Published in final edited form as:

IEEE Trans Signal Process. 2016 February 1; 64(3): 580–591. doi:10.1109/TSP.2015.2477805.

Empirically Estimable Classification Bounds Based on a Nonparametric Divergence Measure

Visar Berisha, Alan Wisler, Alfred O. Hero, and Andreas Spanias

Abstract

Information divergence functions play a critical role in statistics and information theory. In this paper we show that a non-parametric $\mathcal{f}$ -divergence measure can be used to provide improved bounds on the minimum binary classification probability of error for the case when the training and test data are drawn from the same distribution *and* for the case where there exists some mismatch between training and test distributions. We confirm the theoretical results by designing feature selection algorithms using the criteria from these bounds and by evaluating the algorithms on a series of pathological speech classification tasks.

Index Terms

Bayes error rate; classification; divergence measures; non-parametric divergence estimator; domain adaptation

I. Introduction

A number of information-theoretic divergence measures between probability distance functions have been introduced and analyzed in the literature [1]–[5]. They have been extensively used in many signal processing applications involving classification [6], segmentation [7], source separation [8], clustering [9], and other domains.

Among the different divergence functions, the family of $\mathcal{f}$ -divergences or Ali-Silvey distances is perhaps the most widely used in signal processing [10]. This family includes the total variation distance, the Bhattacharya distance [1], the Kullback-Leibler divergence [2], and more generally, the Chernoff α -divergence [3], [4]. Because there exists an indirect relationship between the class of $\mathcal{f}$ -divergences and the minimum achievable error in classification problems [11], this family of divergence measures is particularly useful for this setting. Consider the problem of classifying a multi-dimensional feature vector, $\mathbf{x}$, into one of two classes, $\{0, 1\}$. The conditional distributions are given by $f_0(\mathbf{x})$ and $f_1(\mathbf{x})$ and the prior probabilities are given by p and q , respectively. The classifier that assigns a vector $\mathbf{x}$ to the class with the highest posterior is called Bayes classifier, and the error rate of this classifier is given by:

$$\epsilon^{\text{Bayes}} = \int_{pf_0(\mathbf{x}) \leq qf_1(\mathbf{x})} pf_0(\mathbf{x}) d\mathbf{x} + \int_{qf_1(\mathbf{x}) \leq pf_0(\mathbf{x})} qf_1(\mathbf{x}) d\mathbf{x} \quad (1)$$

This is the minimum classification error rate, or the Bayes error rate (BER), that can be achieved by any classifier. Computing the BER requires evaluating the multi-dimensional integral over regions that can only be determined if one has perfect knowledge of the data distribution. As an alternative to computing the integral, a number of attempts have been made to bound this error using estimable measures of distance between probability functions [3], [12]–[14].

In this paper, we derive a new bound on classification error that is based on a nonparametric probability distance measure that belongs to the family of f -divergences. In the context of binary classification, this new measure has a number of appealing properties: (1) there exists an asymptotically consistent estimator of the divergence measure that does not require density estimates of the two distributions; (2) we show that there exists a local relationship between this new divergence measure and the Chernoff α -divergence; (3) we derive tighter bounds on the BER than those based on the Bhattacharya distance and derive empirical estimates of these bounds using data from the two distributions; (4) we derive bounds on the minimum achievable error rate for the case where training and test data in the classification problem come from different distributions.

A. Related work

There are three lines of research that are related to the work presented in this paper: information theoretic bounds on the Bayes error rate (and related quantities); bounds from the machine learning literature for the scenario where training and test data come from different distributions; and recent work on empirical estimates of the KL divergence.

The total variation (TV) distance is closely related to the Bayes error rate [12]. A number of bounds exist in the literature relating the KL divergence and the TV distance. The well-known Pinsker inequality provides a bound on the total variation distance in terms of the KL divergence [15]. Sharpened inequalities that bound the KL divergence in terms of a polynomial function of the TV distance were derived in [16]. One drawback of the Pinsker-type inequalities is that they become uninformative for completely separable distributions where the KL divergence goes to ∞ (since the TV distance is upper bounded). Vajda's refinement to these bounds addresses this issue [17].

For classification problems, the well-known upper bound on the probability of error based on the Chernoff α -divergence has been used in a number of statistical learning applications [3]. The tightest bound is determined by finding the value of α that minimizes the upper bound. The Bhattacharya (BC) divergence, a special case of the Chernoff α -divergence for $\alpha = \frac{1}{2}$, upper and lower bounds the BER [1], [12]. The BC bounds are often used as motivation for algorithms in the statistical learning literature because these bounds have closed form expressions for many commonly used distributions. In addition, for small differences between the two classes, it has been shown that, in the class of Chernoff α -divergence measures, $\alpha = \frac{1}{2}$ (the BC divergence) results in the tightest upper bound on the probability of error [4].

Beyond the bounds on the BER based on the divergence measures, a number of other bounds exist based on different functionals of the distributions. In [13], the authors derive a new

functional based on a Gaussian-Weighted sinusoid that yields tighter bounds on the BER than other popular approaches. Avi-Itzhak proposes arbitrarily tight bounds on the BER in [14]. Both of these sets of bounds are tighter than the bounds we derive here; however, these bounds cannot be estimated without at least partial knowledge of the underlying distribution. A strength of the bounds proposed in this paper is that they are empirically estimable without knowing a parametric model for the underlying distribution.

In addition to work on bounding the Bayes error rate, recently there have been a number of attempts to bound the error rate in classification problems for the case where the training data and test data are drawn from different distributions (an area known as domain-adaptation or transfer learning in the machine learning literature). In [18], [19], Ben-David *et al.* relate the expected error on the test data to the expected error on the training data, for the case when no labeled test data is available. In [20], the authors derive new bounds for the case where a small subset of labeled data from the test distribution is available. In [21], Mansour *et al.* generalize these bounds to the regression problem. In [22], the authors present a new theoretical analysis of the multi-source domain adaptation problem based on the α -divergence. In contrast to these models, we propose a general non-parametric bound that can be estimated without assuming an underlying model for the data and without restrictions on the hypothesis class.

While previous bounds have proven useful in a number of applications, a drawback shared by most divergence functions (and corresponding bounds) is that they require some knowledge of the underlying distribution for their estimation. For some of the more popular divergence measures, closed form solutions are available for different distribution types [23]. More recently, a number of non-parametric methods have been introduced to estimate information theoretic quantities. Graph-based non-parametric estimators were introduced in [24]. Plug-in estimates of existing divergence measures that require density estimation have also been proposed [25]. More recently, estimates of the KL divergence that rely on estimates of the likelihood ratio instead of direct density estimation have been proposed [26], [27]. In [28] a minimal spanning tree (MST) based estimator of a different kind of $\mathcal{F}$ -divergence measure was investigated. Unlike other divergences, this $\mathcal{F}$ -divergence can be estimated directly from the data without performing density estimation. This estimator was used in [28] to develop a nonparametric estimator for the Fisher information. Whereas that paper analyzes the utility of the proposed $\mathcal{F}$ -divergence for estimation problems, this work focuses on its importance to binary classification tasks.

The rest of this paper is outlined as follows: In section II, we provide an overview of the divergence measure and its consistent estimator. In section III, we derive bounds on the BER based on this probability distance measure and compare the tightness of the bound with Bhattacharya bound and, more generally, with the bound based on the α -divergence. In section IV, we derive bounds on the classification error rate for the case where the training and the test data come from different distributions. In section V, we provide numerical results that confirm the validity of the bounds and describe two practical algorithms for feature learning that aim to minimize the upper bound on the error rate. Section VI contains concluding remarks and a discussion of future work.

II. A Nonparametric Divergence Measure

For parameters $p \in (0, 1)$ and $q = 1 - p$ consider the following divergence measure between distributions f and g with domain $\mathbf{R}^d$:

$$D_p(f, g) = \frac{1}{4pq} \left[\int \frac{(pf(\mathbf{x}) - qg(\mathbf{x}))^2}{pf(\mathbf{x}) + qg(\mathbf{x})} d\mathbf{x} - (p - q)^2 \right] \quad (2)$$

The divergence in (2), first introduced in [28], has the remarkable property that it can be estimated directly without estimation or plug-in of the densities f and g based on an extension of the Friedman-Rafsky (FR) multi-variate two sample test statistic [29]. Let us consider sample realizations from f and g , denoted by $\mathbf{X}_f \in \mathbf{R}^{N_f \times d}$, $\mathbf{X}_g \in \mathbf{R}^{N_g \times d}$. The FR test statistic, $\mathcal{C}(\mathbf{X}_f, \mathbf{X}_g)$, is constructed by first generating a Euclidean minimal spanning tree (MST) on the concatenated data set, $\mathbf{X}_f \cup \mathbf{X}_g$, and then counting the number of edges connecting a data point from f to a data point from g . The test assumes a unique MST for $\mathbf{X}_f \cup \mathbf{X}_g$ - therefore all inter point distances between data points must be distinct. However, this assumption is not restrictive since the MST is unique with probability one when f and g are Lebesgue continuous densities. In Theorem 1, we present an estimator that relies on the FR test statistic and asymptotically converges to $D_p(f, g)$. Note that this theorem combines the results of Theorem 1 and equations (3) and (4) in [28]. The proof of this theorem can be found in Appendix A.

Theorem 1—As $N_f \rightarrow \infty$ and $N_g \rightarrow \infty$ in a linked manner such that $\frac{N_f}{N_f + N_g} \rightarrow p$ and $\frac{N_g}{N_f + N_g} \rightarrow q$,

$$1 - \mathcal{C}(\mathbf{X}_f, \mathbf{X}_g) \frac{N_f + N_g}{2N_f N_g} \rightarrow D_p(f, g).$$

almost surely.

In Fig. 1a and 1b we show two numerical examples in order to visualize the results of Theorem 1 - we plot samples from two distributions, $\mathbf{X}_f \sim f(\mathbf{x})$ and $\mathbf{X}_g \sim g(\mathbf{x})$, and evaluate the value of $\mathcal{C}(\mathbf{X}_f, \mathbf{X}_g)$. In Fig. 1a, both data sets are drawn from the same distribution, $f(\mathbf{x}) = g(\mathbf{x}) = \mathcal{N}([0, 0]^T, \mathbf{I})$. In Fig. 1b, we plot data drawn from $f(\mathbf{x}) = \mathcal{N}([- \frac{\sqrt{2}}{2}, - \frac{\sqrt{2}}{2}]^T, \mathbf{I})$ and $g(\mathbf{x}) = \mathcal{N}([\frac{\sqrt{2}}{2}, \frac{\sqrt{2}}{2}]^T, \mathbf{I})$. $\mathbf{I}$ is the identity matrix. For both data sets, an equal number of points are drawn, therefore $N_f = N_g = N$ and $p = q = \frac{1}{2}$. The dotted line in each figure represents the Euclidean MST associated with $\mathbf{X}_f \cup \mathbf{X}_g$. The green lines represent the edges of the MST connecting points from f to points from g , $\mathcal{C}(\mathbf{X}_f, \mathbf{X}_g)$. We can use this to estimate $D_p(f, g)$ using the results of Theorem 1. It is clear from the figures that this value is much smaller for overlapping distributions (Fig. 1a) than for separable distributions (Fig. 1b). Indeed, as Theorem 1 suggests, in the limit, this statistic converges to the integral used in the divergence measure in (2).

In the ensuing sections we outline some important properties of this divergence measure and develop new bounds for classification using this distance function between distributions.

A. Properties of D_p

The divergence measure in (2) exhibits several properties that make it useful for statistical analysis. It is relatively straightforward to show that the following three properties are satisfied.

1. $0 \leq D_p \leq 1$
2. $D_p = 0 \Leftrightarrow f(\mathbf{x}) = g(\mathbf{x})$.
3. $D_p(f, g) = D_q(g, f)$

The lower bound in the first property follows from the fact that when $f = g$ and $p = q$, the minimum value of D_p is 0. To show that the divergence measure is upper bounded by 1, we first note that

$$\int \frac{(pf(\mathbf{x}) - qg(\mathbf{x}))^2}{pf(\mathbf{x}) + qg(\mathbf{x})} d\mathbf{x} = 1 - 4pq A_p(f, g), \quad (3)$$

where

$$A_p(f, g) = \int \frac{f(\mathbf{x})g(\mathbf{x})}{pf(\mathbf{x}) + qg(\mathbf{x})} d\mathbf{x}.$$

The function $A_p(f, g)$ attains its minimum value of 0, when f and g have no overlapping support (since $f(\mathbf{x}) > 0$ and $g(\mathbf{x}) > 0$ for all $\mathbf{x}$); therefore $D_p = \frac{1}{4pq} [1 - (p - q)^2] = 1$. The second property is closely related to the first: the minimum value $D_p = 0$ only when $f = g$ and $p = q$. The third property follows from commutativity.

The divergence measure in (2) belongs to the class of $\mathcal{F}$ -divergences. Every $\mathcal{F}$ -divergence can be expressed as an average of the ratio of two distributions, weighted by some function

$\phi(t): D_\phi(f, g) = \int \phi\left(\frac{f(\mathbf{x})}{g(\mathbf{x})}\right) g(\mathbf{x}) d\mathbf{x}$. For $D_p(f, g)$, the corresponding function $\phi(t)$ is,

$$\phi(t) = \frac{1}{4pq} \left[\frac{(pt - q)^2}{pt + q} - (2p - 1)^2 \right]. \quad (4)$$

Furthermore, $\phi(t)$ is defined for all $t > 0$, is convex - $\phi''(t) = \frac{2pq}{(pt + q)^3} > 0$, and $\phi(1) = 0$. This is consistent with the requirements of the definition of an $\mathcal{F}$ -divergence [10]. Indeed, for the special case of $\alpha = \frac{1}{2}$, the divergence in (2) becomes the symmetric χ^2 $\mathcal{F}$ -divergence in [5] and is similar to the Rukhin $\mathcal{F}$ -divergence in [30].

III. Bounds on Bayes Classification Error

In this section, we show how D_p in (2) can be used to bound the Bayes error rate (BER) for binary classification. Further, we show that, under certain conditions, this bound is tighter than the well-known Bhattacharya bound commonly used in the machine learning literature and can be empirically estimated from data.

Before deriving the error bounds, for notation convenience, we introduce a slightly modified version of the divergence measure in (2),

$$\begin{aligned}\tilde{D}_p(f, g) &= 1 - 4pq \int \frac{f(\mathbf{x})g(\mathbf{x})}{pf(\mathbf{x}) + qg(\mathbf{x})} d\mathbf{x} \\ &= \int \frac{(pf(\mathbf{x}) - qg(\mathbf{x}))^2}{pf(\mathbf{x}) + qg(\mathbf{x})} d\mathbf{x}.\end{aligned}\quad (5)$$

It is easy to see that $D_p = \frac{\tilde{D}_p}{4pq} - \frac{(p-q)^2}{4pq}$ and when $p = q = 0.5$, $D_p = \tilde{D}_p$. While this function no longer satisfies $\tilde{D}_p(f, g) = 0$, for $f = g$, and therefore is no longer a valid divergence measure, it greatly simplifies the notation of the ensuing error bounds. As with D_p , we can estimate this quantity using the FR test statistic since, under the same conditions as those in Theorem 1,

$$1 - 2 \frac{\mathcal{C}(\mathbf{X}_f, \mathbf{X}_g)}{N_f + N_g} \rightarrow \tilde{D}_p(f, g). \quad (6)$$

Given a binary classification problem with binary labels $y \in \{0, 1\}$ and $\mathbf{x}$ drawn from $f_S(\mathbf{x})$, we denote the conditional distributions for both classes as $f_0(\mathbf{x}) = f_S(\mathbf{x}|y=0)$ and $f_1(\mathbf{x}) = f_S(\mathbf{x}|y=1)$. We draw samples from these distributions with probability p and $q = 1 - p$, respectively, and formulate two data matrices denoted by $\mathbf{X}_0 \in \mathbf{R}^{N_0 \times d}$ and $\mathbf{X}_1 \in \mathbf{R}^{N_1 \times d}$. The Bayes error rate associated with this problem is given in (1). In Theorem 2 below, we show that we can bound this error from above and below using the divergence measure introduced in the previous section. The proof of this theorem can be found in Appendix B.

Theorem 2—For two distributions, $f_0(\mathbf{x})$ and $f_1(\mathbf{x})$, with prior probabilities p and q respectively, the Bayes error rate, ϵ^{Bayes} , is bounded above and below as follows:

$$\frac{1}{2} - \frac{1}{2} \sqrt{\tilde{D}_p(f_0, f_1)} \leq \epsilon^{\text{Bayes}} \leq \frac{1}{2} - \frac{1}{2} \tilde{D}_p(f_0, f_1).$$

Combining the results from Theorem 1 with the results of Theorem 2, we see that we can approximate the upper and lower bounds on the BER from the data matrices $\mathbf{X}_0$ and $\mathbf{X}_1$ as

$$\frac{1}{2} - \frac{1}{2} \sqrt{\tilde{D}_p(f_0, f_1)} \approx \frac{1}{2} - \frac{1}{2} \sqrt{1 - 2 \frac{\mathcal{C}(\mathbf{X}_0, \mathbf{X}_1)}{N_0 + N_1}},$$

and

$$\frac{1}{2} - \frac{1}{2} \tilde{D}_p(f_0, f_1) \approx \frac{\mathcal{C}(\mathbf{X}_0, \mathbf{X}_1)}{N_0 + N_1}.$$

The derived bound is tight for the case $p=q=\frac{1}{2}$. For $f_0(\mathbf{x}) = f_1(\mathbf{x})$, the BER is 0.5. Under these conditions, $\tilde{D}_p(\mathbf{x}) = 0$, and both the upper and lower bound in Theorem 2 go to 0.5. For the case where $f_0(\mathbf{x})$ and $f_1(\mathbf{x})$ are completely separable, the BER is 0, $\tilde{D}_p(\mathbf{x}) = 1$, and both the upper and lower bound go to 0.

A. Relationship to the Chernoff Information Bound

Here we compare the tightness of the bounds on the Bayes error rate based on D_p to the bounds based on the Chernoff information function (CIF) [4], defined as

$$I_\alpha(f_0, f_1) = \int p^\alpha f_0^\alpha(\mathbf{x}) q^{1-\alpha} f_1^{1-\alpha}(\mathbf{x}) d\mathbf{x}.$$

In Theorem 3, we derive an important relationship between the affinity measure, $A_p(f_0, f_1)$, and a scaled version of the CIF. The proof of this theorem can be found in Appendix C.

Theorem 3—The affinity measure, $A_p(f_0, f_1)$, is a lower bound for a scaled version of the Chernoff information function:

$$A_p(f_0, f_1) \leq \int f_0^q(\mathbf{x}) f_1^p(\mathbf{x}) d\mathbf{x}.$$

It is important to note that the second term in Theorem 3 is exactly equal to the CIF for $\alpha = p = q = 1/2$. For this special case, the Chernoff bound reduces to the Bhattacharyya (BC) bound, a widely-used bound on the Bayes error in machine learning that has been used to motivate and develop new algorithms [12], [31], [32]. The popularity of the BC bound is mainly due to the fact that closed form expressions for the bound exist for many of the commonly used distributions. Let us define the Bhattacharyya coefficient as:

$$BC(f_0, f_1) = 2 \int \sqrt{p q f_0(\mathbf{x}) f_1(\mathbf{x})} d\mathbf{x}. \quad (7)$$

The well-known Bhattacharyya bound on the BER is given by

$$\frac{1}{2} - \frac{1}{2} \sqrt{1 - BC^2(f, g)} \leq \varepsilon^{\text{Bayes}} \leq \frac{1}{2} BC(f, g). \quad (8)$$

In Theorem 4 below, we show that, for equiprobable classes, the D_p bound provides tighter upper and lower bounds on the BER when compared to the bound based on the BC

coefficient under all separability conditions. The proof of this theorem can be found in Appendix D.

Theorem 4—For $p=q=\frac{1}{2}$, the D_p upper and lower bounds on the Bayes error rate are tighter than the Bhattacharyya bounds:

$$\frac{1}{2} - \frac{1}{2} \sqrt{1 - BC^2(f_0, f_1)} \leq \frac{1}{2} - \frac{1}{2} \sqrt{\tilde{D}_{\frac{1}{2}}(f_0, f_1)} \leq \varepsilon^{\text{Bayes}} \leq \frac{1}{2} - \frac{1}{2} \tilde{D}_{\frac{1}{2}}(f_0, f_1) \leq \frac{1}{2} BC(f_0, f_1).$$

Using asymptotic analysis of the Chernoff exponent, for small differences between the two classes, it was shown that $\alpha=\frac{1}{2}$ results in the tightest bound on the probability of error - this corresponds to the bound in (8) [4]. Using a variant of this analysis, we derive a local representation of the CIF and relate it to the divergence measure proposed here. In particular, if we let

$$\begin{aligned} pf_0(\mathbf{x}) &= \frac{1}{2}(pf_0(\mathbf{x}) + qf_1(\mathbf{x})) + \frac{1}{2}(pf_0(\mathbf{x}) - qf_1(\mathbf{x})) \\ &= f_{\frac{1}{2}}(\mathbf{x})(1 + \frac{1}{2}\Delta_{\mathbf{x}}), \end{aligned}$$

where $f_{\frac{1}{2}}(\mathbf{x}) = \frac{1}{2}(pf_0(\mathbf{x}) + qf_1(\mathbf{x}))$ and $\Delta_{\mathbf{x}} = (pf_0(\mathbf{x}) - qf_1(\mathbf{x}) / f_{\frac{1}{2}}(\mathbf{x}))$. Similarly,

$$qf_1(\mathbf{x}) = f_{\frac{1}{2}}(\mathbf{x})(1 - \frac{1}{2}\Delta_{\mathbf{x}}).$$

As in [4], after a Taylor series expansion around $p^\alpha f_0^\alpha(\mathbf{x})$ and $q^{1-\alpha} f_1^{1-\alpha}(\mathbf{x})$, the Chernoff information function can be expressed as (see proof of Proposition 5 in [4]):

$$\begin{aligned} I_\alpha(f_0, f_1) &= \int f_{\frac{1}{2}}(\mathbf{x}) \left[1 - (2\alpha - 1) \frac{\Delta_{\mathbf{x}}}{2} - \alpha(1 - \alpha) \left(\frac{\Delta_{\mathbf{x}}}{2} \right)^2 + o(\Delta_{\mathbf{x}}^3) \right] d\mathbf{x} \\ &= \int f_{\frac{1}{2}}(\mathbf{x}) d\mathbf{x} - (2\alpha - 1) \int f_{\frac{1}{2}}(\mathbf{x}) \frac{\Delta_{\mathbf{x}}}{2} d\mathbf{x} - \alpha(1 - \alpha) \int f_{\frac{1}{2}}(\mathbf{x}) \left(\frac{\Delta_{\mathbf{x}}}{2} \right)^2 d\mathbf{x} + o(\Delta^2) \\ &= \frac{1}{2} - (2\alpha - 1)(2p - 1)/2 - \frac{\alpha(1 - \alpha)}{2} \int \frac{(pf_0(\mathbf{x}) - qf_1(\mathbf{x}))^2}{pf_0(\mathbf{x}) + qf_1(\mathbf{x})} d\mathbf{x} + o(\Delta^2) \\ &= (p + \alpha) - 2\alpha p - \frac{\alpha(1 - \alpha)}{2} \tilde{D}_p(f_0, f_1) + o(\Delta^2) \end{aligned}$$

The local equivalence of D_p and I_α is not surprising since all f -divergences are locally equivalent (they induce the same Riemann-Fisher metric on the manifold of densities) [10]. This useful property allows us to estimate the CIF for small differences between f_0 and f_1 using the MST procedure in Section II. Further, we can express the BER in terms of the CIF:

$$\varepsilon^{\text{Bayes}} \leq I_\alpha \approx (p + \alpha) - 2\alpha p - \frac{\alpha(1 - \alpha)}{2} \tilde{D}_p(f_0, f_1).$$

For $p=q=\frac{1}{2}$, this bound reduces to $\varepsilon^{\text{Bayes}} \leq \frac{1}{2} - \frac{\alpha(1-\alpha)}{2} \tilde{D}_{\frac{1}{2}}(f_0, f_1)$. This is very similar to the upper bound in Theorem 2, differing only in the scale of the second term. Further, it is easy to see from this that the bound in Theorem 2 is tighter than the Chernoff bound since $\frac{\alpha(1-\alpha)}{2} < \frac{1}{2}$ for all α . This is not surprising since, locally, $\alpha = 0.5$ yields the tightest bounds on the BER [4]. This corresponds to the BC bound in (8) and we have already shown that new bound is tighter than the BC bound in Theorem 4. This analysis further confirms that result.

In addition to providing tighter bounds on the BER we can estimate the new D_p bound without ever explicitly computing density estimates. We provide a numerical example for comparison. We consider two data samples from two classes, each of which comes from a normally distributed bivariate distribution with varying mean and spherical unit variance. The separation in means between the two class distributions is increased incrementally across 150 trials. The two distributions completely overlap initially, and are almost entirely separated by the final trial. In each trial we calculate the BER analytically using (1), as well as the upper and lower bounds introduced in Theorem 2. We calculate the bounds both analytically (through numerical integration) and empirically (using the results from Theorem 1). In order to demonstrate the tightness of this bound we also plot it against the upper and lower Bhattacharyya error bounds for Gaussian data (the closed form expression of the bound for Gaussian data is known) [12]. Figure 2 displays the true BER along with both error bounds as a function of the Euclidean separation between the means of two bivariate normal distributions of unit variance. We see in this plot that the proposed error bounds are noticeably tighter than the Bhattacharyya error bounds and are well correlated with the true BER. Although the analytically calculated D_p bound never crosses the BC bound, the empirically estimated D_p bound crosses the BC bound for small values of the mean separation. This is due to the variance of the estimator. It is important to note that the estimator used here *asymptotically* converges to the D_p divergence; however this result doesn't necessarily extend to finite data. In fact, for any fixed estimator, there exists a distribution for X and y such that the error converges arbitrarily slowly [33].

IV. Bounds on the Domain Adaptation Error

In this section, we consider a cross-domain binary classification problem and show how the D_p distance can be used to bound the error rate in this setting also. Let us define data from two domains, the source (training) and the target (testing) domain and the corresponding labeling functions for each domain $y_S(\mathbf{x}), y_T(\mathbf{x}) \in \{0, 1\}$ that yields the true class label of a given data point $\mathbf{x}$. The source domain, denoted by the pair $(\mathbf{X}_S, y_S)$, represents the data used to train the machine learning algorithm and the data $(\mathbf{X}_T, y_T)$ represents the data the algorithm will encounter once deployed. Let us further define the conditional distributions $f_{S,0}(\mathbf{x}) = f_S(\mathbf{x}|y_S(\mathbf{x}) = 0)$ and $f_{S,1}(\mathbf{x}) = f_S(\mathbf{x}|y_S(\mathbf{x}) = 1)$. The rows of the source and target data are drawn from $f_S(\mathbf{x})$ and $f_T(\mathbf{x})$. The risk, or the probability that the decision, h , disagrees with the true label is defined as

$$\varepsilon_S(h, y_S) = \mathbb{E}_{f_S(\mathbf{x})} [|h(\mathbf{x}) - y_S|], \quad (9)$$

for the source data. It is similarly defined for the target data. In Theorem 5, we identify a relationship between the error rates on the source and target data. The proof of this theorem can be found in Appendix E.

Theorem 5—Given a hypothesis, h , the target error, $\varepsilon_T(h, y_T)$, can be bounded by the error on the source data, $\varepsilon_S(h, y_S)$, the difference between labels, and a distance measure between source and target distributions as follows:

$$\varepsilon_T(h, y_T) \leq \varepsilon_S(h, y_S) + \mathbf{E}_{f_S(x)} [|y_S - y_T|] + 2\sqrt{\tilde{D}_{\frac{1}{2}}(f_S, f_T)}, \quad (10)$$

where $\tilde{D}_{\frac{1}{2}}(f_S, f_T)$ assumes equiprobable data from the source and target distributions.

The bound in Theorem 5 depends on three terms: the error on the source data, the expected difference in the labeling functions across the two domains, and a measure of the distance between source and target distributions (D_p distance). We expect that the selected training algorithm will seek to minimize the first term; the second term characterizes the difference between labeling functions in the source and target domains; the third term is of particular interest to us - it provides a means of bounding the error on the *target* data as a function of the distance between source and target distributions.

In the *covariate shift* scenario, we assume that there exists no difference between labeling functions (e.g. $y_S(\mathbf{x}) = y_T(\mathbf{x})$) and only the distributions between the source and target data change [19]. Under this assumption, the bound in Theorem 5 reduces to

$$\varepsilon_T(h, y_T) \leq \varepsilon_S(h, y_S) + 2\sqrt{\tilde{D}_{\frac{1}{2}}(f_S, f_T)}. \quad (11)$$

Furthermore, if we assume that the decision rule h attains the Bayes error rate, $\varepsilon^{\text{Bayes}}$, on the source domain, we can use the results from Theorem 2 to rewrite the bound in Theorem 5 using only the D_p distance:

$$\varepsilon_T \leq \frac{1}{2} - \frac{1}{2}\tilde{D}_p(f_{S,0}, f_{S,1}) + 2\sqrt{\tilde{D}_{\frac{1}{2}}(f_S, f_T)}. \quad (12)$$

If we denote the training data matrices by $\mathbf{X}_{S,0} \sim f_{S,0}$ and $\mathbf{X}_{S,1} \sim f_{S,1}$, then we can estimate this upper bound using the FR test statistic by

$$\frac{\mathcal{C}(\mathbf{X}_{S,0}, \mathbf{X}_{S,1})}{N_{S,0} + N_{S,1}} + 2\sqrt{1 - 2\frac{\mathcal{C}(\mathbf{X}_S, \mathbf{X}_T)}{N_S + N_T}}. \quad (13)$$

The result shown in (13) represents an upper bound on the target domain error that can be computed without access to any labels in this domain. This bound provides interesting insight on the importance of invariant representations for classification. The target error is bounded by the sum of the affinity between class distributions in the source domain and the square root of the D_p -distance between domains. Because of the square root and the multiplicative factor, it is clear that the second term in (13) is weighted much more heavily. This stresses the importance of invariant representations in classification. In other words, the bound provides a means of quantifying the relative importance of selecting features that are invariant across domains versus features that provide good separation between classes in the source domain.

V. Numerical Results and Practical Algorithms

Here, we describe a number of numerical experiments that evaluate the bounds in a classification setting. In the first experiment, we evaluate the tightness of the bound on the Bayes error rate in higher dimensions by comparing against two other bounds for an example where the Bayes error rate is known in closed form. In the second and third experiments, we develop new criteria for feature selection based on the derived bounds and compare the probability of correct classification against competing alternatives.

A. Bounding the Bayes Error Rate

Consider the two data sets $\mathcal{D}_1$ and $\mathcal{D}_2$ in Table I, each consisting of data from two 8 dimensional Gaussian distributions. In [34] Fukunaga computed the true Bayes error rate analytically for both of these data sets. Here we compare three different bounds on this error for both datasets - the D_p -based bound, the Mahalanobis bound, and the BC bound. We use the closed-form version of the BC and Mahalanobis bound for Gaussian data [34]. Furthermore, we assume perfect knowledge of the parameters for these two bounds (σ and μ). As a result, this is the best possible case for both of these bounds - the data matches the model and no estimation of the parameters is required.

For both data sets $\mathcal{D}_1$ and $\mathcal{D}_2$, we evaluate the D_p -based upper bound between the two distributions using the graph-based method outlined in Section II for three different sample sizes (100 samples, 500 samples, and 1000 samples - 50 Monte Carlo simulations each). We compare the D_p bound (computed from empirical data without assuming any parametric model of the data distribution) with the Bhattacharyya bound and the Mahalanobis bound. For both data sets, the average D_p -based bound is closer to the true error rate, regardless of the sample size. Again, it is important to stress that this is the best case scenario for the competing bounds since there exists a closed form expression for both bounds for Gaussian data and we assume perfect knowledge of the distribution parameters. Regardless, the empirically-estimated D_p bound is still tighter.

B. Feature Selection using D_p -distance

In machine learning, feature selection algorithms are often used to reduce model complexity and prevent over-fitting [35]. In many scenarios, feature selection can actually improve model performance since the reduced dimensionality leads to a much more densely

populated hypothesis space. This prevents the model from learning irrelevant patterns in the training data that aren't pertinent for a given task and will not generalize to new datasets. This problem is exacerbated in domain adaptation problems where the separation in domains makes misleading patterns in the training data especially problematic. We use the bounds defined in Theorems 2 and 5 to develop new feature selection criteria that aim to directly minimize the BER bound. We consider two different scenarios: (1) one where the training data and the test data come from the same distribution and (2) another where the training data and the test data come from different distributions. For both scenarios, we seek to identify the subset of features, Ω , that will minimize the "worst-case" error. For scenario 1, this results in minimizing the upper bound in Theorem 2:

Algorithm 1

Forward selection algorithm using D_P -distance

Input: Feature data from two different classes in the source domain and unlabelled data from the target domain: $\mathbf{X}_{S,0}$, $\mathbf{X}_{S,1}$, $\mathbf{X}_T$, α

Output: Top k features that minimize Φ .

Ω

Define: $\Omega = \emptyset$

$F = 1 \dots M$

$\mathbf{X}_S = \mathbf{X}_{S,0} \cup \mathbf{X}_{S,1}$

for $j \in 1 \dots k$ **do**

$\Phi = \emptyset$

for $F_i \in F \setminus \Omega$ **do**

$$\Phi(F_i) = \frac{\mathcal{C}(\mathbf{X}_{S,0}(\Omega \cup F_i), \mathbf{X}_{S,1}(\Omega \cup F_i))}{N_{S,0} + N_{S,1}} + 2\alpha \sqrt{1 - 2 \frac{\mathcal{C}(\mathbf{X}_S(\Omega \cup F_i), \mathbf{X}_T(\Omega \cup F_i))}{N_S + N_T}}$$

end for

$$\Omega = \Omega \cup \{\underset{F_i}{\operatorname{argmin}} \Phi(F_i)\}$$

end for

$$\Phi(\Omega) = \frac{\mathcal{C}(\mathbf{X}_1(\Omega), \mathbf{X}_2(\Omega))}{N_1 + N_2}, \quad (14)$$

and, for scenario 2, we minimize the DA bound defined in Theorem 5:

$$\Phi(\Omega) = \frac{\mathcal{C}(\mathbf{X}_{S,0}(\Omega), \mathbf{X}_{S,1}(\Omega))}{N_{S,0} + N_{S,1}} + 2 \sqrt{1 - 2 \frac{\mathcal{C}(\mathbf{X}_S(\Omega), \mathbf{X}_T(\Omega))}{N_S + N_T}}. \quad (15)$$

We integrate the optimization criteria into a forward selection search algorithm in Alg. 1. In this algorithm, we use a parameter α to determine whether or not the algorithm should account for the separation between domains. For traditional machine learning problems α

should be set to 0. For domain adaptation problems, α is set to 1 to minimize the error upper bound, or tuned based on the importance of minimizing the separation between domains. We set α to 1 for all DA experiments reported in this paper - this corresponds directly to the bound in Theorem 5.

We empirically evaluate the feature selection algorithm on a pathological speech database recorded from patients with neurogenic disorders. In particular, we consider the problem of classifying between healthy and dysarthric speech. Dysarthria is a motor speech disorder resulting from an underlying neurological injury. We make use of data collected in the Motor Speech Disorders Laboratory at Arizona State University, consisting of 34 dysarthric speakers and 13 healthy speakers (H). The dysarthria speakers included: 12 speakers with ataxic dysarthria, secondary to cerebellar degeneration (A), 10 mixed flaccid-spastic dysarthria, secondary to amyotrophic lateral sclerosis (ALS), 8 speakers with hypokinetic dysarthria secondary to Parkinson's Disease (PD), and 4 speakers with hyperkinetic dysarthria secondary to Huntington's disease (HD). Each patient provided speech samples, including a reading passage, phrases, and sentences. The speech database consists of approximately 10 minutes of recorded material per speaker. These speech samples were taken from the larger pathological speech database described in [36].

The recordings from each speaker were split into individual sentences by hand and features were extracted at the sentence level. Three different feature sets were used: envelope modulation spectrum (EMS) features, long-term average spectrum (LTAS) features, and ITU-T P.563 features. EMS is a representation of the slow amplitude modulations in a signal and captures aspects of the speech signal related to rhythm. The LTAS features capture atypical average spectral information in the signal. The P.563 features measure atypical and unnatural voice and articulatory quality. For a more detailed discussion of these features, we refer the readers to [37].

In our first experiment we evaluate the FS algorithm based on the criteria in (14). We consider the problem of discriminating between healthy and dysarthric speech based on the features discussed above. For this experiment we form both the training and test sets by randomly drawing 300 dysarthric speech samples and 300 healthy speech samples for each set, ensuring that there is no overlap between training and test data. Using the FS algorithm in Alg. 1, we use the training data to find the top 20 features that maximize the separability between the two groups. We compare this feature selection algorithm against one based on maximizing the Bhattacharyya distance between classes. Using the feature subsets chosen by the two algorithms, we build support vector machine (SVM) classifiers on the training data and evaluate their accuracy on the test data. This experiment is repeated ten times using different randomly generated training and test sets, and the average accuracy is displayed in Figure 3.

The results of this experiment indicate that the initial features selected by the D_F -distance criteria provide faster convergence to the maximum classification rate when compared to those selected by the BC criteria; however, as expected, as additional features are selected, both algorithms eventually converge to roughly the same level of performance. We purposefully restrict ourselves here to a very limited training set (300 samples per class) in

order to evaluate the D_p -distance in a small N setting. Next, we consider the same problem but with a variable number of training samples per class. The results of this experiment are presented in Table III. As the number of training instances increases, the classifier success rate increases for the D_p -based method, however it stays relatively flat for the BC-based method. For very small values of N , the bias/variance associated with the D_p -distance estimator seems to result in features that provide poorer separability when compared to the BC method. Given that the results of this estimator are asymptotic, this is expected. As the number of features increase, both the D_p and BC algorithms converge to approximately the same value.

Next we would like to investigate the efficacy of the FS criteria (15) in a domain adaptation setting. We consider the same problem here - discriminating between healthy and dysarthric individuals; however now we train on data from one disorder and evaluate on data from another disorder. In order to partition the data into dissimilar training and test groups, we start by selecting 300 healthy instances for the training set and 300 (different) healthy instances for the test set. The rest of the training and test data is made up of 300 randomly selected samples from one of the four Dysarthria subtypes: Ataxic, ALS, Huntington's and Parkinson's. Each model is then evaluated on the test sets for each subtype not contained in the training set.

Using each training set-test set combination, we generate feature subsets using the proposed selection algorithm, along with three competing algorithms that are used for comparison. The first algorithm we use for comparison is a standard forward selection algorithm based on the BC distance. This algorithm is used as a baseline for comparison, however because it assumes the training and test data come from the same distribution [6], we expect it to perform poorly relative to the other algorithms. Next we use the same Bhattacharyya FS algorithm, however we account for the separation in domains by using feature normalization, as described in [38], prior to feature selection. We refer to this method as BC with feature normalization (BCFN).

The final domain-invariant feature learning algorithm we compare against is based on Conditional Probability Models (CPM), as described in [39]. This approach attempts to select a sparse mapping that maximizes an objective function that trades off between prediction algorithm performance and the distance between target and source distributions (controlled by a Lagrangian parameter λ). For classification, the logistic regression function is used and a penalization term is added to ensure that the mapping contains minimal contribution from features containing large differences between source and target data. For the specifics of the implementation, we refer the reader to [39]. The same parameter settings are used here. Because this approach utilizes an optimization criteria involving a trade-off between the source-domain separation and the train-test separation, it resembles the proposed FS algorithm more closely than any other method proposed in the literature.

We present the average classification accuracies yielded by the top 20 features from each FS algorithm for each train-test combination in Table IV. The algorithm proposed in this paper achieved the highest classification accuracy in 8 of the 12 trials, while the BC algorithm scored the lowest 8 of 12 trials. The results clearly illustrate the importance of utilizing

domain adaptation in this type of scenario; even an approach as simple as feature normalization yields roughly 8.5 % higher classification accuracy on average. To observe the value of the lower-dimensional subsets generated by each algorithm, we average the accuracy across all twelve trials and display the accuracy as a function of the number of features in Figure 4. We can see in this figure that the performance of the proposed algorithm consistently improves as additional features are added. Because the optimization criteria we have selected minimizes the upper bound on the error, the algorithm has a tendency to pick “safe” features; e.g. using this algorithm invariant features are preferred, even if they are less informative in the source domain.

To better understand how DA helps us build robust models, we look at the top two features returned general and DA FS criteria proposed in this paper. Figure 5a displays the training and test data plotted across the top two features returned by the general FS criteria. We see that these two features represent a strong separation between the two classes in the training set, however this separation is not similarly represented in the test data, and as a result these features will not be beneficial for the target application. Figure 5b displays the data plotted against the top two features returned by the DA FS criteria. Even though the separation between classes in the training data isn’t as noticable as in the features returned by the general criteria, both Dysarthria subtypes manifest themselves very similarly within this feature space, and as a result models built on them will generalize well between these two subtypes.

VI. Conclusion

In this paper we showed that a nonparametric $\mathcal{F}$ -divergence bounds the Bayes classification error rate for two scenarios: the case where training and test data come from the same distribution and the case where training and test data come from different distributions. For the first case, we show that the bound is tighter than the commonly used Bhattacharyya bound on the Bayes error. Our experimental results confirm the theoretical findings - when used as a feature selection criterion in a pathological speech classification problem, the $D_{\mathcal{F}}$ distance yields an improved classification rate with fewer features as compared against popular alternatives.

Future work revolves around analyzing the estimator of the $D_{\mathcal{F}}$ -distance. In particular, understanding the convergence properties of the estimator as a function of the sample size and data dimension will yield insight into the fidelity of the estimation for any given data set. Furthermore, characterizing the bias and variance of this estimator may allow us to apply ensemble estimator methods of [40] to improve estimation accuracy for high dimensional feature space.

Acknowledgments

This research was supported in part by the Office of Naval Research grant N000141410722 (Berisha), National Institute of Health, National Institute on Deafness and Other Communicative Disorders Grant 1R21DC012558 (Berisha), Army Research Office grant W911NF-11-1-0391 (Hero), National Science Foundation grant CCF-1217880, and National Institutes of Health (Hero) grant 2P01CA087634-06A2.

References

1. Bhattacharyya A. On a measure of divergence between two multinomial populations. *Sankhyā: The Indian Journal of Statistics*. 1946:401–406.
2. Kullback S, Leibler RA. On information and sufficiency. *The Annals of Mathematical Statistics*. 1951:79–86.
3. Chernoff H. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *The Annals of Mathematical Statistics*. 1952:493–507.
4. Hero AO, Ma B, Michel O, Gorman J. Alpha-divergence for classification, indexing and retrieval. Communication and Signal Processing Laboratory, Technical Report CSPL-328, U. Mich. 2001
5. Cha SH. Comprehensive survey on distance/similarity measures between probability density functions. *City*. 1(2):1, 2007.
6. Guorong, X.; Peiqi, C.; Minhui, W. Bhattacharyya distance feature selection. *Pattern Recognition*, 1996., *Proceedings of the 13th International Conference on; IEEE*; 1996. p. 195-199.
7. Hamza, AB.; Krim, H. Energy Minimization Methods in Computer Vision and Pattern Recognition. Springer; 2003. Image registration and segmentation by maximizing the Jensen-Rényi divergence; p. 147-163.
8. Hild KE, Erdogmus D, Principe JC. Blind source separation using Renyi's mutual information. *Signal Processing Letters, IEEE*. 2001; 8(6):174–176.
9. Banerjee A, Merugu S, Dhillon IS, Ghosh J. Clustering with Bregman divergences. *The Journal of Machine Learning Research*. 6:1705–1749. 2005.
10. Csiszár I, Shields PC. Information theory and statistics: A tutorial. *Communications and Information Theory*. 2004; 1(4):417–528.
11. Ali S, Silvey SD. A general class of coefficients of divergence of one distribution from another. *Journal of the Royal Statistical Society. Series B (Methodological)*. 1966:131–142.
12. Kailath T. The divergence and Bhattacharyya distance measures in signal selection. *Communication Technology, IEEE Transactions on*. 15(1):52–60. 1967.
13. Hashlamoun WA, Varshney PK, Samarsooriya V. A tight upper bound on the bayesian probability of error. *IEEE Transactions on pattern analysis and machine intelligence*. 16(2):220–224. 1994.
14. Avi-Itzhak H, Diep T. Arbitrarily tight upper and lower bounds on the Bayesian probability of error. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 18(1):89–91. 1996.
15. Csisz I, et al. Information-type measures of difference of probability distributions and indirect observations. *Studia Sci Math Hungar*. 2:299–318. 1967.
16. Kullback S. A lower bound for discrimination information in terms of variation (corresp.). *Information Theory, IEEE Transactions on*. 13(1):126–127. 1967.
17. Vajda I. Note on discrimination information and variation (corresp.). *Information Theory, IEEE Transactions on*. 1970; 16(6):771–773.
18. Ben-David S, Blitzer J, Crammer K, Pereira F, et al. Analysis of representations for domain adaptation. *Advances in neural information processing systems*. 19:137, 2007.
19. Ben-David S, Blitzer J, Crammer K, Kulesza A, Pereira F, Vaughan JW. A theory of learning from different domains. *Machine learning*. 79(1–2):151–175. 2010.
20. Blitzer J, Crammer K, Kulesza A, Pereira F, Wortman J. Learning bounds for domain adaptation,” in. *Advances in neural information processing systems*. 2008:129–136.
21. Mansour Y, Mohri M, Rostamizadeh A. Domain adaptation: Learning bounds and algorithms. 2009arXiv preprint arXiv:0902.3430
22. Mansour, Y.; Mohri, M.; Rostamizadeh, A. Multiple source adaptation and the Rényi divergence. *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence; AUAI Press*; 2009. p. 367-374.
23. Fukunaga, K. Introduction to statistical pattern recognition. Academic press; 1990.
24. Costa JA, Hero AO. Geodesic entropic graphs for dimension and entropy estimation in manifold learning. *Signal Processing, IEEE Transactions on*. 52(8):2210–2221. 2004.
25. Sricharan K, Raich R, Hero AO. Estimation of nonlinear functionals of densities with confidence. *Information Theory, IEEE Transactions on*. 58(7):4135–4159. 2012.

26. Nguyen X, Wainwright MJ, Jordan MI. On surrogate loss functions and f-divergences. *The Annals of Statistics*. 2009:876–904.
27. Nguyen X, Wainwright MJ, Jordan MI. Estimating divergence functionals and the likelihood ratio by convex risk minimization. *Information Theory, IEEE Transactions on*. 56(11):5847–5861. 2010.
28. Berisha V, Hero A. Empirical non-parametric estimation of the fisher information. *Signal Processing Letters, IEEE*. Jul; 2015 22(7):988–992.
29. Friedman JH, Rafsky LC. Multivariate generalizations of the Wald-Wolfowitz and smirnov two-sample tests. *The Annals of Statistics*. 1979:697–717.
30. Rukhin, A. Optimal estimator for the mixture parameter by the method of moments and information affinity. *Trans. 12th Prague Conference on Information Theory*; 1994; p. 214-219.
31. Saon, G.; Padmanabhan, M. *Advances in Neural Information Processing Systems*. MIT Press; 2001. Minimum Bayes error feature selection for continuous speech recognition; p. 75-78.
32. Xuan, G.; Zhu, X.; Chai, P.; Zhang, Z.; Shi, YQ.; Fu, D. Feature selection based on the Bhattacharyya distance. *Pattern Recognition, 2006. ICPR 2006. 18th International Conference on; IEEE; 2006*. p. 1232-1235.
33. Antos A, Devroye L, Györfi L. Lower bounds for Bayes error estimation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*. 21(7):643–645. 1999.
34. Fukunaga, K. *Introduction to statistical pattern recognition*. Academic press; 1990.
35. Liu, H.; Motoda, H. *Computational methods of feature selection*. CRC Press; 2007.
36. Lansford KL, Liss JM. Vowel acoustics in dysarthria: Speech disorder diagnosis and classification. *Journal of Speech, Language, and Hearing Research*. 57(1):57–67. 2014.
37. Berisha, V.; Liss, J.; Sandoval, S.; Utianski, R.; Spanias, A. Modeling pathological speech perception from data with similarity labels. *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on; IEEE; 2014*. p. 915-919.
38. Kinnunen T, Li H. An overview of text-independent speaker recognition: from features to supervectors. *Speech communication*. 52(1):12–40. 2010.
39. Satpal, S.; Sarawagi, S. *Knowledge Discovery in Databases: PKDD 2007*. Springer; 2007. Domain adaptation of conditional probability models via feature subsetting; p. 224-235.
40. Moon K, Hero A. Multivariate f-divergence estimation with confidence,” in. *Advances in Neural Information Processing Systems*. 2014:2420–2428.
41. Henze N, Penrose MD, et al. On the multivariate runs test. *The Annals of Statistics*. 27(1):290–298. 1999.

Appendix A. Proof of Theorem 2

By combining Eq. (2) and (3) from the text we can rewrite

$$D_p = \frac{1}{4pq} [1 - 4pq A_p(f, g) - (p - q)^2] \quad (16)$$

$$= \frac{1 - (p - q)^2}{4pq} - A_p(f, g) \quad (17)$$

$$= 1 - A_p(f, g), \quad (18)$$

where

$$A_p = \int \frac{f(\mathbf{x})g(\mathbf{x})}{pf(\mathbf{x})+qg(\mathbf{x})} d\mathbf{x}. \quad (19)$$

From Theorem 2 in [41], we know that as $N_f \rightarrow \infty$ and $N_g \rightarrow \infty$ in a linked manner such that $\frac{N_f}{N_f+N_g} \rightarrow p$ and $\frac{N_g}{N_f+N_g} \rightarrow q$,

$$\frac{\mathcal{C}(f, g)}{N_f+N_g} \rightarrow 2pqA_p(f, g), \quad (20)$$

almost surely.

Combining the asymptotic relationship in Eq. (20) with the results from Eq. (18), we see that

$$1 - \mathcal{C}(f, g) \frac{N_f+N_g}{2N_fN_g} \rightarrow D_p(f, g), \quad (21)$$

almost surely as $N_f \rightarrow \infty$ and $N_g \rightarrow \infty$ in a linked manner such that $\frac{N_f}{N_f+N_g} \rightarrow p$ and $\frac{N_g}{N_f+N_g} \rightarrow q$.

Appendix B. Proof of Theorem 2

We begin with the realization that the Bayes error rate can be expressed in terms of the total variation (TV) distance between distributions [12]:

$$\varepsilon^{Bayes} = \frac{1}{2} - \frac{1}{2} \int |pf(\mathbf{x}) - qg(\mathbf{x})| d\mathbf{x}. \quad (22)$$

Next, we show that we can bound the TV distance from above and below using $\tilde{D}_p$:

$$\tilde{D}_p = 1 - 4pqA_p(f, g) \quad (23a)$$

$$= 1 - 4pq \int \frac{f(\mathbf{x})g(\mathbf{x})}{pf(\mathbf{x})+qg(\mathbf{x})} d\mathbf{x} \quad (23b)$$

$$= \int [pf(\mathbf{x}) + qg(\mathbf{x})] d\mathbf{x} - 4pq \int \frac{f(\mathbf{x})g(\mathbf{x})}{pf(\mathbf{x}) + qg(\mathbf{x})} d\mathbf{x} \quad (23c)$$

$$= \int \frac{[pf(\mathbf{x}) + qg(\mathbf{x})]^2 - 4pqf(\mathbf{x})g(\mathbf{x})}{pf(\mathbf{x}) + qg(\mathbf{x})} d\mathbf{x} \quad (23d)$$

$$= \int \frac{pf(\mathbf{x})^2 + qg(\mathbf{x})^2 - 2pqf(\mathbf{x})g(\mathbf{x})}{pf(\mathbf{x}) + qg(\mathbf{x})} d\mathbf{x} \quad (23e)$$

$$= \int \frac{[pf(\mathbf{x}) - qg(\mathbf{x})]^2}{pf(\mathbf{x}) + qg(\mathbf{x})} d\mathbf{x} \quad (23f)$$

$$= \int |pf(\mathbf{x}) - qg(\mathbf{x})| \frac{|pf(\mathbf{x}) - qg(\mathbf{x})|}{pf(\mathbf{x}) + qg(\mathbf{x})} d\mathbf{x}. \quad (23g)$$

Since

$$\frac{|pf(\mathbf{x}) - qg(\mathbf{x})|}{pf(\mathbf{x}) + qg(\mathbf{x})} \leq 1 \quad \text{for all } \mathbf{x}, \quad (24)$$

we can simplify (23g) to

$$1 - 4pqA_p(f, g) \leq \int |pf(\mathbf{x}) - qg(\mathbf{x})| d\mathbf{x}. \quad (25)$$

This provides a lower bound on the TV distance based on $\tilde{D}_p$. In order to derive the upper bound we begin with

$$D_{TV}(f, g) = \int |pf(\mathbf{x}) - qg(\mathbf{x})| d\mathbf{x} \quad (26a)$$

$$= \int |pf(\mathbf{x}) - qg(\mathbf{x})| \frac{\sqrt{pf(\mathbf{x}) + qg(\mathbf{x})}}{\sqrt{pf(\mathbf{x}) + qg(\mathbf{x})}} d\mathbf{x} \quad (26b)$$

$$\leq \sqrt{\int \left(\frac{pf(\mathbf{x}) - qg(\mathbf{x})}{\sqrt{pf(\mathbf{x}) + qg(\mathbf{x})}} \right)^2 d\mathbf{x}} \\ \times \sqrt{\int (\sqrt{pf(\mathbf{x}) + qg(\mathbf{x})})^2 d\mathbf{x}} \rightarrow 1$$

(26c)

$$\leq \sqrt{\tilde{D}_p(f, g)}. \quad (26d)$$

By combining the inequalities in (25) and (26d) with the relationship in (22), we see that we can bound the BER by

$$\frac{1}{2} - \frac{1}{2} \sqrt{\tilde{D}_p(f, g)} \leq \varepsilon^{Bayes} \leq \frac{1}{2} - \frac{1}{2} \tilde{D}_p(f, g). \quad (27)$$

Appendix C. Proof of Theorem 3

By the geometric vs harmonic mean inequality,

$$f(\mathbf{x})^q g(\mathbf{x})^p \geq \frac{f(\mathbf{x})g(\mathbf{x})}{pf(\mathbf{x}) + qg(\mathbf{x})}. \quad (28)$$

It immediately follows that $A_p(f, g) = \int f(\mathbf{x})^q g(\mathbf{x})^p$, a scaled Chernoff information function. Thus,

$$A_p(f, g) \leq \int f(\mathbf{x})^q g(\mathbf{x})^p. \quad (29)$$

Appendix D. Proof of Theorem 4

For equiprobable classes ($p=q=\frac{1}{2}$) The upper and lower bounds on the Bayes error rate based on the Bhattacharyya distance are defined by [12]

$$\frac{1 - \sqrt{1 - BC^2(f, g)}}{2} \leq \varepsilon^{Bayes} \leq \frac{BC(f, g)}{2}, \quad (30)$$

where

$$BC(f, g) = \int \sqrt{f(\mathbf{x})g(\mathbf{x})} d\mathbf{x}. \quad (31)$$

To show that the $\tilde{D}_{\frac{1}{2}}$ bound upper bound is tighter than the Bhattacharyya bound we must show that $A_{\frac{1}{2}}(f, g) \leq BC(f, g)$. It is clear that this is the case from Theorem 3. For the $\tilde{D}_{\frac{1}{2}}$ lower bound to be tighter, $BC^2(f, g)$ must be less than equal to $A_{\frac{1}{2}}(f, g)$. We show this to be true using the Cauchy-Schwartz inequality:

$$BC^2(f, g) = \left[\int \sqrt{f(\mathbf{x})g(\mathbf{x})} d\mathbf{x} \right]^2 \quad (32a)$$

$$= \left[\int \frac{\sqrt{f(\mathbf{x})g(\mathbf{x})}}{\sqrt{\frac{1}{2}(f(\mathbf{x})+g(\mathbf{x}))}} \sqrt{\frac{1}{2}(f(\mathbf{x})+g(\mathbf{x}))} d\mathbf{x} \right]^2 \quad (32b)$$

$$\leq \int \frac{f(\mathbf{x})g(\mathbf{x})}{\frac{1}{2}(f(\mathbf{x})+g(\mathbf{x}))} d\mathbf{x} \int \frac{1}{2}(f(\mathbf{x})+g(\mathbf{x})) d\mathbf{x} \xrightarrow{1}$$

(32c)

$$= A_{\frac{1}{2}}(f, g). \quad (32d)$$

Combining both bounds, we see that

$$\frac{1}{2} - \frac{1}{2} \sqrt{1 - BC^2(f, g)} \leq \frac{1}{2} - \frac{1}{2} \sqrt{\tilde{D}_{\frac{1}{2}}(f, g)} \leq \varepsilon^{Bayes} \leq \frac{1}{2} - \frac{1}{2} \tilde{D}_{\frac{1}{2}}(f, g) \leq \frac{1}{2} BC(f, g).$$

Appendix E. Proof of Theorem 5

The proof begins in the same fashion as the result in [19] and then diverges.

$$\varepsilon_T(h, y_T) = \varepsilon_T(h, y_T) + \varepsilon_S(h, y_S) - \varepsilon_S(h, y_S) + \varepsilon_S(h, y_T) - \varepsilon_S(h, y_T) \quad (33a)$$

$$\leq \varepsilon_S(h, y_S) + |\varepsilon_S(h, y_T) - \varepsilon_S(h, y_S)| + |\varepsilon_T(h, y_T) - \varepsilon_S(h, y_T)| \quad (33b)$$

$$\leq \varepsilon_S(h, y_S) + \mathbf{E}_{f_S(\mathbf{x})}[|y_S - y_T|] + \left| \int f_T(\mathbf{x}) |h(\mathbf{x}) - y_T| d\mathbf{x} - \int f_S(\mathbf{x}) |h(\mathbf{x}) - y_T| d\mathbf{x} \right| \quad (33c)$$

$$\leq \varepsilon_S(h, y_S) + \mathbf{E}_{f_S(\mathbf{x})}[|y_S - y_T|] + \int |f_T(\mathbf{x}) - f_S(\mathbf{x})| |h(\mathbf{x}) - y_T| d\mathbf{x} \quad (33d)$$

$$\leq \varepsilon_S(h, y_S) + \mathbf{E}_{f_S(\mathbf{x})}[|y_S - y_T|] + \int |f_T(\mathbf{x}) - f_S(\mathbf{x})| d\mathbf{x} \quad (33e)$$

In (33e), we identify an upper bound on the target error expressed using the TV distance between source and target distributions. Using (26d) this can be expressed in terms of $\tilde{D}_{\frac{1}{2}}$:

$$\varepsilon_T(h, y_T) \leq \varepsilon_S(h, y_S) + E\{|y_S - y_T|\} + 2\sqrt{\tilde{D}_{\frac{1}{2}}(f_T, f_S)} \quad (34)$$

Fig. 1.

Estimation of D_p for the case when (a) $f = g$ and (b) $f \neq g$.

Fig. 2.
The D_p and BC bounds on the Bayes error rate for a bivariate Gaussian example.

Fig. 3.
Average classification accuracy using reduced feature sets

Fig. 4.

Average Classification Accuracy on foreign subtypes using reduced feature sets

(a) Source and target data using top domain-specific features

(b) Source and target data using top domain-invariant features

Fig. 5. Low dimensional representation of datasets (Source Domain: ALS, Target Domain: Parkinson's).

TABLE I

$\mathcal{D}_1$	μ_1	0	0	0	0	0	0	0	0
	σ_1	1	1	1	1	1	1	1	1
	μ_2	2.56	0	0	0	0	0	0	0
	σ_2	1	1	1	1	1	1	1	1
$\mathcal{D}_2$	μ_1	0	0	0	0	0	0	0	0
	σ_1	1	1	1	1	1	1	1	1
	μ_2	3.86	3.10	0.84	0.84	1.64	1.08	0.26	0.01
	σ_2	8.41	12.06	0.12	0.22	1.49	1.77	0.35	2.73

TABLE II

Comparing upper bounds on the Bayes error rate for the multivariate Gaussians defined in Table I.

	Data 1	Data 2
Actual Bayes Error	10%	1.90%
Mahalanobis Bound	18.95%	14.13%
Bhattacharyya Bound	22.04%	4.74%
D_ρ Bound (100 points)	18.23% $\pm$ 3.32%	4.10% $\pm$ 1.10%
D_ρ Bound (500 points)	16.88% $\pm$ 1.51%	2.17% $\pm$ 0.42%
D_ρ Bound (1000 points)	16.46% $\pm$ 1.14%	1.94% $\pm$ 0.29%

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

TABLE III
Average classification accuracies (in percent) of top 10 features selected by D_p and BC divergence

Number of Features	Algorithm	Number of Training Instances				
		100	200	300	400	500
10	BC	86.88	86.93	87.61	87.98	87.22
	D_p	86.36	88.67	89.59	89.20	90.03
15	BC	90.84	90.46	90.51	91.69	90.88
	D_p	88.08	90.66	92.00	92.12	92.72
20	BC	91.10	93.02	93.35	93.98	93.72
	D_p	89.28	92.15	93.20	93.41	94.21

TABLE IV

Classification accuracies of SVM classifier using the top 20 features returned by each feature selection method for each combination of training and test data

Trial	Source	Target	BC	BCFN	CPM	D_p
1	Ataxic	ALS	56.50	73.28	75.82	76.22
2	Ataxic	Huntington's	56.83	72.52	70.12	75.12
3	Ataxic	Parkinson's	49.27	60.75	58.53	64.43
4	ALS	Ataxic	52.95	66.35	54.68	67.15
5	ALS	Huntington's	64.25	73.67	65.50	72.23
6	ALS	Parkinson's	54.32	65.97	69.48	73.60
7	Huntington's	Ataxic	49.95	53.63	43.00	49.30
8	Huntington's	ALS	63.40	64.12	63.17	73.00
9	Huntington's	Parkinson's	59.48	62.22	69.73	76.03
10	Parkinson's	Ataxic	41.13	55.65	42.15	48.23
11	Parkinson's	ALS	62.10	66.30	61.25	67.35
12	Parkinson's	Huntington's	73.67	71.12	64.47	68.98