Direct Estimation of Density Functionals Using a Polynomial Basis

Alan Wisler[®], Visar Berisha, Andreas Spanias, Fellow, IEEE, and Alfred O. Hero, Fellow, IEEE

Abstract—A number of fundamental quantities in statistical signal processing and information theory can be expressed as integral functions of two probability density functions. Such quantities are called density functionals as they map density functions onto the real line. For example, information divergence functions measure the dissimilarity between two probability density functions and are useful in a number of applications. Typically, estimating these quantities requires complete knowledge of the underlying distribution followed by multidimensional integration. Existing methods make parametric assumptions about the data distribution or use nonparametric density estimation followed by high-dimensional integration. In this paper, we propose a new alternative. We introduce the concept of "data-driven basis functions"-functions of distributions whose value we can estimate given only samples from the underlying distributions without requiring distribution fitting or direct integration. We derive a new data-driven complete basis that is similar to the deterministic Bernstein polynomial basis and develop two methods for performing basis expansions of functionals of two distributions. We also show that the new basis set allows us to approximate functions of distributions as closely as desired. Finally, we evaluate the methodology by developing data-driven estimators for the Kullback-Leibler divergences and the Hellinger distance and by constructing empirical estimates of tight bounds on the Bayes error rate.

Index Terms—Divergence estimation, direct estimation, nearest neighbor graphs, Bernstein polynomial.

I. INTRODUCTION

I NFORMATION divergence measures play a central role in the fields of machine learning and information theory. Information divergence functions, functionals that map density functions to IR, have been used in many signal processing applications involving classification [1], segmentation [2], source separation [3], clustering [4], and other domains. In machine learning, a sub-class of these divergences known as

Manuscript received February 20, 2017; revised July 28, 2017 and October 16, 2017; accepted November 6, 2017. Date of publication November 27, 2017; date of current version December 22, 2017. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Chandra Ramabhadra Murthy. The work of V. Berisha was supported in part by the Office of Naval Research under Grants N000141410722 and N000141712826. The work of A. O. Hero was supported by the Army Research Office under Grant W911NF-15-1-0479. (*Corresponding author: Alan Wisler.*)

A. Wisler is with Arizona State University, Tempe, AZ 85281 USA (e-mail: alanwisler@gmail.com).

V. Berisha and A. Spanias are with the Department of Electrical Engineering, Arizona State University, Tempe, AZ 85287 USA (e-mail: visar@asu.edu; spanias@asu.edu).

A. O. Hero is with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109-2122 USA (e-mail: hero@eecs.umich.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TSP.2017.2775587

f-divergences [5], are widely used as surrogate loss functions since they form convex upper bounds on the non-convex 0-1 loss [6].

Although these measures prove useful in a variety of applications, the task of estimating them from multivariate probability distributions using finite sample data can pose a significant challenge. It is especially challenging for continuous distributions, which is the focus of this paper. For algorithms and theory for estimating information measures for discrete distributions the reader is referred to [7]-[12]. For the continuous case treated here, there are three general classes of methods for estimating divergence [13]: 1) parametric methods, 2) non-parametric methods based on density estimation, and 3) non-parametric methods based on direct (or graph-based) estimation. Parametric methods are the most common choice for estimation, and typically offer good convergence rates (1/N) when an accurate parametric model is selected. The fundamental limitation of parametric methods, is that an accurate parametric model is rarely available in real world problems and using an inaccurate parametric model can heavily bias the final estimate. As an alternative, when no parametric form is known, non-parametric density estimates such as kernel density estimation [14], histogram estimation [15], or k-nearest neighbor (k-NN) density estimation [16] are used to characterize the distribution. While these methods are quite powerful in certain scenarios, they are generally high variance, sensitive to outliers, and scale poorly with dimension [13].

An alternative to these two classes of methods is direct (or graph-based) estimation, which exploits the asymptotic properties of minimal graphs in order to *directly* estimate distribution functionals without ever estimating the underlying distributions themselves. These methods have been used to estimate density functionals such as entropy [17]-[19], the information divergence [13], and the D_p -divergence [20]. This class of methods can have faster asymptotic convergence rates [13] and are often simpler to implement than plug-in methods which may have many tuning parameters such as kernel width or histogram bin size. Direct estimators are often inspired by an asymptotic property of a graph-theoretic quantity that can be scaled or modified in order to generate an estimator for a given information-theoretic measure. This approach is customized to a specific form of the density functional and is sometimes difficult to generalize. The estimator for the D_p divergence is an example of this approach [20]. This is in contrast with plug-in estimators, where the same general approach can be used to estimate any distribution functional. Estimators based on influence functions attempt to bridge the gap between the two types of approaches [21]. In that work the authors present a recipe for estimating any smooth functional by using a Von Mises expansion - the

1053-587X © 2017 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

analog of the Taylor expansion for distributions; however that approach still requires that part of the data be used for density estimation.

In this paper, we provide a general approach for estimating a wide range of distribution functionals. We propose decomposing the functional onto a complete set of "data-driven" basis functions; where the term "data-driven" means that the basis is determined directly from the data and does not involve distribution fitting or direct integration. We show that a broad class of distribution functionals can be approximated as closely as desired through linear combinations of our proposed basis, where the weights of the basis expansion are determined through convex optimization. Our approach offers a powerful alternative for estimating information-theoretic distribution functionals. We demonstrate the flexibility of the approach by constructing empirical estimators of bounds on the Bayes error rate using the same basis set.

The remainder of the paper is organized as follows. In the next section, we review the literature in this area. In Section II we provide a detailed description of the problem this paper attempts to solve and establish some of the basic mathematical notation used throughout this paper. In Section III we introduce a set of graph-theoretic basis functions and prove that a wide range of information theoretic quantities can be represented by a linear combination of functions in this set. In Sections IV and V, we explore the limitations of the proposed methodology in the finite sample regime and propose two alternate fitting routines to identify weights to map these basis functions to quantities of interest. In Section VI we empirically investigate how the proposed method can be used to estimate popular divergence measures (the KL-divergence, the Hellinger distance, the D_p divergence), and we compare its performance to various parametric and non-parametric alternatives. In Section VII, we show how the method can be extended to form tighter bounds on the Bayes error rate for binary classification problems. Section VIII offers some concluding remarks.

A. Related Work

A natural method for non-parametric estimation of continuous distribution functionals involves histogram binning followed by plug-in estimation [22], [23]. When the bin-size is adjusted as a function of the number of available samples per bin, this histogram plug-in method is known as Grenander's method of sieves and it enjoys attractive non-parametric convergence rates [24], [25]. While these methods may work well for small data dimension (d = 1, 2), their complexity becomes prohibitive for larger dimensions. Recent work has focused on non-parametric plug-in estimators that are more practical in higher dimensions. These approaches generally only estimate the PDF for the values of samples in the reference data, then calculate the expected value across the sample data in place of numerical integration [26]–[28]. To circumvent the slow convergence associated with these approaches, ensemble methods [27] and methods based on influence functions [21] have been proposed which are capable of achieving the parametric rate $\mathcal{O}(N^{-1})$ MSE convergence if the underlying densities meet certain smoothness conditions [27]. Alternatively, estimates of divergence functions that rely on estimates of the likelihood ratio instead of density estimation have been proposed for

estimating the α -divergence and the L_2 -divergence [6], [29]– [32]. These methods estimate the likelihood ratio of the two density functions and plug that value into the divergence functions. Other approaches that bypass density estimation are the convex optimization approach of [29] to estimate f-divergences and the k-NN graph and minimal spanning tree approaches to estimating Henze-Penrose divergence [20], [33], [34].

Similarly, the approach we propose in this paper bypasses density estimation through a polynomial basis expansion where the basis coefficients are determined through a convex optimization criterion. This provides added flexibility and allows us to easily estimate a large class of distribution functionals and to establish empirical estimates of bounds on the Bayes error rate. Conceptually, this approach is similar to prior work in estimating the entropy of discrete distributions using polynomial approximations [7]–[9], [11].

Bounds on optimal performance are a key component in the statistical signal processing literature. For classification problems, it is often desirable to bound the Bayes error rate (BER) - the minimum achievable error in classification problems. The well-known Chernoff upper bound on the probability of error has been used in a number of statistical signal processing applications [35]. It motivated the Chernoff α -divergence [13]. The Bhattacharyya distance, a special case of the Chernoff α -divergence for $\alpha = \frac{1}{2}$, upper and lower bounds the BER [36], [37]. Beyond the bounds on the BER based on divergence measures, a number of other bounds exist based on other functionals of the distributions [38], [39]. For estimation problems, the Fisher information matrix (FIM) bounds the variance of the optimal unbiased estimator (through it's relationship with the CRLB). The authors have also previously introduced the D_p divergence, a non-parametric f-divergence, and showed that it provides provably tighter bounds on the BER than the BC bound [20]. They extended this work to estimation of the Fisher information in [40].

Our data-driven basis, consisting of Bernstein polynomials, can be used to estimate functionals of distributions and to estimate bounds on Bayes optimal classification performance. Bernstein polynomials of a different form have been used for density estimation [41]–[46]. In contrast to this work, our methods do not rely on density estimation.

II. PROBLEM SETUP

In this section, we will set up the problem and establish the notation that will be used throughout the rest of the paper. We are given a set of data $[\mathbf{X}, \mathbf{y}]$ containing N instances, where each instance is represented by a d-dimensional feature vector \mathbf{x}_i and a binary label y_i . Suppose that this data is sampled from underlying distribution, $f_{\mathbf{x}}(\mathbf{x})$, where

$$f_{\mathbf{x}}(\mathbf{x}) = p_0 f_0(\mathbf{x}) + p_1 f_1(\mathbf{x}) \tag{1}$$

is made up of the two conditional class distributions $f_0(\mathbf{x})$ and $f_1(\mathbf{x})$ for classes 0 and 1, with prior probabilities p_0 and p_1 respectively. If the priors aren't explicitly known, they can be easily estimated from the sample data by measuring the ratio of samples drawn from each class. As a simple application of Bayes theorem, we can define the posterior likelihood of class

1, $\eta(\mathbf{x})$, evaluated at a point $\mathbf{x} = \mathbf{x}^*$, as

$$\eta(\mathbf{x}^*) = P[y = 1 | \mathbf{x} = \mathbf{x}^*] = \frac{P[y = 1] f_{\mathbf{x}}(\mathbf{x}^* | y = 1)}{f_{\mathbf{x}}(\mathbf{x}^*)}$$
$$= \frac{p_1 f_1(\mathbf{x}^*)}{f_{\mathbf{x}}(\mathbf{x}^*)} = \frac{p_1 f_1(\mathbf{x}^*)}{p_0 f_0(\mathbf{x}^*) + p_1 f_1(\mathbf{x}^*)}$$
(2)

We can similarly define the posterior probability for class 0 as

$$P[y = 0 | \mathbf{x} = \mathbf{x}^*] = \frac{p_0 f_0(\mathbf{x}^*)}{p_0 f_0(\mathbf{x}^*) + p_1 f_1(\mathbf{x}^*)},$$
(3)

and since y is binary,

$$P[y = 0 | \mathbf{x} = \mathbf{x}^*] = 1 - \eta(\mathbf{x}^*).$$
(4)

To simplify the notation, we remove the dependence of η on \mathbf{x}^* from portions of the analysis that follow.

Suppose that we wish to estimate some functional $G(f_0, f_1)$ of distributions $f_0(\mathbf{x})$ and $f_1(\mathbf{x})$, which can be expressed in the following form

$$G(f_0, f_1) = \int g(\eta(\mathbf{x})) f_{\mathbf{x}}(\mathbf{x}) d\mathbf{x}.$$
 (5)

Throughout the rest of this paper, we will refer to the $g(\eta)$ in (5) as the posterior mapping function. Many functionals in machine learning and information theory, such as f-divergences and loss functions, can be expressed this way. Consider the family of f-divergences as an example. They are defined as

$$D_{\phi}(f_0, f_1) = \int \phi\left(\frac{f_0(\mathbf{x})}{f_1(\mathbf{x})}\right) f_1(\mathbf{x}) d\mathbf{x}, \tag{6}$$

where $\phi(t)$ is a convex or concave function unique to the given f-divergence. By substituting

$$\frac{f_0(\mathbf{x})}{f_1(\mathbf{x})} = \frac{p_1(1-\eta(\mathbf{x}))}{p_0\eta(\mathbf{x})} \tag{7}$$

and

$$f_1(\mathbf{x}) = \frac{\eta(\mathbf{x})}{p_1} f_{\mathbf{x}}(\mathbf{x})$$
(8)

we can redefine (6) as

$$D_{\phi}(f_0, f_1) = \int \phi\left(\frac{p_1(1-\eta(\mathbf{x}))}{p_0\eta(\mathbf{x})}\right) \frac{\eta(\mathbf{x})}{p_1} f_{\mathbf{x}}(\mathbf{x}) d\mathbf{x}.$$
 (9)

Thus any *f*-divergence can be presented in the form outlined in (5) simply by defining the posterior mapping function $g(\eta)$ as

$$g(\eta) = \frac{\eta}{p_1} \phi\left(\frac{p_1(1-\eta)}{p_0\eta}\right). \tag{10}$$

We propose a procedure for estimating these types of divergence functionals which bypasses density estimation. We do this by representing the functional in terms of the asymptotic limit of a linear combination of graph-theoretic basis functions.

Suppose that there exists a set of basis functions $H_0(\eta), ..., H_k(\eta)$ that can be similarly expressed as

$$H_i(f_0, f_1) = \int h_i(\eta(\mathbf{x})) f_{\mathbf{x}}(\mathbf{x}) d\mathbf{x}.$$
 (11)



Fig. 1. Illustration of two neighborhoods of \mathbf{x}^* for k = 4 and k = 8, instances with y = 0 are blue while instances with y = 1 are red. In the first scenario $\Phi_4(\mathbf{x}^*) = 1$, since only one instance in $\mathcal{N}_4(\mathbf{x}^*)$ is red. In the second scenario $\Phi_8(\mathbf{x}^*) = 3$, since three of the eight instance in $\mathcal{N}_8(\mathbf{x}^*)$ are red.

If we assume that there exists a set of coefficients such that

$$g(\eta) \approx \sum_{i=0}^{\kappa} w_i h_i(\eta), \qquad (12)$$

then consequently

$$G(f_0, f_1) \approx \hat{G}(f_0, f_1) = \sum_{i=0}^k w_i H_i(\eta),$$
(13)

where the sense of approximation is that the ℓ_2 norm of the difference between the right and left hand sides is small. In the following section, we will introduce a set of basis functions that have the desired properties.

III. GRAPH-THEORETIC BASIS FUNCTIONS

Consider the dataset $[\mathbf{X}, \mathbf{y}]$ previously defined. Suppose we select an arbitrary instance \mathbf{x}^* from \mathbf{X} and examine it along with the set of its k-1 nearest neighbors $\mathbf{x}_{NN}^1, \mathbf{x}_{NN}^2, ..., \mathbf{x}_{NN}^{k-1}$ in \mathbf{X} . We can define the neighborhood set $\mathcal{N}_k(\mathbf{x}^*) = [\mathbf{x}^*, \mathbf{x}_{NN}^1, ..., \mathbf{x}_{NN}^{k-1}]$, as the union of \mathbf{x}^* and its k-1 nearest neighbors. Using this we define $\Phi_k(\mathbf{x}^*)$ as the number of instances in the neighborhood set which are drawn from class 1, or alternatively, the sum of y across all points in the neighborhood

$$\Phi_k(\mathbf{x}^*) = \sum_{i:\mathbf{x}_i \in \mathcal{N}_k(\mathbf{x}^*)} y_i.$$
(14)

Fig. 1 provides a simple illustration to help explain how $\Phi_k(\mathbf{x}^*)$ is calculated. Calculating Φ_k is similar to how nearest neighbor classifiers make decisions, but with two important differences:

- 1) The base instance \mathbf{x}^* is considered in the neighborhood indistinguishably from other instances in $\mathcal{N}_k(\mathbf{x}^*)$.
- 2) Where traditional *k*-NN classifiers are concerned only with identifying the majority, we are interested in the exact number of instances drawn from each class.

In essence $\Phi_k(\mathbf{x}^*)$ tells us something about the probability that y = 1 for instances on or near \mathbf{x}^* . Since we are more concerned with the dataset as a whole than the local characteristics in \mathbf{x} ,

we define the statistic $\rho_{r,k,N}(\mathbf{X})$ to be the fraction of instances $\mathbf{x} \in \mathbf{X}$, for which $\Phi_k(\mathbf{x}) = r, r \leq k$. If we define the indicator function $I_{r,k}(\mathbf{x})$ as

$$I_{r,k}(\mathbf{x}) = \begin{cases} 1 & \Phi_k(\mathbf{x}) = r\\ 0 & \text{otherwise,} \end{cases}$$
(15)

then this test statistic $\rho_{r,k,N}(\mathbf{X})$ can be represented by

$$\rho_{r,k,N}(\mathbf{X}) = \frac{1}{N} \sum_{\mathbf{x} \in \mathbf{X}} I_{r,k}(\mathbf{x}).$$
(16)

The function $\rho_{r,k,N}(\mathbf{X})$ is simply the proportion of k-NN neighborhoods that contain exactly r points from class y = 1. This statistic has a number of desirable qualities. We show that this statistic asymptotically converges to a function of the underlying distributions that can be described in the form outlined in (5). The following is proven in Appendix A.

Theorem 1: As the number of samples (N) approaches infinity,

$$\lim_{N \to \infty} \rho_{r,k,N}(\mathbf{X}) \xrightarrow{L^2} \int {\binom{k}{r}} \eta^r(\mathbf{x}) (1 - \eta(\mathbf{x}))^{k-r} f_{\mathbf{x}}(\mathbf{x}) d\mathbf{x}$$

whenever $k/N \rightarrow 0$.

We propose to use the asymptotic form of $\rho_{r,k,N}(\mathbf{X})$ defined in Theorem 1 as a basis function for estimating functionals of the form (5),

$$H_{r,k}(f_0, f_1) = \lim_{N \to \infty} \rho_{r,k,N}(\mathbf{X}) = \int h_{r,k}(\eta(\mathbf{x})) f_{\mathbf{x}}(\mathbf{x}) d\mathbf{x}$$
(17)

where

$$h_{r,k}(\eta) = \binom{k}{r} \eta^r (1-\eta)^{k-r}.$$
 (18)

The function (18) is the *r*th Bernstein basis polynomial of degree k [47]. Bernstein's proof of the Weierstrass Approximation Theorem [48] asserts that any continuous function $g(\eta)$ can be uniformly approximated on $\eta \in [0, 1]$ to any desired accuracy by a linear combination of functions in (18) of the form

$$g(\eta) \approx \sum_{r=0}^{k} g\left(\frac{r}{k}\right) h_{r,k}(\eta).$$
(19)

Combining this result with Theorem 1, we can show that a linear combination of this basis can be used to estimate any function of the form (5).

Theorem 2: For any $G(f_0, f_1)$ that can be expressed in the form

$$G(f_0, f_1) = \int g(\eta(\mathbf{x})) f_{\mathbf{x}}(\mathbf{x}) d\mathbf{x},$$

where $g(\eta)$ is continuous on [0, 1], the approximation

$$\hat{G}_{k,N}(\mathbf{X}) = \sum_{r=0}^{k} g\left(\frac{r}{k}\right) \rho_{r,k,N}(\mathbf{X})$$
(20)

satisfies

$$\lim_{k \to \infty} \lim_{\substack{N \to \infty \\ k/N \to 0}} E\left\lfloor \left(\hat{G}_{k,N}(\mathbf{X}) - G(f_0, f_1) \right)^2 \right\rfloor = 0.$$
(21)

Theorem 2 provides an asymptotically consistent method of estimating a variety of information-theoretic functions that makes no assumptions on the underlying distributions and can be calculated without having to perform density estimation. Throughout the rest of the paper we will refer to the weights g(r/k) in the approximation specified by (20) in Theorem 2 as the Bernstein weights. We next turn to the finite sample properties of the estimator $\rho_{r,k,N}(\mathbf{X})$.

IV. FINITE SAMPLE CONSIDERATIONS

The previous Section investigated the asymptotic properties of linear combinations of the proposed set of empirically estimable basis functions. The asymptotic consistency of the proposed method is valuable, however in real world scenarios, data is inherently a finite resource, and as a result the efficacy of this method is heavily dependent on its convergence characteristics in the finite sample regime. In this section, we will take a detailed look into how restricting both N and k affects our ability to estimate functions of two distributions. To do this, it is necessary to first break down the different sources of error in the proposed methodology.

A. Estimation vs. Approximation Error

The goal of this paper is to empirically estimate the functional $G(f_0, f_1)$ of the two underlying distributions $f_0(\mathbf{x})$ and $f_1(\mathbf{x})$ using a linear combination of directly estimable basis functions

$$\hat{G}_{k,N}(\mathbf{X}) = \sum_{r=0}^{k} w_r \hat{H}_{r,k,N}(\mathbf{X}).$$
(22)

We divide the error of this estimate into two types, the approximation error (e_A) and the estimation error (e_{est}) :

$$e_{T} = G(f_{0}, f_{1}) - \hat{G}_{k,N}(\mathbf{X})$$

= $\underbrace{G(f_{0}, f_{1}) - \hat{G}_{k}(f_{0}, f_{1})}_{=:e_{A}} + \underbrace{\hat{G}_{k}(f_{0}, f_{1}) - \hat{G}_{k,N}(\mathbf{X})}_{=:e_{est}}.$
(23)

This allows us to separate the error in estimating the basis functions from error in fitting to the posterior mapping function. Understanding the trade-off between these two error types will be particularly useful in Section V, where we explore different methods of fitting weights to the desired density functionals.

B. Considerations for Finite k

A finite sample also implies a finite k and impacts the approximation error. Let us consider the Bernstein weighting scheme introduced in (20) for the scenario where the size of the basis set (k) is restricted. Consider the following example problem.

Example: Suppose that we wish to estimate the function

$$g(\eta) = \binom{3}{1}\eta(1-\eta)^2 \tag{24}$$

using the basis set $\beta_{0,3}(\eta), \beta_{1,3}(\eta), \beta_{2,3}(\eta), \beta_{3,3}(\eta)$. Because $g(\eta) = \beta_{1,3}(\eta)$, there exists a set of weights such that

$$\sum_{r=0}^{3} w_r \beta_{r,3}(\eta) = g(\eta),$$
(25)



Fig. 2. Plot of true and estimated basis values vs. r for data drawn from underlying distributions $f_0(\mathbf{x}) \sim \mathcal{N}(\mathbf{0}_3, \mathbf{I}_3)$ and $f_1(\mathbf{x}) \sim \mathcal{N}(\frac{1}{\sqrt{3}}\mathbf{1}_3, \mathbf{I}_3)$.

however, using the Bernstein weighting scheme in (20) yields

$$\hat{g}(\eta) = \sum_{r=0}^{3} g\left(\frac{r}{3}\right) \beta_{r,3}(\eta) = g\left(\frac{0}{3}\right) \beta_{0,3}(\eta) + g\left(\frac{1}{3}\right) \beta_{1,3}(\eta) + g\left(\frac{2}{3}\right) \beta_{2,3}(\eta) + g\left(\frac{3}{3}\right) \beta_{3,3}(\eta) = \frac{4}{3} \eta (1-\eta)^2 + \frac{2}{3} \eta^2 (1-\eta) \neq g(\eta).$$
(26)

It is clear from this example that the Bernstein weighting procedure do not always provide ideal weights when k is restricted. Based on these results, we are motivated to explore alternative weighting procedures in order to improve the performance of this method for the finite sample case. In the following Section, we will introduce a method of finding better weights using convex optimization.

Regardless of how weights are assigned to the approximation, selection of k remains an important factor affecting performance. In order to satisfy Theorem 2, k should be functionally dependent on N, such that N approaches infinity, both $k \to \infty$ and $\frac{k}{N} \to 0$, however this still provides a great degree of freedom in the selection of k.

In general, there are two major competing factors that must be considered when selecting k. The first is that the Weierstrass approximation theorem can exactly represent any posterior mapping function $g(\eta)$ as a linear combination of the proposed basis set only as $k \to \infty$. This provides motivation for choosing a large k-value to ensure the best possible fit of $g(\eta)$. The second factor is that the asymptotic characteristics of $\rho_{r,k,N}(\mathbf{X})$ are dependent on all points in $\mathcal{N}_k(\mathbf{x}^*)$. Moreover the regime for which Theorem 2 holds, requires that $\frac{k}{N} \to 0$, so we are motivated to select k such that $N \gg k$. This means that the selection of k must achieve a compromise in the trade-off between the approximation and estimation errors, since larger values of k will increase the amount of finite sample error made in estimating the individual basis functions, while lower values of k may inhibit our ability to accurately model the desired function in the asymptotic regime.

To illustrate how our ability to estimate the desired set of basis functions varies with k, we calculate the estimated and asymptotic values of $\rho_{r,k,N}(\mathbf{X})$ for k = 10 and k = 100 on data drawn from distributions $f_0(\mathbf{x}) \sim \mathcal{N}(\mathbf{0}_3, \mathbf{I}_3)$ and $f_1(\mathbf{x}) \sim \mathcal{N}(\frac{1}{\sqrt{3}}\mathbf{1}_3, \mathbf{I}_3)$, where $\mathbf{0}_3 = \begin{bmatrix} 0 & 0 & 0 \end{bmatrix}$ and $\mathbf{1}_3 = \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}$, and plot the results in Fig. 2. Estimates are calculated at 3 different sample sizes ($N_0 = N_1 = 100, 1000, 10000$) and each estimate shown in Fig. 2 has been averaged across 500 Monte Carlo trials. While we can estimate the basis set for either k-value with a high degree of accuracy given enough samples, the estimates for k = 10 are noticeably more accurate. In fact, we are able to do about as well with 1000 samples for k = 10 as we are with 10000 samples for k = 100.

V. Optimization Criteria for Fitting Density Functionals

In this section, we propose a convex optimization criterion to identify appropriate weights for fitting information-theoretic functions when k and N are restricted. Inherently, our goal is to minimize the total error, defined in (23), however minimizing this quantity directly isn't feasible since the value of $G(f_0, f_1)$ is unknown. To circumvent this challenge we focus on developing a criterion to minimize the approximation error. We initially develop an optimization criterion that assumes the posterior is uniformly distributed, then propose an alternate method which incorporates an estimate of the posterior density function in order to more accurately model the approximation error.

A. Uniform Optimization Criteria

Recall that the approximation error e_A can be represented as

$$e_A = \int \epsilon(\eta(\mathbf{x})) f_{\mathbf{x}}(\mathbf{x}) d\mathbf{x}, \qquad (27)$$

where

$$\epsilon(\eta) = g(\eta) - \sum_{r=0}^{k} w_r h_{r,k}(\eta).$$
(28)

Since solving (27) requires high-dimensional integration and knowledge of the underlying distributions. However, because η

is a function of \mathbf{x} , $\epsilon(\eta)$ is implicitly a function of \mathbf{x} as well, and by the law of the unconscious statistician [49],

$$e_A = E[\epsilon(\eta)] = \int \epsilon(\eta) f_\eta(\eta) d\eta, \qquad (29)$$

where $f_{\eta}(\eta)$ is the probability density function of the random variable η . Rewriting the error in this form simplifies the region of integration to a well defined space (since $\eta \in [0, 1]$) and circumvents the high dimensionality of \mathbf{x} . While this eliminates some of the challenges in calculating the error it also creates new ones stemming from the fact that $f_{\eta}(\eta)$ is unknown and difficult to estimate due to its implicit dependency on $f_0(\mathbf{x})$ and $f_1(\mathbf{x})$. The task of estimating $f_{\eta}(\eta)$ will be explored in detail in Section V-B, however for the time being we will bypass this challenge and simply attempt to minimize

$$e_A^* = \int |\epsilon(\eta)|^2 d\eta.$$
 (30)

It is worth noting that if $f_{\eta}(\eta)$ is uniformly distributed

$$e_A^* = E\big[|\epsilon(\eta)|^2\big] \ge e_A^2. \tag{31}$$

While there exists an analytical solution for identifying the weights which minimize (30) [50], we use a convex optimization procedure that allows us to also account for the estimation error. If we define a discretized set of posterior values $\tilde{\eta} = [\tilde{\eta}_1, \tilde{\eta}_2, ..., \tilde{\eta}_{\tilde{N}}]$, where $0 \leq \tilde{\eta}_1 < \tilde{\eta}_2 < ... < \eta_{\tilde{N}} \leq 1$, a procedure to identify weights that minimize (30) can be defined as

$$w_0, ..., w_k = \operatorname*{argmin}_{w_0, ..., w_k} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} \left| g(\tilde{\eta}_i) - \sum_{r=0}^k w_r h_{r,k}(\tilde{\eta}_i) \right|^2.$$
(32)

To illustrate the effectiveness of this method, we consider the example problem of trying to estimate the Hellinger distance (a problem we will further explore in Section VI). If we assume both classes have equal prior probability ($p_0 = p_1 = 0.5$), then the posterior mapping function for the squared Hellinger distance is

$$g(\eta) = (\sqrt{\eta} - \sqrt{1 - \eta})^2.$$
 (33)

This function is estimated using this convex weighting procedure as well as the previously described Bernstein weighting procedure, and we compare how well each method models the desired function for values of k varying from 0 to 100. The performances of each method is evaluated by the following formula

$$MSE(\hat{g},g) = \sum_{i=1}^{\tilde{N}} \left| g(\tilde{\eta}_i) - \hat{g}(\tilde{\eta}_i) \right|^2,$$
(34)

and the results are presented in Fig. 3 for a range of k values varying from 1 to 100. This experiment shows that the proposed convex fitting procedure is able to approximate the Hellinger posterior mapping function far more accurately than the Bernstein approximation. This improvement isn't surprising since the proposed method directly minimizes the MSE whereas the Bernstein approximation guarantees consistency only as $k \to \infty$, but it still helps to illustrate the potential for improvement in the Bernstein weights that exists for smaller k.

The expression (34) does not take into account finite sample errors that lead to noisy estimates of the basis functions and thus does not directly reflect our ability to estimate $G(f_0, f_1)$ with a finite sample. The consequences of this could



Fig. 3. Approximation error of each fitting procedure as a function of the number of basis elements k. This is the idealized case where the basis estimation error is zero and the total error is solely due to imperfect approximation of the posterior mapping function $g(\eta)$.

be quite significant. When using a similar approach for entropy estimation, Paninski observed that the good approximation properties were a result of "large oscillating coefficients" which magnify the variance of the corresponding estimator [11]. Additionally, it does not account for the possibility that η is distributed non-uniformly.

To empirically examine the finite sample properties of the two approaches, we repeat the previous experiment, this time estimating the basis functions empirically on samples of data drawn from distributions $f_0(\mathbf{x}) \sim \mathcal{N}(\mathbf{0}_3, \mathbf{I}_3)$ and $f_1(\mathbf{x}) \sim \mathcal{N}(\frac{\mathbf{1}_3}{\sqrt{3}}, \mathbf{I}_3)$. We generate N = 1000 samples (500 samples per class) in each of the 500 iterations of a Monte Carlo simulation, and evaluate the MSE as

$$MSE(G, \hat{G}) = \frac{1}{N_{MC}} \sum_{i=1}^{N_{MC}} \left[G(f_0, f_1) - \hat{G}(f_0, f_1) \right]^2, \quad (35)$$

where N_{MC} represents the number of Monte Carlo iterations. Since we know that the estimation error is scaled by the magnitude of the weights, we also evaluate a modified fitting routine which augments (32) with a regularization term to penalize solutions with large weights,

$$w_0, ..., w_k =$$

$$\underset{w_{0},...,w_{k}}{\operatorname{argmin}} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} \left| g(\tilde{\eta}_{i}) - \sum_{r=0}^{k} w_{r} h_{r,k}(\tilde{\eta}_{i}) \right|^{2} + \frac{\lambda}{k} \sum_{r=0}^{k} w_{r}^{2}, \quad (36)$$

where λ represents a tuning parameter which controls the importance assigned to minimization of the approximation error relative to the estimation error. Intuitively, higher λ values will make sense for smaller data sets to control the variance of the estimator. We set $\lambda = 0.01$ for all experiments conducted in this paper. The results of this experiment are shown in Fig. 4. We immediately see the necessity of the regularization term, as without it the error becomes extremely large for a range of k values. More generally, the inclusion of the regularization term improves the performance at every k value in this experiment. In comparing the Bernstein weights with the convex (regularized) weights, we find that 1) the performance of the convex method is less dependent on the selection of k and 2) the convex



Fig. 4. The total error of each fitting procedure as a function of k, when there is both approximation and estimation error.

weights generally perform better at lower values of k, while the Bernstein weights outperform at higher k. While the peak performance of the Bernstein method is higher than the convex method, there exists no good method of selecting k a priori in order to reliably achieve this performance. In contrast, the convex method with regularization is less sensitive to the value of k selected.

B. Density-Weighted Optimization Criteria

In the optimization criteria introduced in the previous section we implicitly make the assumption that the distribution of the random variable $\eta(\mathbf{x}) \sim f_{\eta}(\eta)$ is uniformly distributed. In this section we will investigate a data-driven estimator for $f_{\eta}(\eta)$. However, before we proceed it is important to clarify what this distribution actually is.

We initially introduced η as the posterior likelihood function for class 1, which we showed in (2) can be represented as a function of the underlying distributions. When this function's input is a known point \mathbf{x}^* , $\eta(\mathbf{x}^*)$ represents the probability that y = 1 given that $\mathbf{x} = \mathbf{x}^*$. However, if the input is a random variable \mathbf{x} , then η is also a random variable, which is distributed according to $f_{\eta}(\eta)$.

Fig. 5 illustrates $f_{\eta}(\eta)$ for two univariate normal distributions $f_0(x) \sim N(0,1)$ and $f_1(x) \sim N(1,1)$. Fig. 5 a displays the two class distributions across x, Fig. 5(b) displays $\eta(x)$ as a function of x, and Fig. 5 c displays $f_{\eta}(\eta)$ as a function of η . We see from this illustration that while, $\eta(x)$ is close to 0 or 1 across most of the region of x that is displayed, $\eta(x)$ remains close to 0.5 in the regions where $f_x(x)$ is greatest. As a results $f_n(\eta)$ is roughly bell-shaped, and the likelihood of η existing at the extremities (close to 0 or 1) is relatively low. Because the probability density in these regions is low, the accuracy of our fit in these regions is less important, and can be given less weight in the fitting routine. Fig. 6 repeats this illustration for two well seperated normal distributions. In this case the distribution of η is such that $f_n(\eta)$ is most dense towards the extremities, and therefore they should be given more weight in the fitting routine. Side by side, these two figures present an interesting contrast. Despite the fact that the set of density functions f_0 and f_1 look quite similar in the two scenarios, the minor difference in separation significantly alters the distribution of η .

In practice the underlying distributions f_0 and f_1 are unknown, and as a result, $f_\eta(\eta)$ is unknown as well. However, if we were able to sample $f_\eta(\eta)$ at $\eta = \tilde{\eta}_i$, a more direct method of minimizing e_T would be to solve

$$\underset{w_{0},...,w_{k}}{\operatorname{argmin}} \sum_{i=1}^{\tilde{N}} \left| g(\tilde{\eta}_{i}) - \sum_{r=0}^{k} w_{r} h_{r,k}(\tilde{\eta}_{i}) \right|^{2} \hat{f}_{\eta}(\tilde{\eta}_{i}) \Delta_{\tilde{\eta}} + \frac{\lambda}{k} \sum_{r=0}^{k} w_{r}^{2},$$
(37)

where $\Delta_{\tilde{\eta}} = \tilde{\eta}_{i+1} - \tilde{\eta}_i$.

Below we show that $\rho_{r,k,N}(\mathbf{X})$, the statistic previously defined in Theorem 1 can be used to sample the PDF of η . This result is stated in Theorem 3.

Theorem 3: As $N \to \infty$ and $k \to \infty$ in a linked manner such that $\frac{k}{N} \to 0$ and $\frac{r}{k} \to \eta^*$

$$k\rho_{r,k,N}(\mathbf{X}) \to f_{\eta}(\eta^*).$$

Theorem 3 is useful as it provides a method of sampling $f_{\eta}(\eta)$ that doesn't depend on estimates of the underlying density functions f_0 and f_1 . Using this result, we can estimate the density of the posterior at point $\tilde{\eta}_i$ as

$$\hat{f}_{\eta}(\tilde{\eta}_i) = \hat{k}_i \rho_{\tilde{r}_i, \tilde{k}_i, N}(\mathbf{X})$$
(38)

where \tilde{r}_i and \tilde{k}_i are integers selected such that $\tilde{\eta}_i = \frac{\tilde{r}_i}{\bar{k}_i}$ for $i = 1, 2, ..., \tilde{N}$. Sampling at exactly $\tilde{\eta}_i$ may not always be possible since the maximum value of k is limited by the sample size, and k determines the resolution of the sampling scheme. Even if it is possible, it may not be desirable to recalculate $\rho_{\tilde{r}_i, \tilde{k}_i, N}(\mathbf{X})$ for different values of k_i because of the computational burden. To overcome these problems we can design our approach such that we utilize the same set of test statistics $\rho_{r,k,N}(\mathbf{X})$ in the estimation of the posterior distribution as are used in the estimation of the basis set. One way to do this is to assign the set of discretized posteriors

$$\tilde{\eta} = [0, \frac{1}{k}, \frac{2}{k}, ..., 1]$$
 (39)

so that it is straightforward to calculate $f_{\eta}(\eta)$ from the known values of $\rho_{r,k,N}(\mathbf{X})$. An alternate approach is to leave $\tilde{\eta}$ unchanged and interpolate $\hat{f}_{\eta}(\eta)$ to determine its value at the desired points. The advantage of this approach is that it doesn't constrain how η is sampled. Throughout the rest of this paper, we will employ the latter method and solve for $\hat{f}_{\eta}(\eta)$ by linearly interpolating between its values on the discretized set (39).

Because this density-weighted fitting routine more directly minimizes the approximation error of the final estimate, we expect it to generally outperform the uniform method, particularly in cases where the density of the posterior is highly non-uniform and where the desired posterior mapping function $g(\eta)$ is difficult to model using the proposed basis set. This hypothesis will be verified in Sections VI and VII, when we empirically evaluate our methods with real data. This improvement comes at a computational cost since the weights are now data-dependent, they must be calculated online, whereas for the uniform method they can be calculated offline and stored. Solving for the k weights can be be implemented with $O(k^3)$ time complexity [51].



Fig. 5. Illustration of the posterior distribution for two close univariate normal distributions.



Fig. 6. Illustration of the posterior distribution for two separated univariate normal distributions.

Current state-of-the-art algorithms for k-NN graphs can be implemented with $\mathcal{O}(N \log N)$ time complexity [52]. As a result the complexity for the Convex (uniform) method is $\mathcal{O}(N \log N)$ assuming that the weights are available. The complexity for the Convex (density weighted) method is $\mathcal{O}(k(N)^3 + N \log N)$ since the weights must be determined for each new dataset.

VI. DIVERGENCE ESTIMATION

In this section, we show how the proposed method can be applied to estimating three f-divergences, the Hellinger distance, the KL-divergence, and the D_p -divergence, directly from data.

A. Hellinger Distance

The Hellinger distance squared is an f-divergence used to quantify the dissimilarity between two probability distributions and is calculated by

$$H^{2}(f_{0}, f_{1}) = \frac{1}{2} \int \left(\sqrt{f_{0}(\mathbf{x})} - \sqrt{f_{1}(\mathbf{x})}\right)^{2} d\mathbf{x}.$$
 (40)

Using the approach proposed in Section V, we can estimate $H^2(f_0, f_1)$ by fitting weights to the posterior mapping function

$$g_H(\eta) = \frac{1}{2} \left(\sqrt{\frac{\eta}{p_1}} - \sqrt{\frac{1-\eta}{p_0}} \right)^2.$$
(41)

To evaluate the efficacy of this method, we conduct four different experiments in which we attempt to estimate the Hellinger distance between two distributions from finite sample data. In the first three experiments, both distributions are normally distributed according to $f(\mathbf{x}) \sim N(\mu \mathbf{1}_d, \boldsymbol{\Sigma}_d)$, where

$$\boldsymbol{\Sigma}_{d} = \begin{bmatrix} \sigma_{1,1} & \sigma_{1,2} & \dots & \sigma_{1,d} \\ \sigma_{2,1} & \sigma_{2,2} & \dots & \sigma_{2,d} \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_{d,1} & \sigma_{d,2} & \dots & \sigma_{d,d} \end{bmatrix},$$
(42)

for $\sigma_{i,j} = \beta^{|i-j|}$. The first experiment evaluates the most basic case where both Gaussians have spherical covariance. The second experiment considers the case where there exists a strong fixed dependency between adjacent dimensions by using an elliptical covariance structure. The third experiment evaluates the case where this dependency between adjacent dimensions is now dependent on which class the data is drawn from. In the fourth experiment, we return to linearly independent dimensions and consider the case where one of the distributions isn't normally distributed, but instead uniformly distributed within a *d*-dimensional hypercube

$$f(\mathbf{x}) = \begin{cases} \frac{1}{(2\beta)^d} & \mathbf{x} \in [\mu - \beta, \mu + \beta]^d \\ 0 & \text{otherwise.} \end{cases}$$
(43)

A detailed description of the distribution types and parameter setting used in each of the four experiments is presented in Table I. In addition to using the proposed method, we also estimate the Hellinger distance using two different plug-in estimators, one based on a parametric density estimator that assumes the data is normally distributed and one based on a k-NN density estimate of the underlying distributions. The proposed method is implemented using $\lambda = 0.01$ and k = 10, while setting $\tilde{\eta} = [0, 0.01, \dots, 1]$. To calculate the k-NN estimate, we use the universal divergence estimation approach described in

	$f_0(\mathbf{x})$			$f_1(\mathbf{x})$		
	Family	μ	β	Family	μ	β
Experiment 1	Normal	0	0	Normal	$\sqrt{\frac{1}{3}}$	0
Experiment 2	Normal	0	0.8	Normal	$\sqrt{\frac{1}{3}}$	0.8
Experiment 3	Normal	0	0.8	Normal	$\sqrt{\frac{1}{3}}$	0.9
Experiment 4	Normal	0	0	Uniform	Ő	3

TABLE I Experiment Overview Table

[53] and implemented in the ITE toolbox [54]. This method allows us to fix k = 10 and still achieve an asymptotically consistent estimate of the divergence.

In each of the first three experiments, the parametric model shows the highest rate of convergence as expected, although in experiment two it is slightly outperformed at smaller sample sizes by the proposed method. In the fourth experiment, when the assumption of Gaussianity in f_1 no longer holds, the parametric solution is significantly biased and as a result, is outperformed by both non-parametric methods at higher sample sizes (N > 2000). Relative to the k-NN plug-in estimator, the proposed method performs slightly worse in experiment 1, slightly better in experiment 2 and significantly better in experiments 3 and 4, with the results being relatively consistent across the various sample sizes. The improvement in performance shown in experiments 3 and 4 suggests that the proposed method offers the greatest benefit when there exists differences in the shapes of the two underlying distributions. Though the density-weighted procedure consistently outperformed the uniform method, the observed improvement was relatively minor in these experiments.

B. Kullback Leibler Divergence

The Kullback Leibler (KL) divergence [55], also sometimes referred to as the KL risk or relative entropy, is an asymmetric measure of divergence between two probability density functions. Using our regular notation the KL-divergence can be calculated by

$$d_{KL}(f_0||f_1) = \int_{-\infty}^{\infty} f_0(\mathbf{x}) \log\left(\frac{f_0(\mathbf{x})}{f_1(\mathbf{x})}\right) d\mathbf{x}.$$
 (44)

While the KL-divergence has the same general purpose as the Hellinger distance, that is to measure the dissimilarity between two probability density functions, it also has several key difference. Firstly since the KL-divergence is asymmetric it doesn't technically qualify as a distance function and $d_{KL}(f_0||f_1)$ isn't necessarily equal to $d_{KL}(f_1||f_0)$. This also means that $d_{KL}(f_0||f_1)$ and $d_{KL}(f_1||f_0)$ will have different posterior mapping functions. We can define the posterior mapping function for $d_{KL}(f_0||f_1)$ as

$$g_{KL}^{0}(\eta) = \frac{1-\eta}{p_0} \log\left(\frac{p_1(1-\eta)}{p_0\eta}\right)$$
 (45)

and the posterior mapping function for $d_{KL}(f_1||f_0)$ as

$$g_{KL}^{1}(\eta) = \frac{\eta}{p_1} \log\left(\frac{p_0\eta}{p_1(1-\eta)}\right)$$
 (46)

such that

$$d_{KL}(f_i||f_{|i-1|}) = \int_{-\infty}^{\infty} g_{KL}^i(\eta(\mathbf{x}))(p_0 f_0(\mathbf{x}) + p_1 f_1(\mathbf{x}))d\mathbf{x}.$$
(47)

It is worth noting that when $p_0 = p_1 = 0.5$

$$g_{KL}^{1}(\eta) = g_{KL}^{0}(1-\eta)$$
(48)

thus $g_{KL}^1(\eta)$ is a reflection of $g_{KL}^0(\eta)$. One challenge presented in modeling the KL-divergence is that the posterior mapping functions are difficult to model due to discontinuities at the end points, since $g_{KL}^0(0) = \infty$ and $g_{KL}^0(1)$ is undefined though

$$\lim_{\eta \to 1^{-}} g_{KL}^{0}(\eta) = 0.$$
(49)

Due to their symmetry g_{KL}^1 has the same problem at the opposite endpoints. To handle this we simply select our discretized set of posteriors $\tilde{\eta}_1, \tilde{\eta}_2, ..., \tilde{\eta}_{\tilde{N}}$ such that $0 < \tilde{\eta}_1 < \tilde{\eta}_2 < ... < \eta_{\tilde{N}} < 1$. For the experiments in this Section, we set

$$\tilde{\eta}_1, \tilde{\eta}_2, \tilde{\eta}_3, ..., \tilde{\eta}_{100}, \tilde{\eta}_{101} = \epsilon, 0.01, 0.02, ..., 0.99, 1 - \epsilon$$
 (50)

where $\epsilon = 10^{-4}$. Using this modified set of discretized posteriors, we repeat the set of four experiments described in Section VI-A to evaluate the proposed methods ability to estimate the KL-divergence. The results of this experiment are displayed in Fig. 8. Like the estimates of the Hellinger distance, the parametric method generally yielded the best performance in the first three experiments, but suffered from a large asymptotic bias in experiment 4. The proposed method once again significantly outperformed the k-NN plug-in estimator in experiments 3 and 4, however the results in experiments 1 and 2 are slightly more nuanced due to the significant difference in performance between the two optimization criteria in these trials. In both of these trials the density-weighted criteria significantly outperforms the uniform method at all sample sizes. In experiment 1 the k-NN plug-in estimator outperforms the regular plug-in estimator at all sample sizes, but is outperformed by the density-weighted method for N > 300. In experiment 2 the k-NN plug-in estimator consistently outperforms both proposed methods, however the improvement over the density-weighted method is marginal.

C. D_p -Divergence

The D_p -divergence is an f-divergence defined by

$$D_{p_0}(f_0, f_1) = \frac{1}{4p_0p_1} \left[\int \frac{(p_0f_0(\mathbf{x}) - p_1f_1(\mathbf{x}))^2}{p_0f_0(\mathbf{x}) + p_1f_1(\mathbf{x})} d\mathbf{x} - (p_0 - p_1)^2 \right].$$
(51)

The D_p -divergence has the unique property of being directly estimable from data using minimum spanning trees [20]. Because of this property, it has been used to form non-parametric estimates of the Fisher information [40] as well as upper bounds on the Bayes error rate in a range of classification problems [20], [56]. The posterior mapping function for the D_p -divergence can be defined as

$$g_{D_p}(\eta) = \frac{(2\eta - 1)^2 - (2p_0 - 1)^2}{4p_0(1 - p_0)}$$
(52)



Fig. 7. Plots of MSE vs. Sample size in estimating the Hellinger distance for the four different experiments outlined in Table I.



Fig. 8. Plots of MSE vs. Sample size in estimating the KL-divergence for the four different experiments outlined in Table I.



Fig. 9. Plots of MSE vs. Sample size in estimating the D_p -divergence for the four different experiments outlined in Table I.

which simplifies to $(2\eta - 1)^2$ when $p_0 = p_1 = 0.5$. We once again repeat the experiments described in Section VI-A to evaluate the proposed methods ability to estimate the D_p -divergence. This experiment provides the unique opportunity to compare the proposed method to a more traditional direct estimation procedure. The results of this experiment are displayed in Fig. 9.

As in the previous experiments, the parametric estimate generally performed the best in the first three experiments, but suffered from a large asymptotic bias in experiment 4. The proposed methods perform better than the MST-based estimation in experiments 1 and 2 but worse in experiments 3 and 4. The relative performance of each method in this experiment was largely consistent across the various sample sizes, though the proposed method seems to be converging slightly faster than the MST method in experiment 4 and could possibly exceed its performance given enough samples. Unlike in the previous experiments, we found no difference in performance between the uniform optimization criteria and the density-weighted criteria in this experiment. This is due to the fact that $g_{D_p}(\eta)$ is a polynomial and can be perfectly represented by the proposed basis set, even when k is truncated. Since we are able to achieve a solution where $\epsilon(\eta) = 0 \forall \eta$, the density weighting has no impact on the final results.



Fig. 10. The Bayes error rate along with the three considered upper bounds displayed as a function of η .

VII. FITTING BOUNDS ON PERFORMANCE

The optimization criteria in the proposed fitting routines gives us the ability to not only approximate information-theoretic functions, but to bound them as well. This is especially useful for forming bounds on the Bayes Error Rate (BER). The Bayes error rate represents the optimal classification performance that is achievable for a given pair of class distributions $f_0(\mathbf{x})$ and $f_1(\mathbf{x})$ with prior probabilities p_0 and p_1 respectively and can be calculated by

$$\epsilon^{\text{Bayes}} = \int_{p_0 f_0(\mathbf{x}) \le p_1 f_1(\mathbf{x})} p_0 f_0(\mathbf{x}) d\mathbf{x} + \int_{p_1 f_1(\mathbf{x}) \le p_0 f_0(\mathbf{x})} p_1 f_1(\mathbf{x}) d\mathbf{x}.$$
(53)

In essence the BER measures the intrinsic difficulty of a particular classification problem based on the data. A thorough understanding of the BER of a particular problem can help design optimal classifiers. Because of the challenges associated with estimating the BER, much of the literature has focused on generating bounds on the BER [20], [37], which are generally formulated in terms of some measure of divergence between the two class distributions. One such bound, the well-known Bhattacharya bound, is given by [37]

$$\frac{1}{2} - \frac{1}{2}\sqrt{1 - BC^2(f_0, f_1)} \le \epsilon^{\text{Bayes}} \le \frac{1}{2}BC(f_0, f_1), \quad (54)$$

where

$$BC(f_0, f_1) = 1 - H^2(f_0, f_1).$$
(55)

While the Hellinger distance here can be estimated via any of the methods discussed in Section VI-A, parametric estimates are most common. Alternatively [20] introduced the bounds

$$\frac{1}{2} - \frac{1}{2}\sqrt{D_{\frac{1}{2}}(f_0, f_1)} \le \epsilon^{\text{Bayes}} \le \frac{1}{2} - \frac{1}{2}D_{\frac{1}{2}}(f_0, f_1) \quad (56)$$

where

$$D_{\frac{1}{2}}(f_0, f_1) = 1 - 2 \int \frac{f_0(\mathbf{x}) f_1(\mathbf{x})}{f_0(\mathbf{x}) + f_1(\mathbf{x})} d\mathbf{x}.$$
 (57)

These bounds have the advantage of being provably tighter than the Bhattacharyya bounds [20]. Furthermore since $D_{\frac{1}{2}}$ represents a particular case of the D_p -divergence, which is estimable directly from data, these bounds bypass the need for density estimation much like the approaches proposed in this paper. While these bounds are significantly tighter than the Bhattacharyya bounds, they still leave room for improvement. Avi-Itzhak proposed arbitrarily tight bounds on the BER in [39], however these bounds require density estimation to be employed in practical problems. In this section, we will use a modified version of the previously described fitting routine in order to investigate how tightly we are able to bound the BER using a linear combination of directly estimable basis functions.

Using the fitting routine described in (36) to bound the BER, requires that we define $g(\eta)$ appropriately for estimation of the BER, and constrain our fit such that

$$\sum_{r=0}^{k} w_r h_{r,k}(\tilde{\eta}_i) \ge g(\tilde{\eta}_i) \quad \forall \tilde{\eta}_i.$$
(58)

We can express (53) as

 ϵ

so $g(\eta) = \min [1 - \eta, \eta]$. Incorporating these changes within the regularized fitting routine described in (36) yields

 $w_0, ..., w_k$

$$= \underset{w_0,\ldots,w_k}{\operatorname{argmin}} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} \left| g(\tilde{\eta}_i) - \sum_{r=0}^k w_r h_{r,k}(\tilde{\eta}_i) \right|^2 + \frac{\lambda}{k} \sum_{r=0}^k w_r^2$$

subject to
$$\sum_{r=0}^k w_r h_{r,k}(\tilde{\eta}_i) \ge g(\tilde{\eta}_i) \quad \forall \tilde{\eta}_i.$$
(60)

Fig. 10 compares the theoretical values of each of these bounds as a function of η . These results indicate that the proposed method offers much tighter theoretical bounds than the other two methods, however this bound is based on the asymptotic properties of the proposed basis set and doesn't consider the limitations of a finite sample estimate.

We evaluate the finite-sample performance of this method by calculating each of the described bounds across the four experiments described in Table I. Fig. 11 displays the ground truth value of the BER, along with the theoretical and estimated values for each of the three bounds (Bhattacharyya, D_p , and convex) across sample sizes ranging from 100 to 10000. The Bhattacharyya bound is calculated based on a parametric plug-in estimator which assumes both class distributions to be normally distributed. The D_p bound is calculated from the Friedman-Rafsky test statistic using the approach described in [20]. Finally the convex bound is calculated as a linear combination of the proposed directly estimable basis functions using weights optimized according to (60). The results of this experiment are largely consistent across the four experiments, the convex method yields the tightest bound, followed by the D_p bound, and finally the Bhattacharyya bound. Except for the Bhattacharyya bound in experiment 4, which is estimated parametrically, all of the bounds appear to converge to their asymptotic solution. While the convex bound generally offers a slightly slower convergence rate than the other two solutions,



Fig. 11. Plots of theoretical and estimate upper bounds on the BER as a function of sample size for the four different experiments outlined in Table I.

TABLE II PARAMETERS FOR 2 8-DIMENSIONAL GAUSSIAN DATA SETS FOR WHICH THE BAYES ERROR RATE IS KNOWN (FROM [57])

\mathcal{D}_0	μ_0	0	0	0	0	0	0	0	0
	σ_0	1	1	1	1	1	1	1	1
	μ_1	2.56	0	0	0	0	0	0	0
	σ_1	1	1	1	1	1	1	1	1
7	μ_0	0	0	0	0	0	0	0	0
	σ_0	1	1	1	1	1	1	1	1
\mathcal{D}_2	μ_1	3.86	3.10	0.84	0.84	1.64	1.08	0.26	0.01
	σ_1	8.41	12.06	0.12	0.22	1.49	1.77	0.35	2.73

TABLE III Comparing Upper Bounds on the Bayes Error Rate for the Multivariate Gaussians Defined in Table II

	Data 1	Data 2
Actual Bayes Error	10%	1.90%
Mahalanobis Bound	$18.90\%\pm0.55\%$	14.07 % \pm 0.45 %
Bhattacharyya Bound	$21.74\%\pm0.87\%$	4.68 % \pm 0.27 %
D_p Bound	$16.51~\% \pm 1.07\%$	3.99 % \pm 0.52 %
Convex Bound	14.17 % \pm 0.86%	$3.87\% \pm 0.43\%$

it remains tighter than the other two bounds across all sample sizes.

In order to further validate this bound we repeat one of the experiments conducted in [20] by evaluating the proposed bound along with the Mahalanobis bound, the Bhattacharyya bound, and the D_p bound on two 8-dimensional Gaussian data sets described in [57]. The mean and standard deviations of f_0 and f_1 for the two data sets are described in Table II, and all dimensions are independent. These data sets allow us to analyze the tightness and validity of the bounds in a higher dimensional setting. For this experiment the sample size was fixed at N = 1000 and only the empirical value of each of the bounds was evaluated. Table III displays the mean and standard deviation of each bound calculated across 500 Monte Carlo iterations for each of the two data sets. In both data sets the convex method provides the tightest bounds on the BER.

VIII. CONCLUSION

This paper introduces a novel method for estimating density functionals which utilizes a set of directly estimable basis functions. The most appealing feature of the proposed method is its flexibility. Where previous methods of direct estimation are generally only applicable to a specific quantity, we show that the basis set can be used to generate an asymptotically consistent estimate of a broad class of density functionals, including all f-divergences and the Bayes error rate. We validate these findings by experimentally evaluating the proposed method's ability to estimate three different divergences (the KL-divergence, the Hellinger distance, and the D_p -divergence) for four pairs of multivariate probability density functions. The results reveal that the proposed method performs competitively with other non-parametric divergence estimation methods, and seems to outperform them in cases where the data from the two distributions have different covariance structures or belong to different families. Additionally we demonstrate how the method can be modified to generate empirically-estimable bounds on the Bayes error rate that are much tighter than existing bounds.

Future work could focus on studying the finite-sample properties of the basis set proposed in this paper, since this represents a major source of error for the proposed methodology. An improved understanding of these properties could enable us to refine the regularization term in our optimization criteria to more accurately model each weights contribution to the estimation error or to develop ensemble methods, like those in [26], for estimating the individual basis functions. Another worthwhile future direction would be on determining an orthogonal version of the Bernstein basis expansion since this would simplify solving for the weights of the expansion.

Appendix A

PROOF OF THEOREM 1

Aspects of this proof mirror the methods used to prove the asymptotic convergence of the k-NN error rate in pages 62–70 of [58], please consult this text for additional details. Construction of this proof requires the introduction of an auxilary variable **u**. We define the augmented data set $[\mathbf{X}, \mathbf{y}, \mathbf{u}]$, where **X** is as defined previously, **u** is a set of i.i.d. random variables which are uniformly distributed in [0, 1], and **y** is defined such that

$$y_i = \begin{cases} 1 & u_i \le \eta(\mathbf{x}_i) \\ 0 & \text{otherwise.} \end{cases}$$
(61)

Now, let us consider an arbitrary point x^* in the support of $f_x(x)$. For each instance x_i , we can define an alternate label

$$y'_{i} = \begin{cases} 1 & u_{i} \leq \eta(\mathbf{x}^{*}) \\ 0 & \text{otherwise.} \end{cases}$$
(62)

Unlike the real labels y_i , these alternate labels y'_i have no dependency on \mathbf{x}_i ; they depend only on the fixed point \mathbf{x}^* . From

these alternate labels we construct the alternate statistic

$$\Phi'_k(\mathbf{x}^*) = \sum_{i:\mathbf{x}_i \in \mathcal{N}_k(\mathbf{x}^*)} y'_i.$$
(63)

Since $\Phi'_k(\mathbf{x}^*)$ is simply the sum of k i.i.d. Bernoulli random variables, we can express the probability that $\Phi'_k(\mathbf{x}^*) = r$ as

$$P[\Phi_k'(\mathbf{x}) = r | \mathbf{x} = \mathbf{x}^*] = \binom{k}{r} \eta^r (\mathbf{x}^*) (1 - \eta(\mathbf{x}^*))^{k-r}.$$
 (64)

We can upper bound the likelihood that $\Phi'_k(\mathbf{x}^*) \neq \Phi_k(\mathbf{x}^*)$ by

$$P\left[\Phi_{k}'(\mathbf{x}^{*}) \neq \Phi_{k}(\mathbf{x}^{*})\right] \leq \sum_{i:\mathbf{x}_{i} \in \mathcal{N}_{k}(\mathbf{x}^{*})} P[y_{i} \neq y_{i}'].$$
(65)

Using the definitions of y_i and y'_i above, (65) can be expressed in terms of the difference between the posterior likelihood at \mathbf{x}^* vs. \mathbf{x}_i as

$$P\left[\Phi_{k}'(\mathbf{x}^{*}) \neq \Phi_{k}(\mathbf{x}^{*})\right] \leq \sum_{x_{i} \in \mathcal{N}_{k}(\mathbf{x}^{*})} E\left[\left|\eta(\mathbf{x}^{*}) - \eta(\mathbf{x}_{i})\right|\right].$$
(66)

Using Lemma 5.4 in [58], we can show that

x

$$\sum_{i \in \mathcal{N}_k(\mathbf{x}^*)} E\big[|\eta(\mathbf{x}^*) - \eta(\mathbf{x}_i)|\big] \to 0$$
(67)

as $N \to \infty$ whenever $\frac{k}{N} \to 0$. Combining (66) and (67)

$$P\left[\Phi'_{k}(\mathbf{x}^{*}) \neq \Phi_{k}(\mathbf{x}^{*})\right] \to 0$$
(68)

as $N \to \infty$ whenever $\frac{k}{N} \to 0$. Furthermore, since convergence in probability implies convergence in distribution [59],

$$\lim_{N \to \infty} P[\Phi_k(\mathbf{x}^*) = r] = P[\Phi'_k(\mathbf{x}^*) = r].$$
(69)

Using (69), we can simplify the expectation of $I_{r,k}(\mathbf{x})$ to

$$\lim_{N \to \infty} E[I_{r,k}(\mathbf{x})] = \lim_{N \to \infty} E[E[I_{r,k}(\mathbf{x})|\mathbf{x} = \mathbf{x}']]$$
$$= E[P[\Phi'_k(\mathbf{x}) = r|\mathbf{x} = \mathbf{x}']]$$
$$= \int {\binom{k}{r}} \eta(\mathbf{x}')^r (1 - \eta(\mathbf{x}'))^{k-r} f_{\mathbf{x}}(\mathbf{x}') d\mathbf{x}' = \rho^* \quad (70)$$

where \mathbf{x}' represents a random variable independent of all instances \mathbf{x}_i and distributed according to $f_{\mathbf{x}}(\mathbf{x}')$. Now, let us evaluate the expression

$$E[(\rho_{r,k,N}(\mathbf{X}) - \rho^*)^2] = E[(\rho_{r,k,N}(\mathbf{X}))^2] - 2E[\rho^*\rho_{r,k,N}(\mathbf{X})] + E[(\rho^*)^2].$$
(71)

Beginning with the first term in (71)

$$E[(\rho_{r,k,N}(\mathbf{X}))^{2}] = E\left[\frac{1}{N^{2}}\sum_{\mathbf{x}_{i}\in\mathbf{X}}\sum_{\mathbf{x}_{j}\in\mathbf{X}}I_{r,k}(\mathbf{x}_{i})I_{r,k}(\mathbf{x}_{j})\right]$$
$$= \frac{1}{N^{2}}E\left[\sum_{\mathbf{x}_{i}\in\mathbf{X}}I_{r,k}(\mathbf{x}_{i}) + \frac{1}{N^{2}}\sum_{\mathbf{x}_{i},\mathbf{x}_{j}\in\mathbf{X}}\sum_{i\neq j}I_{r,k}(\mathbf{x}_{i})I_{r,k}(\mathbf{x}_{j})\right]$$
$$= \frac{1}{N}E[I_{r,k}(\mathbf{x}_{i})] + \frac{N^{2}-N}{N^{2}}E[I_{r,k}(\mathbf{x}_{i})I_{r,k}(\mathbf{x}_{j})]$$
(72)

Note that the random variables $I_{r,k}(\mathbf{x}_i)$ and $I_{r,k}(\mathbf{x}_j)$ are conditionally independent if B is true, where $B = \{\mathbf{x}_i \in \mathbf{X} \setminus \mathcal{N}_k(\mathbf{x}_j) \cap \mathbf{x}_j \in \mathbf{X} \setminus \mathcal{N}_k(\mathbf{x}_i)\}$. Since every instance in $\mathbf{X} \setminus \mathbf{x}_i$ is equally likely to be in $\mathcal{N}_k(\mathbf{x}_i)$, we can bound the probability of the complement of B by $P[\overline{B}] \leq 2k/N$. From this we know that as $N \to \infty$, $P[B] \to 1$ and as a result

$$\lim_{N \to \infty} E\left[I_{r,k}(\mathbf{x}_i)I_{r,k}(\mathbf{x}_j)\right] = \lim_{N \to \infty} E\left[I_{r,k}(\mathbf{x}_i)\right] E\left[I_{r,k}(\mathbf{x}_j)\right]$$
$$= (\rho^*)^2.$$
(73)

Plugging this into (72) yields

$$\lim_{N \to \infty} E[(\rho_{r,k,N}(\mathbf{X}))^2] = \lim_{N \to \infty} \frac{1}{N} \rho^* + \frac{N-1}{N} (\rho^*)^2.$$
(74)

Similarly the second term simplifies to

$$\lim_{N \to \infty} E[\rho_{r,k,N}(\mathbf{X})\rho^*] = \lim_{N \to \infty} \rho^* E\left[\frac{1}{N} \sum_{\mathbf{x}_i \in \mathbf{X}} I_{r,k}(\mathbf{x}_i)\right] = (\rho^*)^2$$
(75)

and the third term clearly equals $(\rho^*)^2$. Substituting these results into (71) yields

$$\lim_{N \to \infty} E\left[(\rho_{r,k,N}(\mathbf{X}) - \rho^*)^2 \right] = \lim_{N \to \infty} \frac{1}{N} \rho^* + \frac{N-1}{N} (\rho^*)^2$$
$$- 2(\rho^*)^2 + (\rho^*)^2$$
$$= \lim_{N \to \infty} \frac{\rho^* - (\rho^*)^2}{N} = 0.$$
(76)

APPENDIX B PROOF OF THEOREM 2

Starting with the expression

$$\lim_{k \to \infty} \lim_{\substack{N \to \infty \\ k/N \to 0}} \hat{G}_{k,N}(\mathbf{X}) = \lim_{k \to \infty} \lim_{\substack{N \to \infty \\ k/N \to 0}} \sum_{r=0}^{k} g\left(\frac{r}{k}\right) \rho_{r,k,N}(\mathbf{X})$$
(77)

we first evaluate the limit with respect to N. Since these conditions mirror those of Theorem 1 and g(r/k) is independent of N, we can rewrite (77) as

$$\lim_{k \to \infty} \sum_{r=0}^{k} g\left(\frac{r}{k}\right) \int {\binom{k}{r}} \eta^{r} (1-\eta)^{k-r} f_{\mathbf{x}}(\mathbf{x}) d\mathbf{x}$$
$$= \int \left[\lim_{k \to \infty} \sum_{r=0}^{k} g\left(\frac{r}{k}\right) {\binom{k}{r}} \eta^{r} (1-\eta)^{k-r} \right] f_{\mathbf{x}}(\mathbf{x}) d\mathbf{x}, \quad (78)$$

which according to Weierstrass' Approximation Theorem simplifies to

$$= \int g(\eta) f_{\mathbf{x}}(\mathbf{x}) d\mathbf{x}$$
$$= G(f_0, f_1). \tag{79}$$

Therefore

$$\lim_{k \to \infty} \lim_{\substack{N \to \infty \\ k/N \to 0}} E\left[\left(\hat{G}_{k,N}(\mathbf{X}) - G(f_0, f_1)\right)^2\right] = 0.$$
(80)

APPENDIX C PROOF OF THEOREM 3

Starting with the augmented data set defined in Appendix IX, remember that $\Phi'_k(\mathbf{x}^*)$ is the sum of k i.i.d. Bernoulli random variables. Therefore $\frac{1}{k}\Phi'_k(\mathbf{x}^*)$ represents the arithmetic mean of these output values, and

$$\lim_{k \to \infty} \frac{1}{k} \Phi'_k(\mathbf{x}^*) = \lim_{k \to \infty} \frac{1}{k} \sum_{i: \mathbf{x}_i \in \mathcal{N}_k(\mathbf{x}^*)} y'_i = \eta(\mathbf{x}^*).$$
(81)

Now, since $\Phi'_k(\mathbf{x})$ must be an integer, its probability of equaling r can be expressed as

$$P[\Phi'_{k}(\mathbf{x}) = r] = P\left[r - 1 < \Phi'_{k}(\mathbf{x}) \le r\right].$$
 (82)

Taking the limit of this expression w.r.t. k, allows us to use (81) to form

$$\lim_{k \to \infty} P\left[r - 1 < \Phi_k'(\mathbf{x}) \le r\right] = \lim_{k \to \infty} P\left[\frac{r - 1}{k} < \eta \le \frac{r}{k}\right]$$
$$= \lim_{k \to \infty} F_\eta\left(\frac{r}{k}\right) - F_\eta\left(\frac{r - 1}{k}\right).$$
(83)

Now if we multiply each side by k, the right hand side takes the form of Newton's difference quotient and can be simplified to the probability density function

$$\lim_{k \to \infty} kP[\Phi'_k(\mathbf{x}) = r] = \lim_{k \to \infty} \frac{F_\eta(\frac{r}{k}) - F_\eta(\frac{r}{k} - \frac{1}{k})}{\frac{1}{k}}$$
$$= \lim_{k \to \infty} \frac{d}{d\eta} F_\eta\left(\frac{r}{k}\right) = \lim_{k \to \infty} f_\eta\left(\frac{r}{k}\right).$$
(84)

Since $\rho_{r,k,N} \to P[\Phi'_k(\mathbf{x}) = r]$ as N and k approach infinity in a linked manner such that $k/N \to 0$ and $\frac{r}{k} \to \eta^*$

$$k\rho_{r,k,N}(\mathbf{X}) \to f_{\eta}(\eta^*).$$
 (85)

REFERENCES

- P. J. Moreno, P. P. Ho, and N. Vasconcelos, "A Kullback-Leibler divergence based kernel for SVM classification in multimedia applications," in *Proc. Adv. Neural Inf. Process. Syst.*, 2003, pp. 1385–1392.
- [2] A. B. Hamza and H. Krim, "Image registration and segmentation by maximizing the Jensen-Rényi divergence," in *Energy Minimization Methods in Computer Vision and Pattern Recognition*. New York, NY, USA: Springer-Verlag, 2003, pp. 147–163.
- [3] K. E. Hild, D. Erdogmus, and J. C. Principe, "Blind source separation using Rényi's mutual information," *IEEE Signal Process. Lett.*, vol. 8, no. 6, pp. 174–176, Jun. 2001.
- [4] A. Banerjee, S. Merugu, I. S. Dhillon, and J. Ghosh, "Clustering with Bregman divergences," J. Mach. Learn. Res., vol. 6, pp. 1705–1749, 2005.
- [5] S. Ali and S. D. Silvey, "A general class of coefficients of divergence of one distribution from another," *J. Roy. Statist. Soc. B (Methodol.)*, vol. 28, pp. 131–142, 1966.
- [6] X. Nguyen, M. J. Wainwright, and M. I. Jordan, "On surrogate loss functions and f-divergences," Ann. Statist., vol. 37, pp. 876–904, 2009.
- [7] Y. Wu and P. Yang, "Minimax rates of entropy estimation on large alphabets via best polynomial approximation," *IEEE Trans. Inf. Theory*, vol. 62, no. 6, pp. 3702–3720, Jun. 2016.
- [8] J. Jiao, K. Venkat, Y. Han, and T. Weissman, "Minimax estimation of functionals of discrete distributions," *IEEE Trans. Inf. Theory*, vol. 61, no. 5, pp. 2835–2885, May 2015.

- [9] Y. Wu and P. Yang, "Optimal entropy estimation on large alphabets via best polynomial approximation," in *Proc. IEEE Int. Symp. Inf. Theory*, 2015, pp. 824–828.
- [10] G. Valiant and P. Valiant, "Estimating the unseen: An n/log (n)-sample estimator for entropy and support size, shown optimal via new CLTs," in *Proc. 43rd Annu. ACM Symp. Theory Comput.*, 2011, pp. 685–694.
- [11] L. Paninski, "Estimation of entropy and mutual information," *Neural Comput.*, vol. 15, no. 6, pp. 1191–1253, 2003.
- [12] L. Paninski, "Estimating entropy on m bins given fewer than m samples," *IEEE Trans. Inf. Theory*, vol. 50, no. 9, pp. 2200–2203, Sep. 2004.
- [13] A. O. Hero, B. Ma, O. Michel, and J. Gorman, "Alpha-divergence for classification, indexing and retrieval," Commun. Signal Process. Lab., Univ. Michigan, Ann Arbor, MI, USA, Tech. Rep. CSPL-328, 2001.
- [14] I. Ahmad and P.-E. Lin, "A nonparametric estimation of the entropy for absolutely continuous distributions (corresp.)," *IEEE Trans. Inf. Theory*, vol. 22, no. 3, pp. 372–375, May 1976.
- [15] L. Györfi and E. C. Van der Meulen, "Density-free convergence properties of various estimators of entropy," *Comput. Statist. Data Anal.*, vol. 5, no. 4, pp. 425–436, 1987.
- [16] A. J. Izenman, "Review papers: Recent developments in nonparametric density estimation," *J. Amer. Statist. Assoc.*, vol. 86, no. 413, pp. 205–224, 1991.
- [17] L. Kozachenko and N. N. Leonenko, "Sample estimate of the entropy of a random vector," *Probl. Peredachi Inf.*, vol. 23, no. 2, pp. 9–16, 1987.
- [18] A. O. Hero III and O. Michel, "Robust entropy estimation strategies based on edge weighted random graphs," in *Proc. SPIE's Int. Symp. Opt. Sci. Eng. Instrum.*, 1998, pp. 250–261.
- [19] D. Pál, B. Póczos, and C. Szepesvári, "Estimation of rényi entropy and mutual information based on generalized nearest-neighbor graphs," in *Proc. Adv. Neural Inf. Process. Syst.*, 2010, pp. 1849–1857.
- [20] V. Berisha, A. Wisler, A. O. Hero III, and A. Spanias, "Empirically estimable classification bounds based on a nonparametric divergence measure," *IEEE Trans. Signal Process.*, vol. 64, no. 3, pp. 580–591, Feb. 2016.
- [21] K. Kandasamy, A. Krishnamurthy, B. Poczos, L. Wasserman, and J. Robins, "Nonparametric von Mises estimators for entropies, divergences and mutual informations," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 397–405.
- [22] J. Silva and S. S. Narayanan, "Information divergence estimation based on data-dependent partitions," *J. Statist. Planning Inference*, vol. 140, no. 11, pp. 3180–3198, 2010.
- [23] G. A. Darbellay *et al.*, "Estimation of the information by an adaptive partitioning of the observation space," *IEEE Trans. Inf. Theory*, vol. 45, no. 4, pp. 1315–1321, May 1999.
- [24] S. Geman and C.-R. Hwang, "Nonparametric maximum likelihood estimation by the method of sieves," *Ann. Statist.*, vol. 10, pp. 401–414, 1982.
- [25] U. Grenander and G. Ulf, "Abstract inference," Tech. Rep. 04; QA276, G7, 1981.
- [26] K. Moon and A. Hero, "Multivariate f-divergence estimation with confidence," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 2420–2428.
- [27] K. R. Moon, K. Sricharan, K. Greenewald, and A. O. Hero, "Improving convergence of divergence functional ensemble estimators," in *Proc. IEEE Int. Symp. Inf. Theory*, 2016, pp. 1133–1137.
- [28] B. Póczos, L. Xiong, and J. Schneider, "Nonparametric divergence estimation with applications to machine learning on distributions," arXiv:1202.3758, 2012.
- [29] X. Nguyen, M. J. Wainwright, and M. I. Jordan, "Estimating divergence functionals and the likelihood ratio by convex risk minimization," *IEEE Trans. Inf. Theory*, vol. 56, no. 11, pp. 5847–5861, Nov. 2010.
- [30] Q. Wang, S. R. Kulkarni, and S. Verdú, "Divergence estimation for multidimensional densities via k-nearest-neighbor distances," *IEEE Trans. Inf. Theory*, vol. 55, no. 5, pp. 2392–2405, May 2009.
- [31] Q. Wang, S. R. Kulkarni, and S. Verdú, "Universal estimation of information measures for analog sources," *Found. Trends Commun. Inf. Theory*, vol. 5, no. 3, pp. 265–353, 2009.
- [32] B. Póczos and J. G. Schneider, "On the estimation of alpha-divergences," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2011, pp. 609–617.
- [33] J. H. Friedman and L. C. Rafsky, "Multivariate generalizations of the Wald-Wolfowitz and Smirnov two-sample tests," *Ann. Statist.*, vol. 7, pp. 697–717, 1979.
- [34] N. Henze and M. D. Penrose, "On the multivariate runs test," Ann. Statist., vol. 27, no. 1, pp. 290–298, 1999.
- [35] H. Chernoff, "A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations," *Ann. Math. Statist.*, vol. 23, pp. 493– 507, 1952.

- [36] A. Bhattacharyya, "On a measure of divergence between two multinomial populations," *Sankhyā, Indian J. Statist.*, vol. 7, pp. 401–406, 1946.
- [37] T. Kailath, "The divergence and Bhattacharyya distance measures in signal selection," *IEEE Trans. Commun. Technol.*, vol. 15, no. 1, pp. 52–60, Feb. 1967.
- [38] W. A. Hashlamoun, P. K. Varshney, and V. Samarasooriya, "A tight upper bound on the Bayesian probability of error," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 16, no. 2, pp. 220–224, Feb. 1994.
- [39] H. Avi-Itzhak and T. Diep, "Arbitrarily tight upper and lower bounds on the Bayesian probability of error," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 18, no. 1, pp. 89–91, Jan. 1996.
- [40] V. Berisha and A. O. Hero, "Empirical non-parametric estimation of the Fisher information," *IEEE Signal Process. Lett.*, vol. 22, no. 7, pp. 988– 992, Jul. 2015.
- [41] A. Leblanc, "On estimating distribution functions using Bernstein polynomials," Ann. Inst. Statist. Math., vol. 64, no. 5, pp. 919–943, 2012.
- [42] B. C. Turnbull and S. K. Ghosh, "Unimodal density estimation using Bernstein polynomials," *Comput. Statist. Data Anal.*, vol. 72, pp. 13–29, 2014.
- [43] S. Ghosal, "Convergence rates for density estimation with Bernstein polynomials," Ann. Statist., vol. 29, pp. 1264–1280, 2001.
- [44] G. Igarashi and Y. Kakizawa, "On improving convergence rate of Bernstein polynomial density estimator," *J. Nonparametric Statist.*, vol. 26, no. 1, pp. 61–84, 2014.
- [45] A. Tenbusch, "Two-dimensional Bernstein polynomial density estimators," *Metrika*, vol. 41, no. 1, pp. 233–253, 1994.
- [46] G. J. Babu and Y. P. Chaubey, "Smooth estimation of a distribution and density function on a hypercube using Bernstein polynomials for dependent random vectors," *Statist. Probab. Lett.*, vol. 76, no. 9, pp. 959–969, 2006.
- [47] G. G. Lorentz, *Bernstein Polynomials*. Providence, RI, USA: Amer. Math. Soc., 2012.
- [48] S. Bernstein, "Démonstration du théorème de Weierstrass fondée sur le calcul des probabilités," *Comm. Soc. Math. Kharkow, Ser.*, vol. 2, no. 13, pp. 49–194, 1912.
- [49] J. A. Gubner, Probability and Random Processes for Electrical and Computer Engineers. Cambridge, U.K.: Cambridge Univ. Press, 2006.
- [50] R. T. Farouki, "Legendre–Bernstein basis transformations," J. Comput. Appl. Math., vol. 119, no. 1, pp. 145–160, 2000.
- [51] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani, "Least angle regression," Ann. Statist., vol. 32, no. 2, pp. 407–499, 2004.
- [52] S. Arya, D. M. Mount, N. S. Netanyahu, R. Silverman, and A. Y. Wu, "An optimal algorithm for approximate nearest neighbor searching fixed dimensions," *J. ACM*, vol. 45, no. 6, pp. 891–923, 1998.
- [53] B. Póczos et al., "Nonparametric kernel estimators for image classification," IEEE Conf. Comput. Vision Pattern Recognit., 2012.
- [54] Z. Szabó, "Information theoretical estimators toolbox," J. Mach. Learn. Res., vol. 15, pp. 283–287, 2014.
- [55] S. Kullback, *Information Theory and Statistics*. North Chelmsford, MA, USA: Courier, 1997.
- [56] A. Wisler, V. Berisha, K. Ramamurthy, D. Wei, and A. Spanias, "Empirically-estimable multi-class performance bounds," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2016, pp. 2594–2598.
- [57] K. Fukunaga, Introduction to Statistical Pattern Recognition. New York, NY, USA: Academic, 1990.
- [58] L. Devroye, L. Györfi, and G. Lugosi, A Probabilistic Theory of Pattern Recognition, vol. 31. New York, NY, USA: SpringerVerlag, 2013.
- [59] A. W. Van der Vaart, Asymptotic Statistics, vol. 3. Cambridge, U.K.: Cambridge Univ. Press, 2000.



Alan Wisler received the B.S. and M.S. degrees in electrical engineering from the University of Texas at Dallas, Dallas, TX, USA, in 2011 and 2012, respectively, and the Ph.D. degree in electrical engineering from Arizona State University, Tempe, AZ, USA, in 2016. He is currently a Postdoctoral Fellow at the New Zealand Institute of Language, Brain, and Behavior. His research interests span a range of topics in machine learning, signal processing, and information theory. His current research is primarily focused on the study and development of nonparametric estima-

tors and their applications to statistical learning challenges in speech processing.



Visar Berisha received the Ph.D. degree from Arizona State University, Tempe, AZ, USA, in 2007.

From 2007 to 2009, he was a member of the technical staff at Lincoln Laboratory, Massachusetts Institute of Technology. Following his appointment at Lincoln Labs, he was a Principal Engineer at Raytheon Co. Since 2013, he has been an Assistant Professor at Arizona State University with a joint appointment with the School of Electrical Computer and Energy Engineering and the Department of Speech and Hearing Sciences. His current research interests

include computational models of speech and audio perception, pathological speech processing, and statistical signal processing



Andreas Spanias (F'03) is a Professor in the School of Electrical, Computer, and Energy Engineering, Arizona State University. He is also the Director of the Sensor Signal and Information Processing (Sen-SIP) Center and the Founder of the SenSIP Industry Consortium (now an NSF I/UCRC site). His research interests are in the areas of adaptive signal processing, speech processing, and sensor systems. He and his student team developed the computer simulation software Java-DSP and its award winning iPhone/iPad and Android versions. He is the author of two text

books: Audio Processing and Coding (Wiley) and DSP: An Interactive Approach (2nd ed.). He served as an Associate Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING and as the General Co-Chair for the IEEE ICASSP-99. He also served as the IEEE Signal Processing Vice-President for Conferences. He received the 2002 IEEE Donald G. Fink Paper Prize Award. He served as the Distinguished Lecturer for the IEEE Signal Processing Society in 2004. He is a Series Editor for the Morgan and Claypool lecture series on Algorithms and Software.



Alfred O. Hero (F'07) received the B.S. (summa cum laude) degree from Boston University, Boston, MA, USA, in 1980, and the Ph.D. degree from Princeton University, Princeton, NJ, USA, in 1984, both in electrical engineering. Since 1984, he has been with the University of Michigan, Ann Arbor, where he is the R. Jamison and Betty Williams Professor of Engineering and the Co-Director of the Michigan Institute for Data Science. His primary appointment is in the Department of Electrical Engineering and Computer Science and he also has appointments, by

courtesy, in the Department of Biomedical Engineering and the Department of Statistics. From 2008 to 2013, he held the Digiteo Chaire d'Excellence at the Ecole Superieure d'Electricite, Gif-sur-Yvette, France. His several research articles have received best paper awards. He received the University of Michigan Distinguished Faculty Achievement Award (2011). He received the IEEE Signal Processing Society Meritorious Service Award (1998), the IEEE Third Millenium Medal (2000), and the IEEE Signal Processing Society Technical Achievement Award (2014). He was the President of the IEEE Signal Processing Society (2006-2008) and was on the Board of Directors of the IEEE (2009-2011) where he served as the Director of Division IX (Signals and Applications). He served on the IEEE TAB Nominations and Appointments Committee (2012-2014). He is currently a member of the Big Data Special Interest Group of the IEEE Signal Processing Society. Since 2011, he has been a member of the Committee on Applied and Theoretical Statistics of the US National Academies of Science. His recent research interests include the data science of high-dimensional spatio-temporal data, statistical signal processing, and machine learning. Of particular interest are applications to networks, including social networks, multimodal sensing and tracking, database indexing and retrieval, imaging, biomedical signal processing, and biomolecular signal processing.