/(\frac{\beta^2}{\text{Prec}} + \frac{1}{\text{TPR}})$	[1, 19]	$\frac{(1+\beta^2)pu}{p+\beta^2(pu+(1-p)(1-v))}$	Assumption A
G-TP/PR	$\sqrt{\text{TPR} \cdot \text{Prec}}$	[3]	$\sqrt{\frac{pu^2}{pu+(1-p)(1-v)}}$	Assumption A
G-Mean (GM)	$\sqrt{\text{TPR} \cdot \text{TNR}}$	[2, 3]	\sqrt{uv}	Assumption B
H-Mean (HM)	$2/(\frac{1}{\text{TPR}} + \frac{1}{\text{TNR}})$	[4]	$\frac{2uv}{u+v}$	Assumption B
Q-Mean (QM)	$1 - ((1 - \text{TPR})^2 + (1 - \text{TNR})^2)/2$	[5]	$1 - \frac{(1-u)^2 + (1-v)^2}{2}$	Assumption B

ology to show statistical consistency of these methods (under a mild assumption on the underlying distribution) for any performance measure that can be expressed as a continuous function of the TPR and TNR and the class proportion, and for which the Bayes optimal classifier is the class probability function thresholded at a suitable point. We use our proof template to derive consistency results for the F-measure (using a recent result by [15] on the Bayes optimal classifier for F-measure), and the geometric mean of TPR and precision; to our knowledge, these are the first such results for these performance measures. Using our template, we also obtain a recent consistency result by Menon et al. [16] for the arithmetic mean of TPR and TNR. In addition, we show that for continuous distributions, the optimal classifier for any performance measure that is a continuous and monotonically increasing function of TPR and TNR is necessarily of the requisite thresholded form, thus establishing consistency of the plug-in algorithms for all such performance measures. Experiments on real and synthetic data confirm our theoretical findings, and show that the plug-in methods considered here are competitive with the state-of-the-art SVM^{perf} method [12] for non-decomposable measures.

Related Work. Much of the work on non-decomposable performance measures in binary classification settings has focused on the F-measure; these include the empirical plug-in algorithm considered here [6], cost-weighted versions of SVM [9], methods that optimize convex and non-convex approximations to F-measure [10–14], and decision-theoretic methods that learn a class probability estimate and compute predictions that maximize the expected F-measure on a test set [7–9]. While there has been considerable amount of work on consistency of algorithms for univariate performance measures [16, 20–22], theoretical results on non-decomposable measures have been limited to characterizing the Bayes optimal classifier for F-measure [15, 23, 24], and some consistency results for F-measure for certain idealized versions of the empirical plug-in and decision theoretic methods that have access to the true class probability [15]. There has also been some work on algorithms that optimize F-measure in multi-label classification settings [25, 26] and consistency results for these methods [26, 27], but these results do not apply to the binary classification setting that we consider here; in particular, in a binary classification setting, the F-measure that one seeks to optimize is a single number computed over the entire training set, while in a multi-label setting, the goal is to optimize the mean F-measure computed over multiple labels on individual instances.

Organization. We start with some preliminaries in Section 2. Section 3 presents our main result on consistency of plug-in algorithms for non-decomposable performance measures that are functions of TPR and TNR. Section 4 contains application of our proof template to the AM, F $_\beta$ and G-TP/PR measures, and Section 5 contains results under continuous distributions for performance measures that are monotonic in TPR and TNR. Section 6 describes our experimental results on real and synthetic data sets. Proofs not provided in the main text can be found in the Appendix.

2 Preliminaries

Problem Setup. Let \mathcal{X} be any instance space. Given a training sample $S = ((x_1, y_1), \dots, (x_n, y_n)) \in (\mathcal{X} \times \{\pm 1\})^n$, our goal is to learn a binary classifier $\hat{h}_S : \mathcal{X} \rightarrow \{\pm 1\}$ to make predictions for new instances drawn from \mathcal{X} . Assume all examples (both training and test) are drawn iid according to some unknown probability distribution D on $\mathcal{X} \times \{\pm 1\}$. Let $\eta(x) = \mathbf{P}(y = 1|x)$ and $p = \mathbf{P}(y = 1)$ (both under D). We will be interested in settings where the performance of \hat{h}_S is measured via a *non-decomposable performance measure* $\mathcal{P} : \{\pm 1\}^{\mathcal{X}} \rightarrow \mathbb{R}_+$, which cannot be expressed as a sum or expectation of losses on individual examples.

Non-decomposable performance measures. Let us first define the following quantities associated with a binary classifier $h : \mathcal{X} \rightarrow \{\pm 1\}$:

$$\begin{aligned} \text{True Positive Rate / Recall} & \quad \text{TPR}_D[h] = \mathbf{P}(h(x) = 1 \mid y = 1) \\ \text{True Negative Rate} & \quad \text{TNR}_D[h] = \mathbf{P}(h(x) = -1 \mid y = -1) \\ \text{Precision} & \quad \text{Prec}_D[h] = \mathbf{P}(y = 1 \mid h(x) = 1) = \frac{p \text{TPR}_D[h]}{p \text{TPR}_D[h] + (1-p)(1 - \text{TNR}_D[h])}. \end{aligned}$$

In this paper, we will consider non-decomposable performance measures that can be expressed as a function of the TPR and TNR and the class proportion p . Specifically, let $\Psi : [0, 1]^3 \rightarrow \mathbb{R}_+$; then the Ψ -performance of h w.r.t. D , which we will denote as $\mathcal{P}_D^\Psi[h]$, will be defined as:

$$\mathcal{P}_D^\Psi[h] = \Psi(\text{TPR}_D[h], \text{TNR}_D[h], p).$$

For example, for $\beta > 0$, the F_β -measure of h can be defined through the function $\Psi_{F_\beta} : [0, 1]^3 \rightarrow \mathbb{R}_+$ given by $\Psi_{F_\beta}(u, v, p) = \frac{(1+\beta^2)pu}{p+\beta^2(pu+(1-p)(1-v))}$, which gives $\mathcal{P}_D^{F_\beta}[h] = (1 + \beta^2) / \left(\frac{\beta^2}{\text{Prec}_D[h]} + \frac{1}{\text{TPR}_D[h]} \right)$. Table 1 gives several examples of non-decomposable performance measures that are used in practice. We will also find it useful to consider empirical versions of these performance measures calculated from a sample S , which we will denote as $\widehat{\mathcal{P}}_S^\Psi[h]$:

$$\widehat{\mathcal{P}}_S^\Psi[h] = \Psi(\widehat{\text{TPR}}_S[h], \widehat{\text{TNR}}_S[h], \widehat{p}_S), \quad (1)$$

where $\widehat{p}_S = \frac{1}{n} \sum_{i=1}^n \mathbf{1}(y_i = 1)$ is an empirical estimate of p , and

$$\widehat{\text{TPR}}_S[h] = \frac{1}{\widehat{p}_S n} \sum_{i=1}^n \mathbf{1}(h(x_i) = 1, y_i = 1); \quad \widehat{\text{TNR}}_S[h] = \frac{1}{(1 - \widehat{p}_S)n} \sum_{i=1}^n \mathbf{1}(h(x_i) = -1, y_i = -1)$$

are the empirical TPR and TNR respectively.¹

Ψ -consistency. We will be interested in the optimum value of \mathcal{P}_D^Ψ over all classifiers:

$$\mathcal{P}_D^{\Psi,*} = \sup_{h: \mathcal{X} \rightarrow \{\pm 1\}} \mathcal{P}_D^\Psi[h].$$

In particular, one can define the Ψ -regret of a classifier h as:

$$\text{regret}_D^\Psi[h] = \mathcal{P}_D^{\Psi,*} - \mathcal{P}_D^\Psi[h].$$

A learning algorithm is then said to be Ψ -consistent if the Ψ -regret of the classifier \widehat{h}_S output by the algorithm on seeing training sample S converges in probability to 0:²

$$\text{regret}_D^\Psi[\widehat{h}_S] \xrightarrow{P} 0.$$

Class of Threshold Classifiers. We will find it useful to define for any function $f : \mathcal{X} \rightarrow [0, 1]$, the set of classifiers obtained by assigning a threshold to f : $\mathcal{T}_f = \{\text{sign} \circ (f - t) \mid t \in [0, 1]\}$, where $\text{sign}(u) = 1$ if $u > 0$ and -1 otherwise. For a given f , we shall also define the thresholds corresponding to maximum population and empirical measures respectively (when they exist) as:

$$t_{D,f,\Psi}^* \in \underset{t \in [0,1]}{\text{argmax}} \mathcal{P}_D^\Psi[\text{sign} \circ (f - t)]; \quad \widehat{t}_{S,f,\Psi} \in \underset{t \in [0,1]}{\text{argmax}} \widehat{\mathcal{P}}_S^\Psi[\text{sign} \circ (f - t)].$$

Plug-in Algorithms and Result of Ye et al. (2012). In this work, we consider a family of plug-in algorithms, which divide the input sample S into samples (S_1, S_2) , use a suitable class probability estimation (CPE) algorithm to learn a class probability estimator $\widehat{\eta}_{S_1} : \mathcal{X} \rightarrow [0, 1]$ from S_1 , and output a classifier $\widehat{h}_S(x) = \text{sign}(\widehat{\eta}_{S_1}(x) - \widehat{t}_{S_2, \widehat{\eta}_{S_1}, \Psi})$, where $\widehat{t}_{S_2, \widehat{\eta}_{S_1}, \Psi}$ is a threshold that maximizes the empirical performance measure on S_2 (see Algorithm 1). We note that this approach is different from the idealized plug-in method analyzed by Ye et al. (2012) in the context of F-measure optimization, where a classifier is learned by assigning an empirical threshold to the ‘true’ class probability function η [15]; the consistency result therein is useful only if precise knowledge of η is available to a learning algorithm, which is not the case in most practical settings.

L_1 -consistency of a CPE algorithm. Let \mathcal{C} be a CPE algorithm, and for any sample S , denote $\widehat{\eta}_S = \mathcal{C}(S)$. We will say \mathcal{C} is L_1 -consistent w.r.t. a distribution D if $\mathbf{E}_x[|\widehat{\eta}_S(x) - \eta(x)|] \xrightarrow{P} 0$.

¹In the setting considered here, the goal is to maximize a (non-decomposable) function of expectations; we note that this is different from the decision-theoretic setting in [15], where one looks at the expectation of a non-decomposable performance measure on n examples, and seeks to maximize its limiting value as $n \rightarrow \infty$.

²We say $\phi(S)$ converges in probability to $a \in \mathbb{R}$, written as $\phi(S) \xrightarrow{P} a$, if $\forall \epsilon > 0$, $\mathbf{P}_{S \sim D^n}(|\phi(S) - a| \geq \epsilon) \rightarrow 0$ as $n \rightarrow \infty$.

Algorithm 1 Plug-in with Empirical Threshold for Performance Measure $\mathcal{P}^\Psi : 2^{\mathcal{X}} \rightarrow \mathbb{R}_+$

- 1: **Input:** $S = ((x_1, y_1), \dots, (x_n, y_n)) \in (\mathcal{X} \times \{\pm 1\})^n$
 - 2: **Parameter:** $\alpha \in (0, 1)$
 - 3: Let $S_1 = ((x_1, y_1), \dots, (x_{n_1}, y_{n_1}))$, $S_2 = ((x_{n_1+1}, y_{n_1+1}), \dots, (x_n, y_n))$, where $n_1 = \lceil n\alpha \rceil$
 - 4: Learn $\widehat{\eta}_{S_1} = \mathcal{C}(S_1)$, where $\mathcal{C} : \cup_{n=1}^{\infty} (\mathcal{X} \times \{\pm 1\})^n \rightarrow [0, 1]^{\mathcal{X}}$ is a suitable CPE algorithm
 - 5: $\widehat{t}_{S_2, \widehat{\eta}_{S_1}, \Psi} \in \operatorname{argmax}_{t \in [0, 1]} \widehat{\mathcal{P}}_{S_2}^\Psi[\operatorname{sign} \circ (\widehat{\eta}_{S_1} - t)]$
 - 6: **Output:** Classifier $\widehat{h}_S(x) = \operatorname{sign}(\widehat{\eta}_{S_1}(x) - \widehat{t}_{S_2, \widehat{\eta}_{S_1}, \Psi})$
-

3 A Generic Proof Template for Ψ -consistency of Plug-in Algorithms

We now give a general result for showing consistency of the plug-in method in Algorithm 1 for any performance measure that can be expressed as a continuous function of TPR and TNR, and for which the Bayes optimal classifier is obtained by suitably thresholding the class probability function.

Assumption A. We will say that a probability distribution D on $\mathcal{X} \times \{\pm 1\}$ satisfies Assumption A w.r.t. Ψ if $t_{D, \eta, \Psi}^*$ exists and is in $(0, 1)$, and the cumulative distribution functions of the random variable $\eta(x)$ conditioned on $y = 1$ and on $y = -1$, $\mathbf{P}(\eta(x) \leq z \mid y = 1)$ and $\mathbf{P}(\eta(x) \leq z \mid y = -1)$, are continuous at $z = t_{D, \eta, \Psi}^*$.³

Note that this assumption holds for any distribution D for which $\eta(x)$ conditioned on $y = 1$ and on $y = -1$ is continuous, and also for any D for which $\eta(x)$ conditioned on $y = 1$ and on $y = -1$ is mixed, provided the optimum threshold $t_{D, \eta, \Psi}^*$ for \mathcal{P}^Ψ exists and is not a point of discontinuity.

Under the above assumption, and assuming that the CPE algorithm used in Algorithm 1 is L_1 -consistent (which holds for any algorithm that uses a regularized empirical risk minimization of a proper loss [16, 28]), we have our main consistency result.

Theorem 1 (Ψ -consistency of Algorithm 1). Let $\Psi : [0, 1]^3 \rightarrow \mathbb{R}_+$ be continuous in each argument. Let D be a probability distribution on $\mathcal{X} \times \{\pm 1\}$ that satisfies Assumption A w.r.t. Ψ , and for which the Bayes optimal classifier is of the form $h^{\Psi, *}(x) = \operatorname{sign} \circ (\eta(x) - t_{D, \eta, \Psi}^*)$. If the CPE algorithm \mathcal{C} in Algorithm 1 is L_1 -consistent, then Algorithm 1 is Ψ -consistent w.r.t. D .

Before we prove the above theorem, we will find it useful to state the following lemmas. In our first lemma, we state that the TPR and TNR of a classifier constructed by thresholding a suitable class probability estimate at a fixed $c \in (0, 1)$ converge respectively to the TPR and TNR of the classifier obtained by thresholding the true class probability function η at c .

Lemma 2 (Convergence of TPR and TNR for fixed threshold). Let D be a distribution on $\mathcal{X} \times \{\pm 1\}$. Let $\widehat{\eta}_S : \mathcal{X} \rightarrow [0, 1]$ be generated by an L_1 -consistent CPE algorithm. Let $c \in (0, 1)$ be an a priori fixed constant such that the cumulative distribution functions $\mathbf{P}(\eta(x) \leq z \mid y = 1)$ and $\mathbf{P}(\eta(x) \leq z \mid y = -1)$ are continuous at $z = c$. We then have

$$\begin{aligned} \operatorname{TPR}_D[\operatorname{sign} \circ (\widehat{\eta}_S - c)] &\xrightarrow{P} \operatorname{TPR}_D[\operatorname{sign} \circ (\eta - c)]; \\ \operatorname{TNR}_D[\operatorname{sign} \circ (\widehat{\eta}_S - c)] &\xrightarrow{P} \operatorname{TNR}_D[\operatorname{sign} \circ (\eta - c)]. \end{aligned}$$

As a corollary to the above lemma, we have a similar result for \mathcal{P}^Ψ .

Lemma 3 (Convergence of \mathcal{P}^Ψ for fixed threshold). Let $\Psi : [0, 1]^3 \rightarrow \mathbb{R}_+$ be continuous in each argument. Under the statement of Lemma 2, we have

$$\mathcal{P}_D^\Psi[\operatorname{sign} \circ (\widehat{\eta}_S - c)] \xrightarrow{P} \mathcal{P}_D^\Psi[\operatorname{sign} \circ (\eta - c)].$$

We next state a result showing convergence of the empirical performance measure to its population value for a fixed classifier, and a uniform convergence result over a class of thresholded classifiers.

Lemma 4 (Concentration result for \mathcal{P}^Ψ). Let $\Psi : [0, 1]^3 \rightarrow \mathbb{R}_+$ be continuous in each argument. Then for any fixed $h : \mathcal{X} \rightarrow \{\pm 1\}$, and $\epsilon > 0$,

$$\mathbf{P}_{S \sim D^n} \left(\left| \mathcal{P}_S^\Psi[h] - \widehat{\mathcal{P}}_S^\Psi[h] \right| \geq \epsilon \right) \rightarrow 0 \text{ as } n \rightarrow \infty.$$

³For simplicity, we assume that $t_{D, \eta, \Psi}^*$ is in $(0, 1)$; our results easily extend to the case when $t_{D, \eta, \Psi}^* \in [0, 1]$.

Lemma 5 (Uniform convergence of \mathcal{P}^Ψ over threshold classifiers). Let $\Psi : [0, 1]^3 \rightarrow \mathbb{R}_+$ be continuous in each argument. For any $f : \mathcal{X} \rightarrow [0, 1]$ and $\epsilon > 0$,

$$\mathbf{P}_{S \sim D^n} \left(\bigcup_{\theta \in \mathcal{T}_f} \left\{ \left| \mathcal{P}_D^\Psi[\theta] - \widehat{\mathcal{P}}_S^\Psi[\theta] \right| \geq \epsilon \right\} \right) \rightarrow 0 \quad \text{as } n \rightarrow \infty.$$

We are now ready to prove our main theorem.

Proof of Theorem 1. Recall $t_{D, \eta, \Psi}^* \in \operatorname{argmax}_{t \in [0, 1]} \mathcal{P}_D^\Psi[\operatorname{sign} \circ (\eta - t)]$ exists by Assumption A. In the following, we shall use t^* in the place of $t_{D, \eta, \Psi}^*$ and \widehat{t}_{S_2, S_1} in the place of $\widehat{t}_{S_2, \widehat{\eta}_{S_1}, \Psi}$. We have

$$\begin{aligned} \operatorname{regret}_D^\Psi[h_S] &= \operatorname{regret}_D^\Psi[\operatorname{sign} \circ (\widehat{\eta}_{S_1} - \widehat{t}_{S_2, S_1})] \\ &= \mathcal{P}_D^{\Psi, *} - \mathcal{P}_D^\Psi[\operatorname{sign} \circ (\widehat{\eta}_{S_1} - \widehat{t}_{S_2, S_1})] \\ &= \mathcal{P}_D^\Psi[\operatorname{sign} \circ (\eta - t^*)] - \mathcal{P}_D^\Psi[\operatorname{sign} \circ (\widehat{\eta}_{S_1} - \widehat{t}_{S_2, S_1})], \end{aligned}$$

which follows from the assumption on the Bayes optimal classifier for \mathcal{P}^Ψ . Adding and subtracting empirical and population versions of \mathcal{P}^Ψ computed on certain classifiers,

$$\begin{aligned} \operatorname{regret}_D^\Psi[\operatorname{sign} \circ (\widehat{\eta}_{S_1} - \widehat{t}_{S_2, S_1})] &= \underbrace{\mathcal{P}_D^\Psi[\operatorname{sign} \circ (\eta - t^*)] - \mathcal{P}_D^\Psi[\operatorname{sign} \circ (\widehat{\eta}_{S_1} - t^*)]}_{\text{term}_1} \\ &\quad + \underbrace{\mathcal{P}_D^\Psi[\operatorname{sign} \circ (\widehat{\eta}_{S_1} - t^*)] - \widehat{\mathcal{P}}_{S_2}^\Psi[\operatorname{sign} \circ (\widehat{\eta}_{S_1} - \widehat{t}_{S_2, S_1})]}_{\text{term}_2} \\ &\quad + \underbrace{\widehat{\mathcal{P}}_{S_2}^\Psi[\operatorname{sign} \circ (\widehat{\eta}_{S_1} - \widehat{t}_{S_2, S_1})] - \mathcal{P}_D^\Psi[\operatorname{sign} \circ (\widehat{\eta}_{S_1} - \widehat{t}_{S_2, S_1})]}_{\text{term}_3}. \end{aligned}$$

We now show convergence for each of the above terms. Applying Lemma 3 with $c = t^*$ (by Assumption A, $t^* \in (0, 1)$ and satisfies the necessary continuity assumption), we have $\text{term}_1 \xrightarrow{P} 0$. For term_2 , from the definition of threshold \widehat{t}_{S_2, S_1} (see Algorithm 1), we have

$$\text{term}_2 \leq \mathcal{P}_D^\Psi[\operatorname{sign} \circ (\widehat{\eta}_{S_1} - t^*)] - \widehat{\mathcal{P}}_{S_2}^\Psi[\operatorname{sign} \circ (\widehat{\eta}_{S_1} - t^*)]. \quad (2)$$

Then for any $\epsilon > 0$,

$$\begin{aligned} \mathbf{P}_{S \sim D^n} (\text{term}_2 \geq \epsilon) &= \mathbf{P}_{S_1 \sim D^{n_1}, S_2 \sim D^{n-n_1}} (\text{term}_2 \geq \epsilon) \\ &= \mathbf{E}_{S_1} \left[\mathbf{P}_{S_2 | S_1} (\text{term}_2 \geq \epsilon) \right] \\ &\leq \mathbf{E}_{S_1} \left[\mathbf{P}_{S_2 | S_1} \left(\left| \mathcal{P}_D^\Psi[\operatorname{sign} \circ (\widehat{\eta}_{S_1} - t^*)] - \widehat{\mathcal{P}}_{S_2}^\Psi[\operatorname{sign} \circ (\widehat{\eta}_{S_1} - t^*)] \right| \geq \epsilon \right) \right] \\ &\rightarrow 0 \end{aligned}$$

as $n \rightarrow \infty$, where the third step follows from Eq. (2), and the last step follows by applying, for a fixed S_1 , the concentration result in Lemma 4 with $h = \operatorname{sign} \circ (\widehat{\eta}_{S_1} - t^*)$ (given continuity of Ψ). Finally, for term_3 , we have for any $\epsilon > 0$,

$$\begin{aligned} \mathbf{P}_S (\text{term}_3 \geq \epsilon) &= \mathbf{E}_{S_1} \left[\mathbf{P}_{S_2 | S_1} \left(\widehat{\mathcal{P}}_{S_2}^\Psi[\operatorname{sign} \circ (\widehat{\eta}_{S_1} - \widehat{t}_{S_2, S_1})] - \mathcal{P}_D^\Psi[\operatorname{sign} \circ (\widehat{\eta}_{S_1} - \widehat{t}_{S_2, S_1})] \geq \epsilon \right) \right] \\ &\leq \mathbf{E}_{S_1} \left[\mathbf{P}_{S_2 | S_1} \left(\bigcup_{\theta \in \mathcal{T}_{\widehat{\eta}_{S_1}}} \left\{ \left| \widehat{\mathcal{P}}_{S_2}^\Psi[\theta] - \mathcal{P}_D^\Psi[\theta] \right| \geq \epsilon \right\} \right) \right] \\ &\rightarrow 0 \end{aligned}$$

as $n \rightarrow \infty$, where the last step follows by applying the uniform convergence result in Lemma 5 over the class of thresholded classifiers $\mathcal{T}_{\widehat{\eta}_{S_1}} = \{\operatorname{sign} \circ (\widehat{\eta}_{S_1} - t) \mid t \in [0, 1]\}$ (for a fixed S_1). \square

4 Consistency of Plug-in Algorithms for AM, F_β , and G-TP/PR

We now use the result in Theorem 1 to establish consistency of the plug-in algorithms for the arithmetic mean of TPR and TNR, the F_β -measure, and the geometric mean of TPR and precision.

4.1 Consistency for AM-measure

The arithmetic mean of TPR and TNR (AM) or one minus the balanced error rate (BER) is a widely-used performance measure in class imbalanced binary classification settings [17–19]:

$$\mathcal{P}_D^{\text{AM}}[h] = \frac{\text{TPR}_D[h] + \text{TNR}_D[h]}{2}.$$

It can be shown that Bayes optimal classifier for the AM-measure is of the form $h^{\text{AM},*}(x) = \text{sign} \circ (\eta(x) - p)$ (see for example [16]), and that the threshold chosen by the plug-in method in Algorithm 1 for the AM-measure is an empirical estimate of p . In recent work, Menon et al. show that this plug-in method is consistent w.r.t. the AM-measure [16]; their proof makes use of a decomposition of the AM-measure in terms of a certain cost-sensitive error and a result of [22] on regret bounds for cost-sensitive classification. We now use our result in Theorem 1 to give an alternate route for showing AM-consistency of this plug-in method.⁴

Theorem 6 (Consistency of Algorithm 1 w.r.t. AM-measure). Let $\Psi = \Psi_{\text{AM}}$. Let D be a distribution on $\mathcal{X} \times \{\pm 1\}$ that satisfies Assumption A w.r.t. Ψ_{AM} . If the CPE algorithm \mathcal{C} in Algorithm 1 is L_1 -consistent, then Algorithm 1 is AM-consistent w.r.t. D .

Proof. We apply Theorem 1 noting that $\Psi_{\text{AM}}(u, v, p) = (u+v)/2$ is continuous in all its arguments, and that the Bayes optimal classifier for \mathcal{P}^{AM} is of the requisite thresholded form. \square

4.2 Consistency for F_β -measure

The F_β -measure or the (weighted) harmonic mean of TPR and precision is a popular performance measure used in information retrieval [1]:

$$\mathcal{P}_D^{F_\beta}[h] = \frac{(1 + \beta^2)\text{TPR}_D[h]\text{Prec}_D[h]}{\beta^2\text{TPR}_D[h] + \text{Prec}_D[h]} = \frac{(1 + \beta^2)p\text{TPR}_D[h]}{p + \beta^2(p\text{TPR}_D[h] + (1 - p)(1 - \text{TNR}_D[h]))},$$

where $\beta \in (0, 1)$ controls the trade-off between TPR and precision. In a recent work, Ye et al. [15] show that the optimal classifier for the F_β -measure is the class probability η thresholded suitably.

Lemma 7 (Optimality of threshold classifiers for F_β -measure; Ye et al. (2012) [15]). For any distribution D over $\mathcal{X} \times \{\pm 1\}$ that satisfies Assumption A w.r.t. Ψ , the Bayes optimal classifier for \mathcal{P}^{F_β} is of the form $h^{F_\beta,*}(x) = \text{sign} \circ (\eta(x) - t_{D,\eta,F_\beta}^*)$.

As noted earlier, the authors in [15] show that an idealized plug-in method that applies an empirically determined threshold to the ‘true’ class probability η is consistent w.r.t. the F_β -measure. This result is however useful only when the ‘true’ class probability is available to a learning algorithm, which is not the case in most practical settings. On the other hand, the plug-in method considered in our work constructs a classifier by assigning an empirical threshold to a suitable ‘estimate’ of the class probability. Using Theorem 1, we now show that this method is consistent w.r.t. the F_β -measure.

Theorem 8 (Consistency of Algorithm 1 w.r.t. F_β -measure). Let $\Psi = \Psi_{F_\beta}$ in Algorithm 1. Let D be a distribution on $\mathcal{X} \times \{\pm 1\}$ that satisfies Assumption A w.r.t. Ψ_{F_β} . If the CPE algorithm \mathcal{C} in Algorithm 1 is L_1 -consistent, then Algorithm 1 is F_β -consistent w.r.t. D .

Proof. We apply Theorem 1 noting that $\Psi_{F_\beta}(u, v, p) = \frac{(1+\beta^2)pu}{p+\beta^2(pu+(1-p)(1-v))}$ is continuous in each argument, and that (by Lemma 7) the Bayes optimal classifier for \mathcal{P}^{F_β} is of the requisite form. \square

4.3 Consistency for G-TP/PR

The geometric mean of TPR and precision (G-TP/PR) is another performance measure proposed for class imbalanced classification problems [3]:

$$\mathcal{P}_D^{\text{G-TP/PR}}[h] = \sqrt{\text{TPR}_D[h]\text{Prec}_D[h]} = \sqrt{\frac{p\text{TPR}_D[h]^2}{p\text{TPR}_D[h] + (1 - p)(1 - \text{TNR}_D[h])}}.$$

⁴Note that the plug-in classification threshold chosen for the AM-measure is the same independent of the class probability estimate used; our consistency results will therefore apply in this case even if one uses, as in [16], the same sample for both learning a class probability estimate, and estimating the plug-in threshold.

We first show that the optimal classifier for G-TP/PR is obtained by thresholding the class probability function η at a suitable point; our proof uses a technique similar to the one for the F_β -measure in [15].

Lemma 9 (Optimality of threshold classifiers for G-TP/PR). For any distribution D on $\mathcal{X} \times \{\pm 1\}$ that satisfies Assumption A w.r.t. Ψ , the Bayes optimal classifier for $\mathcal{P}^{\text{G-TP/PR}}$ is of the form $h^{\text{G-TP/PR},*}(x) = \text{sign}(\eta(x) - t_{D,\eta,\text{G-TP/PR}}^*)$.

Theorem 10 (Consistency of Algorithm 1 w.r.t. G-TP/PR). Let $\Psi = \Psi_{\text{G-TP/PR}}$. Let D be a distribution on $\mathcal{X} \times \{\pm 1\}$ that satisfies Assumption A w.r.t. $\Psi_{\text{G-TP/PR}}$. If the CPE algorithm \mathcal{C} in Algorithm 1 is L_1 -consistent, then Algorithm 1 is G-TP/PR-consistent w.r.t. D .

Proof. We apply Theorem 1 noting that $\Psi_{\text{G-TP/PR}}(u, v, p) = \sqrt{\frac{pu^2}{pu+(1-p)(1-v)}}$ is continuous in each argument, and that (by Lemma 9) the Bayes optimal classifier for $\mathcal{P}^{\text{G-TP/PR}}$ is of the requisite form. \square

5 Consistency of Plug-in Algorithms for Non-decomposable Performance Measures that are Monotonic in TPR and TNR

The consistency results seen so far apply to any distribution that satisfies a mild continuity condition at the optimal threshold for a performance measure, and have crucially relied on the specific functional form of the measure. In this section, we shall see that under a stricter continuity assumption on the distribution, the empirical plug-in algorithm can be shown to be consistent w.r.t. any performance measure that is a continuous and monotonically increasing function of TPR and TNR.

Assumption B. We will say that a probability distribution D on $\mathcal{X} \times \{\pm 1\}$ satisfies Assumption B w.r.t. Ψ if $t_{D,\eta,\Psi}^*$ exists and is in $(0, 1)$, and the cumulative distribution function of the random variable $\eta(x)$, $\mathbf{P}(\eta(x) \leq z)$, is continuous at all $z \in (0, 1)$.

Distributions that satisfy the above assumption also satisfy Assumption A. We show that under this assumption, the optimal classifier for any performance measure that is monotonically increasing in TPR and TNR is obtained by thresholding η , and this holds irrespective of the specific functional form of the measure. An application of Theorem 1 then gives us the desired consistency result.

Lemma 11 (Optimality of threshold classifiers for monotonic Ψ under distributional assumption). Let $\Psi : [0, 1]^3 \rightarrow \mathbb{R}_+$ be monotonically increasing in its first two arguments. Then for any distribution D on $\mathcal{X} \times \{\pm 1\}$ that satisfies Assumption B, the Bayes optimal classifier for \mathcal{P}^Ψ is of the form $h^{\Psi,*}(x) = \text{sign}(\eta(x) - t_{D,\eta,\Psi}^*)$.

Theorem 12 (Consistency of Algorithm 1 for monotonic Ψ under distributional assumption). Let $\Psi : [0, 1]^3 \rightarrow \mathbb{R}_+$ be continuous in each argument, and monotonically increasing in its first two arguments. Let D be a distribution on $\mathcal{X} \times \{\pm 1\}$ that satisfies Assumption B. If the CPE algorithm \mathcal{C} in Algorithm 1 is L_1 -consistent, then Algorithm 1 is Ψ -consistent w.r.t. D .

Proof. We apply Theorem 1 by using the continuity assumption on Ψ , and noting that, by Lemma 11 and monotonicity of Ψ , the Bayes optimal classifier for \mathcal{P}^Ψ is of the requisite form. \square

The above result applies to all performance measures listed in Table 1, and in particular, to the geometric, harmonic, and quadratic means of TPR and TNR [2–5], for which the Bayes optimal classifier need not be of the requisite thresholded form for a general distribution (see Appendix C).

6 Experiments

We performed two types of experiments. The first involved synthetic data, where we demonstrate diminishing regret of the plug-in method in Algorithm 1 with growing sample size for different performance measures; since the data is generated from a known distribution, exact calculation of regret is possible here. The second involved real data, where we show that the plug-in algorithm is competitive with the state-of-the-art SVM^{perf} algorithm for non-decomposable measures (SVMPerf) [12]; we also include for comparison a plug-in method with a fixed threshold of 0.5 (Plug-in (0-1)). We consider three performance measures here: F_1 -measure, G-TP/PR and G-Mean (see Table 1).

Synthetic data. We generated data from a known distribution (class conditionals are multivariate Gaussians with mixing ratio p and equal covariance matrices) for which the optimal classifier for




Figure 1: Experiments on synthetic data with $p = 0.5$: regret as a function of number of training examples using various methods for the F_1 , G-TP/PR and G-mean performance measures.




Figure 2: Experiments on synthetic data with $p = 0.1$: regret as a function of number of training examples using various methods for the F_1 , G-TP/PR and G-Mean performance measures.




Figure 3: Experiments on real data: results for various methods (using linear models) on four data sets in terms of F_1 , G-TP/PR and G-Mean performance measures. Here N , d , p refer to the number of instances, number of features and fraction of positives in the data set respectively.

each performance measure considered here is linear, making it sufficient to learn a linear model; the distribution satisfies Assumption B w.r.t. each performance measure. We used regularized logistic regression as the CPE method in Algorithm 1 in order to satisfy the L_1 -consistency condition in Theorem 1 (see Appendix A.1 and A.4 for details). The experimental results are shown in Figures 1 and 2 for $p = 0.5$ and $p = 0.1$ respectively. In each case, the regret for the empirical plug-in method (Plug-in (F1), Plug-in (G-TP/PR) and Plug-in (GM)) goes to zero with increasing training set size, validating our consistency results; SVM^{perf} often fails to exhibit diminishing regret for $p = 0.1$; and as expected, Plug-in (0-1), with its apriori fixed threshold, fails to be consistent in most cases.

Real data. We ran the three algorithms described earlier over data sets drawn from the UCI ML repository [29] and a cheminformatics data set obtained from [30], and report their performance on held-out test sets. Figure 3 contains results for four data sets averaged over 10 random train-test splits of the original data (also see Appendix A.2 and A.3). Clearly, in most cases, the empirical plug-in performs comparable to SVM^{perf} and outperforms Plug-in (0-1). Moreover, the empirical plug-in was found to exhibit faster run-times than the SVM^{perf} method (see Figure 5 in Appendix).

7 Conclusions

We have presented a general method for proving consistency of plug-in algorithms that assign an empirical threshold to a suitable class probability estimate for a variety of non-decomposable performance measures for binary classification that can be expressed as a continuous function of TPR and TNR, and for which the Bayes optimal classifier is the class probability function thresholded suitably. We use our template to show consistency for the AM, F_β and G-TP/PR measures, and under a continuous distribution, for any performance measure that is continuous and monotonic in TPR and TNR. Our experiments suggest that these algorithms yield performance comparable to the state-of-the-art SVM^{perf} method, while being faster than this method in practice.

Acknowledgments

HN thanks support from a Google India PhD Fellowship. SA gratefully acknowledges support from DST, Indo-US Science and Technology Forum, and an unrestricted gift from Yahoo.

References

- [1] C. D. Manning, P. Raghavan, and H. Schütze. *Introduction to Information Retrieval*. Cambridge University Press, 2008.
- [2] M. Kubat and S. Matwin. Addressing the curse of imbalanced training sets: One-sided selection. In *ICML*, 1997.
- [3] S. Daskalaki, I. Kopanas, and N. Avouris. Evaluation of classifiers for an uneven class distribution problem. *Applied Artificial Intelligence*, 20:381–417, 2006.
- [4] K. Kennedy, B.M. Namee, and S.J. Delany. Learning without default: a study of one-class classification and the low-default portfolio problem. In *ICAICS*, 2009.
- [5] S. Lawrence, I. Burns, A. Back, A-C. Tsoi, and C.L. Giles. Neural network classification and prior class probabilities. In *Neural Networks: Tricks of the Trade*, pages 1524:299–313. 1998.
- [6] Y. Yang. A study of thresholding strategies for text categorization. In *SIGIR*, 2001.
- [7] D.D. Lewis. Evaluating and optimizing autonomous text classification systems. In *SIGIR*, 1995.
- [8] K.M.A. Chai. Expectation of F-measures: Tractable exact computation and some empirical observations of its properties. In *SIGIR*, 2005.
- [9] D.R. Musicant, V. Kumar, and A. Ozgur. Optimizing F-measure with support vector machines. In *FLAIRS*, 2003.
- [10] S. Gao, W. Wu, C-H. Lee, and T-S. Chua. A maximal figure-of-merit learning approach to text categorization. In *SIGIR*, 2003.
- [11] M. Jansche. Maximum expected F-measure training of logistic regression models. In *HLT*, 2005.
- [12] T. Joachims. A support vector method for multivariate performance measures. In *ICML*, 2005.
- [13] Z. Liu, M. Tan, and F. Jiang. Regularized F-measure maximization for feature selection and classification. *BioMed Research International*, 2009, 2009.
- [14] P.M. Chinta, P. Balamurugan, S. Shevade, and M.N. Murty. Optimizing F-measure with non-convex loss and sparse linear classifiers. In *IJCNN*, 2013.
- [15] N. Ye, K.M.A. Chai, W.S. Lee, and H.L. Chieu. Optimizing F-measures: A tale of two approaches. In *ICML*, 2012.
- [16] A.K. Menon, H. Narasimhan, S. Agarwal, and S. Chawla. On the statistical consistency of algorithms for binary classification under class imbalance. In *ICML*, 2013.
- [17] J. Cheng, C. Hatzis, H. Hayashi, M-A. Krogel, S. Morishita, D. Page, and J. Sese. KDD Cup 2001 report. *ACM SIGKDD Explorations Newsletter*, 3(2):47–64, 2002.
- [18] R. Powers, M. Goldszmidt, and I. Cohen. Short term performance forecasting in enterprise systems. In *KDD*, 2005.
- [19] Q. Gu, L. Zhu, and Z. Cai. Evaluation measures of the classification performance of imbalanced data sets. In *Computational Intelligence and Intelligent Systems*, volume 51, pages 461–471. 2009.
- [20] T. Zhang. Statistical behaviour and consistency of classification methods based on convex risk minimization. *Annals of Mathematical Statistics*, 32:56–134, 2004.
- [21] P.L. Bartlett, M.I. Jordan, and J.D. McAuliffe. Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473):138–156, 2006.
- [22] C. Scott. Calibrated asymmetric surrogate losses. *Electronic Journal of Statistics*, 6:958–992, 2012.
- [23] M. Zhao, N. Edakunni, A. Pocock, and G. Brown. Beyond Fano’s inequality: Bounds on the optimal F-score, BER, and cost-sensitive risk and their implications. *Journal of Machine Learning Research*, 14(1):1033–1090, 2013.
- [24] Z.C. Lipton, C. Elkan, and B. Naryanaswamy. Optimal thresholding of classifiers to maximize F1 measure. In *ECML/PKDD*, 2014.
- [25] J. Petterson and T. Caetano. Reverse multi-label learning. In *NIPS*, 2010.
- [26] K. Dembczynski, W. Waegeman, W. Cheng, and E. Hüllermeier. An exact algorithm for F-measure maximization. In *NIPS*, 2011.
- [27] K. Dembczynski, A. Jachnik, W. Kotlowski, W. Waegeman, and E. Huellermeier. Optimizing the F-measure in multi-label classification: Plug-in rule approach versus structured loss minimization. In *ICML*, 13.
- [28] S. Agarwal. Surrogate regret bounds for the area under the ROC curve via strongly proper losses. In *COLT*, 2013.
- [29] A. Frank and A. Asuncion. UCI machine learning repository, 2010. URL: <http://archive.ics.uci.edu/ml>.
- [30] Robert N. Jorissen and Michael K. Gilson. Virtual screening of molecular databases using a support vector machine. *Journal of Chemical Information and Modeling*, 45:549–561, 2005.

On the Statistical Consistency of Plug-in Classifiers for Non-decomposable Performance Measures

Appendix

A Experimental Details

A.1 Synthetic data experiments

We use data drawn from a distribution D over $(\mathcal{X} = \mathbb{R}^{10}) \times \{\pm 1\}$ that satisfies Assumption B (and therefore assumption A); recall that our plug-in consistency results for the F_1 and G-TP/PR measures apply to distributions that satisfy Assumption A (see Section 4), and the consistency result for G-Mean holds for distributions that satisfy Assumption B (see Section 5). The specifics of our experiments mirror those used in [16] and are listed here for completeness: positive examples ($y = 1$) are drawn from $\mathcal{N}(\mu, \Sigma)$ with probability $p \in (0, 1)$ and negative examples ($y = -1$) drawn from $\mathcal{N}(-\mu, \Sigma)$ with probability $(1 - p)$; μ is drawn uniformly at random from $\{\pm 1\}^{10}$ and $\Sigma \in \mathbb{R}^{10 \times 10}$ is drawn from a Wishart distribution with 20 degrees of freedom and a randomly drawn invertible positive semidefinite scale matrix. As pointed out earlier, the optimal classifier for each performance measure considered here under this distribution is linear, making it sufficient to learn a linear model (see Section A.4).

We evaluate the statistical regret of the empirical plug-in method (Algorithm 1 with $\alpha = 0.5$) and compare it against SVM^{perf} with linear kernel (SVMPerf) adapted to optimize the performance measures considered here⁵, and the Plug-in algorithm with a default threshold 0.5 (Plug-in (0-1)). The empirical plug-in algorithm (denoted for the three performance measures respectively as Plug-in (F1), Plug-in (G-TP/PR) and Plug-in (GM)) randomly splits the input data S (drawn from D) into samples S_1 and S_2 for the purposes of learning a class probability estimate and choosing an appropriate threshold respectively; we use regularized (linear) logistic regression for learning a class probability estimate from S_1 , with the regularization parameter set to $1/\sqrt{|S_1|}$, in order to satisfy the L_1 -consistency requirement in Theorem 1 (see [16, 20] for details). The Plug-in (0-1) method learns a class probability estimate using the entire input data (S_1, S_2) , with the regularization parameter set to $1/\sqrt{|S|}$. The SVM^{perf} algorithm also uses the entire input data (S_1, S_2) , with the regularization parameter selected via 5-fold cross-validation over the training sample; for the experiments involving $p = 0.5$, the regularization parameter was chosen from the range $\{10^{-3}, 10^{-2}, \dots, 10^1\}$ and for the experiments involving $p = 0.1$, this parameter was chosen from $\{10^{-7}, 10^{-2}, \dots, 10^2\}$.⁶ The tolerance parameter for SVM^{perf} was set to 0.1 for all experiments.

A.2 Real data experiments

For experiments with real data sets, we report the performance of the learned classifiers on separately held test data (we perform a random 2:1 train-test split of the original data, preserving class proportions). For the empirical plug-in algorithm, the parameter α in Algorithm 1 was set to 0.8. The regularization parameter for SVM^{perf} was chosen from the range $\{10^{-7}, 10^{-2}, \dots, 10^4\}$ and that for logistic regression was chosen from the range $\{10^{-3}, 10^{-2}, \dots, 10^3\}$ using 5-fold cross validation over the corresponding training sample. The tolerance parameter for SVM^{perf} was set to 0.1 for all experiments.

A.3 Additional results on real data

Table 2 summarizes all the real data sets that have been used in our experiments (both in Section 6 and this section). Figure 4 shows the test performances of the plug-in (with both the empirically chosen threshold and default threshold) and SVM^{perf} methods w.r.t. F_1 , G-TP/PR and G-Mean

⁵We used the SVM^{perf} routine provided in http://www.cs.cornell.edu/people/tj/svm_light/svm_perf.html for the F_1 -measure; we made necessary modifications to this code (as prescribed in [12]) to optimize G-TP/PR and G-Mean.

⁶Here, we cannot set the regularization parameter to $1/\sqrt{|S|}$ since the theoretical prescriptions of [20] are not applicable to multivariate extension of hinge loss optimized by SVM^{perf}.

Table 2: Summary of real data sets used in this study

Data sets	#examples	#features	$p = \mathbf{P}(y = 1)$
car	1728	21	0.038
chemo-a1a	2111	1021	0.024
nursery	12960	27	0.026
letter	18668	16	0.034
spambase	4210	57	0.398
splice	3005	240	0.226
thyroid	7129	21	0.977




Figure 4: Experiments on real data: results for empirical Plug-in, SVM^{perf} and Plug-in (0-1) methods (with linear models) on several UCI data sets in terms of F_1 , G-TP/PR and G-Mean performance measures. Here N, d, p refer to the number of instances, number of features and fraction of positives in the data set respectively.

measures over data sets included in Table 2 that were not already covered in Figure 3. Table 3 lists experiment results for all datasets, all algorithms and all performance measures. It can be observed that the empirical Plug-in is competitive with SVM^{perf} and outperforms the Plug-in (0-1) method in most cases. Figure 5 shows comparison of the average run-times of these methods (across the 10 training sets) on different performance measures and data sets. It is clearly seen that the empirical Plug-in method analyzed in this work offers significant speedups compared to SVM^{perf} in terms of the time taken for cross-validation, while yielding comparable training time.

A.4 Regret calculation for synthetic data

As mentioned in Section A.1, the distribution D over $\mathbb{R}^{10} \times \{\pm 1\}$ that we consider consists of multivariate Gaussian class conditional distributions, with positive instances being drawn from $\mathcal{N}(\mu, \Sigma)$ and negative instances being drawn from $\mathcal{N}(-\mu, \Sigma)$. We denote the probability density functions (pdfs) corresponding to $x|y = 1$ and $x|y = -1$ as f_+ and f_- respectively.

We first show that any classifier obtained by thresholding the the class probability function η under the above distribution is linear. For any $x \in \mathbb{R}^{10}$, we have

$$\begin{aligned} \eta(x) &= \mathbf{P}(y = 1|x) = \frac{\mathbf{P}(x|y = 1) \cdot \mathbf{P}(y = 1)}{\mathbf{P}(x|y = 1) \cdot \mathbf{P}(y = 1) + \mathbf{P}(x|y = -1) \cdot \mathbf{P}(y = -1)} \\ &= \frac{p \cdot f_+(x)}{p \cdot f_+(x) + (1-p) f_-(x)} \\ &= \frac{1}{1 + e^{-f(x)}}, \end{aligned}$$

where $f(x) = \ln\left(\frac{p \cdot f_+(x)}{(1-p) f_-(x)}\right) = 2\mu^T \Sigma^{-1} x + \ln\left(\frac{p}{1-p}\right)$ turns out to be a linear function of x . As a result, any thresholded classifier of the form $\text{sign} \circ (\eta(x) - c)$, for some $c \in (0, 1)$, can be written as a linear classifier: $\text{sign} \circ (f(x) - \ln(c/(1-c)))$.

We next describe how one can compute the Ψ -regret of a linear classifier $h : \mathbb{R}^{10} \rightarrow \{\pm 1\}$ under the given distribution:

$$\text{regret}_D^\Psi[h] = \mathcal{P}_D^{\Psi,*} - \mathcal{P}_D^\Psi[h].$$

In particular, we shall describe how the values of $\mathcal{P}_D^\Psi[h]$ and $\mathcal{P}_D^{\Psi,*}$ in the above expression can be computed for the given distribution D .

Table 3: Results from experiments on all real datasets. For each dataset and performance measure, the algorithm outputs with the highest and second-highest mean performance are highlighted in boldface and italics respectively.

Data sets	Algorithm	F ₁	G-TP/PR	G-Mean
car	Emp. Plug-in	0.8520 ± 0.0574	0.8602 ± 0.0506	<i>0.9461 ± 0.0427</i>
	SVM ^{perf}	<i>0.8518 ± 0.0402</i>	<i>0.8491 ± 0.0355</i>	0.9864 ± 0.0050
	Plug-in (0-1)	0.7475 ± 0.0474	0.7692 ± 0.0390	0.7801 ± 0.0439
chemo-ala	Emp. Plug-in	0.7063 ± 0.1058	<i>0.7234 ± 0.0896</i>	<i>0.8477 ± 0.0534</i>
	SVM ^{perf}	0.7288 ± 0.0653	0.7233 ± 0.0629	0.9207 ± 0.0269
	Plug-in (0-1)	<i>0.7086 ± 0.1050</i>	0.7352 ± 0.0884	0.7534 ± 0.0874
nursery	Emp. Plug-in	<i>0.7299 ± 0.0256</i>	<i>0.7252 ± 0.0417</i>	<i>0.9598 ± 0.0285</i>
	SVM ^{perf}	0.7655 ± 0.0195	0.7662 ± 0.0182	0.9611 ± 0.0113
	Plug-in (0-1)	0.4450 ± 0.0468	0.5278 ± 0.0342	0.5380 ± 0.0378
letter	Emp. Plug-in	<i>0.6667 ± 0.0255</i>	<i>0.6694 ± 0.0244</i>	0.9026 ± 0.0133
	SVM ^{perf}	0.7101 ± 0.0094	0.7105 ± 0.0098	<i>0.8945 ± 0.0155</i>
	Plug-in (0-1)	0.0817 ± 0.0182	0.2055 ± 0.0234	0.2055 ± 0.0234
spambase	Emp. Plug-in	<i>0.8470 ± 0.0096</i>	<i>0.8474 ± 0.0096</i>	<i>0.8737 ± 0.0082</i>
	SVM ^{perf}	0.9076 ± 0.0079	0.9069 ± 0.0098	0.9229 ± 0.0073
	Plug-in (0-1)	0.8039 ± 0.0122	0.8089 ± 0.0117	0.8289 ± 0.0101
splice	Emp. Plug-in	<i>0.9287 ± 0.0100</i>	<i>0.9292 ± 0.0099</i>	0.9611 ± 0.0097
	SVM ^{perf}	0.9233 ± 0.0133	0.9234 ± 0.0142	0.9496 ± 0.0060
	Plug-in (0-1)	0.9376 ± 0.0093	0.9376 ± 0.0093	<i>0.9583 ± 0.0058</i>
thyroid	Emp. Plug-in	<i>0.9937 ± 0.0015</i>	<i>0.9936 ± 0.0015</i>	<i>0.9461 ± 0.0155</i>
	SVM ^{perf}	0.9950 ± 0.0009	0.9948 ± 0.0008	0.9753 ± 0.0100
	Plug-in (0-1)	0.9887 ± 0.0002	0.9887 ± 0.0002	0.1756 ± 0.0749




Figure 5: Experiments on real data: average CPU run-time performance (in seconds) of empirical Plug-in, SVM^{perf} and Plug-in (0-1) methods (with linear models) on several UCI data sets. In each plot, the portion in black at the top of each bar corresponds to the training time, while the remaining portion corresponds to the time taken for cross-validation.

We start with the procedure outlined in [16] for calculating the performance measure \mathcal{P}_D^Ψ for any linear classifier $h(x) = \text{sign} \circ (w^\top x + b)$, where (for our purpose) $w \in \mathbb{R}^{10}$ and $b \in \mathbb{R}$. The TPR of h is given by:

$$\text{TPR}_D[h] = \mathbf{P}(h(x) = 1 | y = 1) = \int_{x | w^\top x + b \geq 0} f_+(x) dx.$$

It can be seen that $w^\top x | y = 1$ follows the normal distribution $\mathcal{N}(w^\top x, w^\top \Sigma w)$, and therefore by change of variables, we have

$$\text{TPR}_D[h] = \int_{-b}^{\infty} g_+(x) dx,$$

where g_+ is the pdf corresponding to $\mathcal{N}(w^\top x, w^\top \Sigma w)$. Likewise, the TNR for h is given by:

$$\text{TNR}_D[h] = \int_{-\infty}^{-b} g_-(x) dx,$$

where g_- is the pdf corresponding to $\mathcal{N}(-w^\top x, w^\top \Sigma w)$. This way, given w and b , both TPR and TNR are straightforward to determine, and consequently so is any performance measure that is a function Ψ of these quantities.

We next describe how the optimal value of the given performance measure $\mathcal{P}_D^{\Psi,*}$ can be computed. Since the given distribution satisfies Assumptions A and B, the optimal classifier for all performance measures considered in this work is obtained by suitably thresholding the class probability function η ; hence the optimal value $\mathcal{P}_D^{\Psi,*}$ for the given measure can be computed by performing a line search over $(0, 1)$ and picking the threshold $c^* \in (0, 1)$ for which the linear classifier $\text{sign} \circ (f(x) - c^*)$ maximizes the performance measure.

B Complete Proofs for Lemmas

B.1 Complete proof for Lemma 2

Proof. First, we simplify what we need to prove. We need to show that for a fixed $c \in (0, 1)$,

$$\begin{aligned}
& \text{TPR}_D[\text{sign} \circ (\hat{\eta}_S(x) - c)] \xrightarrow{P} \text{TPR}_D[\text{sign} \circ (\eta(x) - c)] \text{ and} \\
& \quad \text{TNR}_D[\text{sign} \circ (\hat{\eta}_S(x) - c)] \xrightarrow{P} \text{TNR}_D[\text{sign} \circ (\eta(x) - c)] \\
\iff & \quad \mathbf{P}(\text{sign} \circ (\hat{\eta}_S(x) - c) = 1 \mid y = 1) \xrightarrow{P} \mathbf{P}(\text{sign} \circ (\eta(x) - c) = 1 \mid y = 1) \text{ and} \\
& \quad \mathbf{P}(\text{sign} \circ (\hat{\eta}_S(x) - c) = -1 \mid y = -1) \xrightarrow{P} \mathbf{P}(\text{sign} \circ (\eta(x) - c) = -1 \mid y = -1) \\
\iff & \quad \mathbf{P}(\hat{\eta}_S(x) > c \mid y = 1) \xrightarrow{P} \mathbf{P}(\eta(x) > c \mid y = 1) \text{ and} \\
& \quad \mathbf{P}(\hat{\eta}_S(x) \leq c \mid y = -1) \xrightarrow{P} \mathbf{P}(\eta(x) \leq c \mid y = -1) \\
\iff & \quad \mathbf{P}(\hat{\eta}_S(x) \leq c \mid y = 1) \xrightarrow{P} \mathbf{P}(\eta(x) \leq c \mid y = 1) \text{ and} \\
& \quad \mathbf{P}(\hat{\eta}_S(x) \leq c \mid y = -1) \xrightarrow{P} \mathbf{P}(\eta(x) \leq c \mid y = -1) \\
\iff & \quad \mathbf{P}_{x|y=1}(\hat{\eta}_S(x) \leq c) \xrightarrow{P} \mathbf{P}_{x|y=1}(\eta(x) \leq c) \text{ and} \\
& \quad \mathbf{P}_{x|y=-1}(\hat{\eta}_S(x) \leq c) \xrightarrow{P} \mathbf{P}_{x|y=-1}(\eta(x) \leq c). \tag{3}
\end{aligned}$$

We now analyze the L_1 -consistency guarantee assumed in the statement of Lemma 2, namely $\mathbf{E}_x[|\hat{\eta}_S(x) - \eta(x)|] \xrightarrow{P} 0$. We begin by expanding this term.

$$\begin{aligned}
\mathbf{E}_x[|\hat{\eta}_S(x) - \eta(x)|] &= p \cdot \mathbf{E}_x[|\hat{\eta}_S(x) - \eta(x)| \mid y = 1] + (1 - p) \cdot \mathbf{E}_x[|\hat{\eta}_S(x) - \eta(x)| \mid y = -1] \\
&= p \cdot \mathbf{E}_{x|y=1}[|\hat{\eta}_S(x) - \eta(x)|] + (1 - p) \cdot \mathbf{E}_{x|y=-1}[|\hat{\eta}_S(x) - \eta(x)|].
\end{aligned}$$

Using the above expansion and the given guarantee on $\hat{\eta}_S$ (along with $p \in (0, 1)$), we obtain $\mathbf{E}_{x|y=1}[|\hat{\eta}_S(x) - \eta(x)|] \xrightarrow{P} 0$ and $\mathbf{E}_{x|y=-1}[|\hat{\eta}_S(x) - \eta(x)|] \xrightarrow{P} 0$ ⁷. Applying Markov inequality for the random variable $|\hat{\eta}_S(x) - \eta(x)|$ for a fixed S , we have for any $\epsilon_1 > 0$,

$$\begin{aligned}
\mathbf{P}_{x|y=1}(|\hat{\eta}_S(x) - \eta(x)| \geq \epsilon_1) &\leq \frac{\mathbf{E}_{x|y=1}[|\hat{\eta}_S(x) - \eta(x)|]}{\epsilon_1} \\
\text{and } \mathbf{P}_{x|y=-1}(|\hat{\eta}_S(x) - \eta(x)| \geq \epsilon_1) &\leq \frac{\mathbf{E}_{x|y=-1}[|\hat{\eta}_S(x) - \eta(x)|]}{\epsilon_1},
\end{aligned}$$

which in turn yields for a fixed $\epsilon_1 > 0$,

$$\mathbf{P}_{x|y=1}(|\hat{\eta}_S(x) - \eta(x)| \geq \epsilon_1) \xrightarrow{P} 0; \tag{4}$$

$$\mathbf{P}_{x|y=-1}(|\hat{\eta}_S(x) - \eta(x)| \geq \epsilon_1) \xrightarrow{P} 0, \tag{5}$$

where recall that the convergence in probability is w.r.t. to a random draw of S according to D^n .

⁷Here, we make use of the fact that for any two sequences of non-negative random variables X_n and Y_n for which $X_n + Y_n \xrightarrow{P} 0$, we have $X_n \xrightarrow{P} 0$ and $Y_n \xrightarrow{P} 0$.

In the rest of the proof, we shall make use of (a) the fact that Eq. (4) and (5) hold for arbitrarily small values of ϵ_1 and (b) our assumption that $\mathbf{P}(\eta(x) \geq c | y = 1)$ and $\mathbf{P}(\eta(x) \geq c | y = -1)$ are continuous at $c \in (0, 1)$ to establish the desired result. We start proving the result w.r.t. $x|y = 1$. For a fixed S and a fixed $\epsilon_2 > 0$, we have

$$\begin{aligned} & \mathbf{P}_{x|y=1}(\widehat{\eta}_S(x) \leq c) \\ &= \mathbf{P}_{x|y=1}(\widehat{\eta}_S(x) \leq c, \eta(x) \leq c + \epsilon_2) + \mathbf{P}_{x|y=1}(\widehat{\eta}_S(x) \leq c, \eta(x) > c + \epsilon_2) \\ &\leq \mathbf{P}_{x|y=1}(\eta(x) \leq c + \epsilon_2) + \mathbf{P}_{x|y=1}(|\widehat{\eta}_S(x) - \eta(x)| \geq \epsilon_2), \end{aligned} \quad (6)$$

and

$$\begin{aligned} & \mathbf{P}_{x|y=1}(\eta(x) \leq c - \epsilon_2) \\ &= \mathbf{P}_{x|y=1}(\widehat{\eta}_S(x) \leq c, \eta(x) \leq c - \epsilon_2) + \mathbf{P}_{x|y=1}(\widehat{\eta}_S(x) > c, \eta(x) \leq c - \epsilon_2) \\ &\leq \mathbf{P}_{x|y=1}(\widehat{\eta}_S(x) \leq c) + \mathbf{P}_{x|y=1}(|\widehat{\eta}_S(x) - \eta(x)| \geq \epsilon_2). \end{aligned} \quad (7)$$

Consequently from Eq. (6) and (7), we get

$$\begin{aligned} & \mathbf{P}_{x|y=1}(\eta(x) \leq c - \epsilon_2) - \mathbf{P}_{x|y=1}(|\widehat{\eta}_S(x) - \eta(x)| \geq \epsilon_2) \leq \mathbf{P}_{x|y=1}(\widehat{\eta}_S(x) \leq c) \\ & \text{and } \mathbf{P}_{x|y=1}(\widehat{\eta}_S(x) \leq c) \leq \mathbf{P}_{x|y=1}(\eta(x) \leq c + \epsilon_2) + \mathbf{P}_{x|y=1}(|\widehat{\eta}_S(x) - \eta(x)| \geq \epsilon_2). \end{aligned}$$

Subtracting the term $\mathbf{P}_{x|y=1}(\eta(x) \leq c)$ from both sides in each of the above inequalities and combining the resulting inequalities then gives us

$$\begin{aligned} & |\mathbf{P}_{x|y=1}(\widehat{\eta}_S(x) \leq c) - \mathbf{P}_{x|y=1}(\eta(x) \leq c)| \leq \\ & \max \left\{ \underbrace{\mathbf{P}_{x|y=1}(|\widehat{\eta}_S(x) - \eta(x)| \geq \epsilon_2) + \mathbf{P}_{x|y=1}(\eta(x) \leq c + \epsilon_2) - \mathbf{P}_{x|y=1}(\eta(x) \leq c)}_{\text{term}_1}, \right. \\ & \left. \underbrace{\mathbf{P}_{x|y=1}(|\widehat{\eta}_S(x) - \eta(x)| \geq \epsilon_2) - \mathbf{P}_{x|y=1}(\eta(x) \leq c - \epsilon_2) + \mathbf{P}_{x|y=1}(\eta(x) \leq c)}_{\text{term}_2} \right\}. \end{aligned} \quad (8)$$

Keeping S fixed, we now allow $\epsilon_2 \rightarrow 0$. In particular, by our assumption that $\mathbf{P}_{x|y=1}(\eta(x) \leq c)$ is continuous at c , for the terms inside the above ‘max’, we have:

$$\begin{aligned} \lim_{\epsilon_2 \rightarrow 0} \text{term}_1 &= \lim_{\epsilon_2 \rightarrow 0} \mathbf{P}_{x|y=1}(|\widehat{\eta}_S(x) - \eta(x)| \geq \epsilon_2); \\ \lim_{\epsilon_2 \rightarrow 0} \text{term}_2 &= \lim_{\epsilon_2 \rightarrow 0} \mathbf{P}_{x|y=1}(|\widehat{\eta}_S(x) - \eta(x)| \geq \epsilon_2). \end{aligned}$$

Thus for a fixed S , the following holds from Eq. (8):

$$0 \leq |\mathbf{P}_{x|y=1}(\widehat{\eta}_S(x) \leq c) - \mathbf{P}_{x|y=1}(\eta(x) \leq c)| \leq \lim_{\epsilon_2 \rightarrow 0} \mathbf{P}_{x|y=1}(|\widehat{\eta}_S(x) - \eta(x)| \geq \epsilon_2).$$

Now, from an application of Eq. (4) (which holds for arbitrarily small ϵ_1), we obtain the following convergence in probability over a random draw of S from D^n :

$$|\mathbf{P}_{x|y=1}(\widehat{\eta}_S(x) \leq c) - \mathbf{P}_{x|y=1}(\eta(x) \leq c)| \xrightarrow{P} 0,$$

which in turn, implies

$$\mathbf{P}_{x|y=1}(\widehat{\eta}_S(x) \leq c) \xrightarrow{P} \mathbf{P}_{x|y=1}(\eta(x) \leq c).$$

This is the desired relation w.r.t $x|y = 1$ (as seen in Eq. (3)). The desired result w.r.t. $x|y = -1$ follows likewise. \square

B.2 Complete proof for Lemma 4

Proof. Define for $i, j \in \{-1, 1\}$:

$$\widehat{p}_{i,j,n}[h] = \sum_{k=1}^n \mathbf{1}(y_k = i, h(x_k) = j) / n \quad \text{and} \quad p_{ij}[h] = \mathbf{E}_D[\mathbf{1}(y = i, h(x) = j)].$$

For a fixed $h \in \mathcal{T}_f$, by the weak law of large numbers (WLLN), we have $\forall i, j$:

$$\widehat{p}_{i,j,n}[h] \xrightarrow{P} p_{i,j}[h],$$

where the convergence in probability is over draw of $S \sim D^n$. Also, $\widehat{p}_S \xrightarrow{P} p$ (again by WLLN).

Given that

$$\widehat{\text{TPR}}_S[h] = \frac{1}{\widehat{p}_S} \widehat{p}_{1,1,n}[h] \quad \text{and} \quad \widehat{\text{TNR}}_S[h] = \frac{1}{(1 - \widehat{p}_S)} \widehat{p}_{-1,-1,n}[h],$$

we thus have

$$\widehat{\text{TPR}}_S[h] \xrightarrow{P} \frac{p_{1,1}[h]}{p} = \text{TPR}_D[h] \quad \text{and} \quad \widehat{\text{TNR}}_S[h] \xrightarrow{P} \frac{p_{-1,-1}[h]}{1-p} = \text{TNR}_D[h].$$

In turn, by continuity of Ψ , we obtain

$$\widehat{\mathcal{P}}_S^\Psi[h] \xrightarrow{P} \mathcal{P}_D^\Psi[h].$$

□

B.3 Complete proof for Lemma 5

Proof. Recall that any fixed $h \in \mathcal{T}_f$ is of the form $\text{sign} \circ (f(x) - c)$ for some constant $c \in (0, 1)$. Let $p_{i,j}[h]$ and $\widehat{p}_{i,j,n}[h]$ be as defined in the proof of Lemma 4 (Section B.2). Since $\text{VC-dimension}(\mathcal{T}_f) = 1$, by standard VC-dimension based uniform convergence arguments, we can argue that for all $i, j \in \{\pm 1\}$, given any $\epsilon' > 0$,

$$\mathbf{P}_{S \sim D^n} \left(\bigcup_{h \in \mathcal{T}_f} |\widehat{p}_{i,j,n}[h] - p_{i,j}[h]| \geq \epsilon' \right) \rightarrow 0.$$

We also have $\widehat{p}_S \xrightarrow{P} p$ (by WLLN).

Now, as in the proof of Lemma 4, observing that TPR and TNR are continuous functions of the above quantities, it can be shown using an appropriate choice of $\epsilon' > 0$ in the above expressions, and by an application of union bound, that for any given $\epsilon > 0$,

$$\begin{aligned} \mathbf{P}_{S \sim D^n} \left(\bigcup_{h \in \mathcal{T}_f} \left\{ |\text{TPR}_D[h] - \widehat{\text{TPR}}_S[h]| \geq \epsilon \right\} \right) &\rightarrow 0 \\ \text{and} \quad \mathbf{P}_{S \sim D^n} \left(\bigcup_{h \in \mathcal{T}_f} \left\{ |\text{TNR}_D[h] - \widehat{\text{TNR}}_S[h]| \geq \epsilon \right\} \right) &\rightarrow 0. \end{aligned}$$

Once again, by continuity of Ψ , we have:

$$\mathbf{P}_{S \sim D^n} \left(\bigcup_{h \in \mathcal{T}_f} \left\{ |\mathcal{P}_D^\Psi[h] - \widehat{\mathcal{P}}_S^\Psi[h]| \geq \epsilon \right\} \right) \rightarrow 0 \text{ as } n \rightarrow \infty,$$

as desired. □

B.4 Complete proof for Lemma 9

Proof. Our proof is similar to that of Theorem 4 in [15]. Recall $\mathcal{T}_\eta = \{\text{sign} \circ (\eta - t) \mid t \in [0, 1]\}$ and let $h^* = \sup_{h \in \mathcal{T}_\eta} \mathcal{P}_D^{\text{G-TP/PR}}[h]$ (the existence of this classifier is guaranteed by Assumption A).

We shall show that for any $h \notin \mathcal{T}_\eta$, $\exists \tilde{h} \in \mathcal{T}_\eta$ such that $\mathcal{P}_D^{\text{G-TP/PR}}[\tilde{h}] \geq \mathcal{P}_D^{\text{G-TP/PR}}[h]$, thus giving us $\mathcal{P}_D^{\text{G-TP/PR}}[h^*] \geq \mathcal{P}_D^{\text{G-TP/PR}}[\tilde{h}] \geq \mathcal{P}_D^{\text{G-TP/PR}}[h]$; this would imply that the optimal predictor for G-TP/PR is indeed of the desired threshold form.

For any $h \notin \mathcal{T}_\eta$, upon arranging all instances $x \in \mathcal{X}$ in non-increasing order of η , we can find disjoint subsets $A, B, C \subseteq \mathcal{X}$, with $\sup_{x \in A} \eta(x) \leq \inf_{x \in B} \eta(x) \leq \sup_{x \in B} \eta(x) \leq \inf_{x \in C} \eta(x)$, such that: $A \cup C = \{x \in \mathcal{X} \mid h(x) = 1\}$ and $B \subseteq \{x \in \mathcal{X} \mid h(x) = -1\}$. We now define two new classifiers:

$$h_A(x) = \begin{cases} -1 & \text{if } x \in A \\ h(x) & \text{o/w} \end{cases} \quad \text{and} \quad h_B(x) = \begin{cases} 1 & \text{if } x \in B \\ h(x) & \text{o/w} \end{cases}.$$

We now claim that one of these newly defined classifiers must be at least as good as h w.r.t. G-TP/PR (this claim is proved below).

Claim. Either $\mathcal{P}_D^{\text{G-TP/PR}}[h_A] \geq \mathcal{P}_D^{\text{G-TP/PR}}[h]$ or $\mathcal{P}_D^{\text{G-TP/PR}}[h_B] \geq \mathcal{P}_D^{\text{G-TP/PR}}[h]$.

According to the above claim, any classifier that is not of the form $\text{sign} \circ (\eta(x) - c)$ is only as good as one of h_B or h_A w.r.t. G-TP/PR. We could now imagine one of h_A or h_B as the new h and make repeated use of the above exchange argument to eventually arrive at a classifier \tilde{h} in \mathcal{T}_η with $\mathcal{P}_D^{\text{G-TP/PR}}[\tilde{h}] \geq \mathcal{P}_D^{\text{G-TP/PR}}[h]$, as desired.

It remains to be shown that the above claim is true.

Proof of Claim. Let us assume the contrary, that $\mathcal{P}_D^{\text{G-TP/PR}}[h] > \mathcal{P}_D^{\text{G-TP/PR}}[h_A]$ and $\mathcal{P}_D^{\text{G-TP/PR}}[h] > \mathcal{P}_D^{\text{G-TP/PR}}[h_B]$, and arrive at a contradiction. Let $a = \mathbf{P}(x \in A)$, $b = \mathbf{P}(x \in B)$ and $c = \mathbf{P}(x \in C)$, and assume without loss of generality that $a, b > 0$. Let $\alpha = \mathbf{E}_x[\eta(x) \mid x \in A]$, $\beta = \mathbf{E}_x[\eta(x) \mid x \in B]$ and $\gamma = \mathbf{E}_x[\eta(x) \mid x \in C]$. It is clearly seen that $0 \leq \alpha \leq \beta \leq \gamma$. With the above definitions, we have, $\text{TPR}_D(h) = \frac{a\alpha + c\gamma}{p}$ and $\text{Prec}_D(h) = \frac{a\alpha + c\gamma}{a+c}$, and in turn, $\mathcal{P}_D^{\text{G-TP/PR}}(h) = \sqrt{\frac{(a\alpha + c\gamma)^2}{p(a+c)}}$, while $\mathcal{P}_D^{\text{G-TP/PR}}(h_B) = \sqrt{\frac{(a\alpha + b\beta + c\gamma)^2}{p(a+b+c)}}$ and $\mathcal{P}_D^{\text{G-TP/PR}}(h_A) = \sqrt{\frac{(c\gamma)^2}{p(c)}}$.

By our contradiction hypothesis,

$$\sqrt{\frac{(a\alpha + c\gamma)^2}{p(a+c)}} > \sqrt{\frac{(a\alpha + b\beta + c\gamma)^2}{p(a+b+c)}} \quad \text{and} \quad \sqrt{\frac{(a\alpha + c\gamma)^2}{p(a+c)}} > \sqrt{\frac{(c\gamma)^2}{p(c)}},$$

which implies

$$(a+b+c)(a\alpha + c\gamma)^2 > (a+c)(a\alpha + b\beta + c\gamma)^2 \tag{9}$$

$$\text{and } c(a\alpha + c\gamma)^2 > (a+c)(c\gamma)^2. \tag{10}$$

Now, from Eq. (10), we have

$$c(a^2\alpha^2 + c^2\gamma^2 + 2a\alpha c\gamma) > ac^2\gamma^2 + c^3\gamma^2 \quad \text{or} \quad aca^2 + 2c^2\alpha\gamma > c^2\gamma^2, \tag{11}$$

Next, from Eq. (9), we have

$$(a+b+c)(a^2\alpha^2 + c^2\gamma^2 + 2a\alpha c\gamma) > (a+c)(a^2\alpha^2 + b^2\beta^2 + c^2\gamma^2 + 2ab\alpha\beta + 2bc\beta\gamma + 2aca\gamma),$$

which can be simplified as

$$b(a^2\alpha^2 + c^2\gamma^2 + 2a\alpha c\gamma) > (a+c)(b^2\beta^2 + 2ab\alpha\beta + 2bc\beta\gamma).$$

Using the upper bound for the term $c^2\gamma^2$ from Eq. (11) in the above inequality, we get

$$\begin{aligned} b(a^2\alpha^2 + aca^2 + 2c^2\alpha\gamma + 2aca\gamma) &> (a+c)(b^2\beta^2 + 2ab\alpha\beta + 2bc\beta\gamma) \\ \implies b(a+c)(a\alpha^2 + 2c\alpha\gamma) &> (a+c)(b^2\beta^2 + 2ab\alpha\beta + 2bc\beta\gamma) \\ \implies a\alpha^2 + 2c\alpha\gamma &> b\beta^2 + 2a\alpha\beta + 2c\beta\gamma. \end{aligned}$$

Using $\beta \geq \alpha$, we can now lower bound the right hand side in the above inequality to get

$$a\alpha^2 + 2c\alpha\gamma > b\beta^2 + 2a\alpha^2 + 2c\alpha\gamma \implies 0 > a\alpha^2 + b\beta^2,$$

which is a contradiction since $a, b > 0$ and $\alpha, \beta \geq 0$. This proves the claim. \square

B.5 Complete proof for Lemma 11

Proof. Recall $\mathcal{T}_\eta = \{\text{sign} \circ (\eta - t) \mid t \in [0, 1]\}$ and let $h^* = \sup_{h \in \mathcal{T}_\eta} \mathcal{P}_D^\Psi[h]$ (the existence of this classifier is guaranteed by Assumption B). We shall now use an exchange argument (that makes use of Assumption B) to show that for any $h \notin \mathcal{T}_\eta$, $\exists \tilde{h} \in \mathcal{T}_\eta$ such that $\mathcal{P}_D^\Psi[\tilde{h}] \geq \mathcal{P}_D^\Psi[h]$, thus implying $\mathcal{P}_D^\Psi[h^*] \geq \mathcal{P}_D^\Psi[\tilde{h}] \geq \mathcal{P}_D^\Psi[h]$; this would imply that the optimal predictor for \mathcal{P}^Ψ is indeed of the

desired threshold form. In particular, we shall show that $\text{TPR}_D[\tilde{h}] \geq \text{TPR}_D[h]$ and $\text{TNR}_D[\tilde{h}] \geq \text{TNR}_D[h]$, which by the monotonicity assumption on Ψ yields $\mathcal{P}_D^\Psi[\tilde{h}] \geq \mathcal{P}_D^\Psi[h]$.

For any $h \notin \mathcal{T}_\eta$, upon arranging all instances $x \in \mathcal{X}$ in non-increasing order of η , we can find disjoint subsets $A \subseteq \{x \in \mathcal{X} \mid h(x) = 1\}$ and $B \subseteq \{x \in \mathcal{X} \mid h(x) = -1\}$ such that $\sup_{x \in A} \eta(x) \leq \inf_{x \in B} \eta(x)$. Let $a = \mathbf{P}(x \in A)$ and $b = \mathbf{P}(x \in B)$; assume without loss of generality, $a, b > 0$.

Let us consider the case where $a \geq b$; here we choose a set $A' \subseteq A$ with $\mathbf{P}(x \in A') = b$ (such a set exists by the continuity assumption in Assumption B), and define a classifier h' as

$$h'(x) = \begin{cases} 1 & \text{if } x \in B \\ -1 & \text{if } x \in A' \\ h(x) & \text{o/w} \end{cases}.$$

We shall now show that $\text{TPR}_D[h'] \geq \text{TPR}_D[h]$ and $\text{TNR}_D[h'] \geq \text{TNR}_D[h]$. In particular,

$$\begin{aligned} & \text{TPR}_D[h'] - \text{TPR}_D[h] \\ &= \mathbf{P}(h'(x) = 1 \mid y = 1) - \mathbf{P}(h(x) = 1 \mid y = 1) \\ &= \frac{1}{p} \mathbf{E}_x[\eta(x) \mathbf{1}(h'(x) = 1)] - \frac{1}{p} \mathbf{E}_x[\eta(x) \mathbf{1}(h(x) = 1)] \\ &= \frac{1}{p} \left[\mathbf{E}_x[\eta(x) \mathbf{1}(h'(x) = 1, h(x) = -1)] - \mathbf{E}_x[\eta(x) \mathbf{1}(h(x) = 1, h'(x) = -1)] \right] \\ &= \frac{1}{p} \left[\mathbf{E}_x[\eta(x) \mathbf{1}(x \in B)] - \mathbf{E}_x[\eta(x) \mathbf{1}(x \in A')] \right] \quad (\text{by definition of } h') \\ &\geq \frac{1}{p} \left[\left(\inf_{x \in B} \eta(x) \right) \mathbf{P}(x \in B) - \left(\sup_{x \in A'} \eta(x) \right) \mathbf{P}(x \in A') \right] \\ &= \frac{b}{p} \left[\inf_{x \in B} \eta(x) - \sup_{x \in A'} \eta(x) \right] \quad (\text{using } \mathbf{P}(x \in B) = \mathbf{P}(x \in A') = b) \\ &\geq \frac{b}{p} \left[\inf_{x \in B} \eta(x) - \sup_{x \in A} \eta(x) \right] \quad (\text{using } A' \subseteq A) \\ &\geq 0, \end{aligned}$$

where the last step follows from the definition of sets A and B ; in a similar manner, one can show that $\text{TNR}_D[h'] - \text{TNR}_D[h] \geq 0$.

For the case when $a < b$, we choose a set $B' \subset B$ with $\mathbf{P}(x \in B') = a$, and define h' as

$$h'(x) = \begin{cases} 1 & \text{if } x \in B' \\ -1 & \text{if } x \in A \\ h(x) & \text{o/w} \end{cases}.$$

Similar to the previous case, one can show that $\text{TPR}_D[h'] \geq \text{TPR}_D[h]$ and $\text{TNR}_D[h'] \geq \text{TNR}_D[h]$.

In both these cases, we have by monotonicity of Ψ that $\mathcal{P}_D^\Psi[h'] \geq \mathcal{P}_D^\Psi[h]$. Note that unless $a = b$, $h' \notin \mathcal{T}_\eta$. Hence, when $a \neq b$, we can view h' as the new h , and apply the above exchange argument repeatedly to eventually arrive at $\tilde{h} \in \mathcal{T}_\eta$ with $\mathcal{P}_D^\Psi[\tilde{h}] \geq \mathcal{P}_D^\Psi[h]$, as desired. \square

C Example Distribution Where the Optimal Classifier for G-mean, H-mean and Q-mean is Not Threshold-based

We now present an example of a distribution under which the optimal classifier for the G-Mean, H-Mean and Q-Mean performance measures (see Table 1) is not of the requisite thresholded form, i.e. not of the form $\text{sign} \circ (\eta(x) - c)$ for any $c \in (0, 1)$.

Let $\mathcal{X} = \{x_1, x_2, x_3\}$. For some a constant $\epsilon \in (0, 1/2)$, consider the following distribution D over $\mathcal{X} \times \{\pm 1\}$:

	$\mathbf{P}(x)$	$\eta(x) = \mathbf{P}(y = 1 x)$
x_1	0.25	$1/2 - \epsilon$
x_2	0.5	$1/2$
x_3	0.25	$1/2 + \epsilon$

Table 4: Example distribution D over $\mathcal{X} \times \{\pm 1\}$, where the optimal classifier for the G-mean, H-mean and Q-mean performance measures is not threshold-based.

Consider the following binary classifiers defined on \mathcal{X} :

$$\begin{aligned} \tilde{h}_0(x) &= \begin{cases} 1 & \text{if } x \in \{x_1, x_2, x_3\} \\ -1 & \text{o/w} \end{cases} \\ \tilde{h}_1(x) &= \begin{cases} 1 & \text{if } x \in \{x_1, x_2\} \\ -1 & \text{o/w} \end{cases} \\ \tilde{h}_2(x) &= \begin{cases} 1 & \text{if } x \in \{x_1\} \\ -1 & \text{o/w} \end{cases} \\ \tilde{h}_3(x) &= \begin{cases} -1 & \text{if } x \in \{x_1, x_2, x_3\} \\ 1 & \text{o/w} \end{cases} \\ h_4(x) &= \begin{cases} 1 & \text{if } x \in \{x_2\} \\ -1 & \text{o/w} \end{cases}, \end{aligned}$$

where the first four classifiers constitute the set of all classifiers on \mathcal{X} of the form $\text{sign} \circ (\eta - c)$ for $c \in (0, 1)$ (indicated by a ‘ \sim ’), while the last one is not of a thresholded form. We next list out in Table 5 the values of the G-mean, H-mean and Q-mean performance measures for these classifiers. It can be seen that for distributions defined using a small value of $\epsilon \in (0, 0.5)$, for each of G-Mean, H-Mean and Q-Mean, the classifier h_4 offers a higher performance measure value than any of the threshold-based classifiers. Clearly, threshold-based classifiers are not optimal under distributions of the above form with small values of ϵ .

	TPR	TNR	G-Mean $\sqrt{\text{TPR} \cdot \text{TNR}}$	H-Mean $2/(\frac{1}{\text{TPR}} + \frac{1}{\text{TNR}})$	Q-Mean $1 - ((1 - \text{TPR})^2 + (1 - \text{TNR})^2)/2$
\tilde{h}_0	1	0	0	0	$1/2$
\tilde{h}_1	$3/4 + \epsilon/2$	$1/4 + \epsilon/2$	$\sqrt{3}/4 + O(\sqrt{\epsilon})$	$3/8 + O(\epsilon)$	$11/16 + O(\epsilon^2)$
\tilde{h}_2	$1/4 + \epsilon/2$	$3/4 + \epsilon/2$	$\sqrt{3}/4 + O(\sqrt{\epsilon})$	$3/8 + O(\epsilon)$	$11/16 + O(\epsilon^2)$
\tilde{h}_3	0	1	0	0	$1/2$
h_4	$1/2$	$1/2$	$1/2$	$1/2$	$3/4$

Table 5: Performance measures G-mean, H-mean and Q-mean evaluated for classifiers $\tilde{h}_0, \tilde{h}_1, \tilde{h}_2, \tilde{h}_3$ and h_4 under the example distribution D in Table 4. Here $\epsilon \in (0, 0.5)$. For small values of ϵ , classifier h_4 offers the best value w.r.t. all measures (highlighted in bold).