

6-2014

Generalized Exponential Concentration Inequality for Renyi Divergence Estimation

Shashank Singh
Carnegie Mellon University

Barnabas Poczos
Carnegie Mellon University, bapoczos@cs.cmu.edu

Follow this and additional works at: http://repository.cmu.edu/machine_learning



Part of the [Theory and Algorithms Commons](#)

Published In

Journal of Machine Learning Research : Workshop and Conference Proceedings, 32, 333-341.

This Conference Proceeding is brought to you for free and open access by the School of Computer Science at Research Showcase @ CMU. It has been accepted for inclusion in Machine Learning Department by an authorized administrator of Research Showcase @ CMU. For more information, please contact research-showcase@andrew.cmu.edu.

Generalized Exponential Concentration Inequality for Rényi Divergence Estimation

Shashank Singh

Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA 15213 USA

SSS1@ANDREW.CMU.EDU

Barnabás Póczos

Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA 15213 USA

BAPOCZOS@CS.CMU.EDU

Abstract

Estimating divergences in a consistent way is of great importance in many machine learning tasks. Although this is a fundamental problem in nonparametric statistics, to the best of our knowledge there has been no finite sample exponential inequality convergence bound derived for any divergence estimators. The main contribution of our work is to provide such a bound for an estimator of Rényi- α divergence for a smooth Hölder class of densities on the d -dimensional unit cube $[0, 1]^d$. We also illustrate our theoretical results with a numerical experiment.

1. Introduction

There are several important problems in machine learning and statistics that require the estimation of the distance or divergence between distributions. In the past few decades many different kinds of divergences have been defined to measure the discrepancy between distributions. They include the Kullback–Leibler (KL) (Kullback & Leibler, 1951), Rényi- α (Rényi, 1961; 1970), Tsallis- α (Villmann & Haase, 2010), Bregman (Bregman, 1967), L_p , maximum mean discrepancy (Borgwardt et al., 2006), Csiszár’s- f divergence (Csiszár, 1967) and many others.

Most machine learning algorithms operate on finite dimensional feature vectors. Using divergence estimators one can develop machine learning algorithms (such as regression, classification, clustering, and others) that can operate on sets and distributions (Póczos et al., 2012; Oliva et al., 2013). Under certain conditions, divergences can estimate entropy and mutual information. Entropy estimators are important in goodness-of-fit testing (Goria et al., 2005),

parameter estimation in semi-parametric models (Wolsztynski et al., 2005), studying fractal random walks (Alemany & Zanette, 1994), and texture classification (Hero et al., 2002a;b). Mutual information estimators have been used in feature selection (Peng & Dind, 2005), clustering (Aghagolzadeh et al., 2007), optimal experimental design (Lewi et al., 2007), fMRI data processing (Chai et al., 2009), prediction of protein structures (Adami, 2004), boosting and facial expression recognition (Shan et al., 2005). Both entropy estimators and mutual information estimators have been used for independent component and subspace analysis (Learned-Miller & Fisher, 2003; Szabó et al., 2007), as well as for image registration (Kybic, 2006; Hero et al., 2002a;b). For further applications, see Leonenko et al. (2008).

In this paper we will focus on the estimation of Rényi- α divergences. This important class contains the Kullback–Leibler divergence as the $\alpha \rightarrow 1$ limit case and can also be related to the Tsallis- α , Jensen-Shannon, and Hellinger divergences.¹ It can be shown that many information theoretic quantities (including entropy, conditional entropy, and mutual information) can be computed as special cases of Rényi- α divergence.

In our framework, we assume that the underlying distributions are not given explicitly. Only two finite, independent and identically distributed (i.i.d.) samples are given from some unknown, continuous, nonparametric distributions.

Although many of the above mentioned divergences were defined a couple of decades ago, interestingly there are still many open questions left to be answered about the properties of their estimators. In particular, even simple questions, such as rates are unknown for many estimators, and to the best of our knowledge no finite sample exponential concentration bounds have ever been derived for divergence

Proceedings of the 31st International Conference on Machine Learning, Beijing, China, 2014. JMLR: W&CP volume 32. Copyright 2014 by the author(s).

¹Some of the divergences mentioned in the paper are distances as well. To simplify the treatment we will call all of them divergences.

estimators.

Our main contribution is to derive an exponential concentration bound for a particular consistent, nonparametric, Rényi- α divergence estimator. We illustrate the behaviour of the estimator with a numerical experiment.

Organization

In the next section we discuss related work (Section 2). In Section 3 we formally define the Rényi- α divergence estimation problem, and introduce the notation and the assumptions used in the paper. Section 4 presents the divergence estimator that we study in the paper. Section 5 contains our main theoretical contributions concerning the exponential concentration bound of the divergence estimator. Section 6 contains the proofs of our theorems. To illustrate the behaviour of the estimator, we provide a simple numerical experiment in Section 7. We draw conclusions in Section 8.

2. Related Work

Probably the closest work to our contribution is Liu et al. (2012), who derived exponential-concentration bound for an estimator of the two-dimensional Shannon entropy. We generalize these results in several aspects:

1. The estimator of Liu et al. (2012) operates in the unit square $[0, 1]^2$. Our estimator operates on the d -dimensional unit hypercube $[0, 1]^d$.
2. In Liu et al. (2012) the exponential concentration inequality was proven for densities in the Hölder class $\Sigma_\kappa(2, L, 2)$, whereas our inequality applies for densities in the Hölder class $\Sigma_\kappa(\beta, L, r)$ for any fixed $\beta \geq 0, r \geq 1$ (see Section 3.3 for definitions of these Hölder classes).
3. While Liu et al. (2012) estimated the Shannon entropy using one i.i.d. sample set, in this paper we estimate the Rényi- α divergence using two i.i.d. sample sets.

To the best of our knowledge, only very few consistent non-parametric estimators exist for Rényi- α divergences: Póczos & Schneider (2011) proposed a k -nearest neighbour based estimator and proved the weak consistency of the estimator, but they did not study the convergence rate of the estimator.

Wang et al. (2009) provided an estimator for the $\alpha \rightarrow 1$ limit case only, i.e., for the KL-divergence. They did not study the convergence rate either, and there is also an apparent error in this work; they applied the reverse Fatou lemma under conditions when it does not hold. This error originates in the work Kozachenko & Leonenko (1987) and can

also be found in other works. Recently, Pérez-Cruz (2008) has proposed another consistency proof for this estimator, but it also contains some errors: he applies the strong law of large numbers under conditions when it does not hold and also assumes that convergence in probability implies almost sure convergence.

Hero et al. (2002a;b) also investigated the Rényi divergence estimation problem but assumed that one of the two density functions is known. Gupta & Srivastava (2010) developed algorithms for estimating the Shannon entropy and the KL divergence for certain parametric families.

Recently, Nguyen et al. (2010) developed methods for estimating f -divergences using their variational characterization properties. They estimate the likelihood ratio of the two underlying densities and plug that into the divergence formulas. This approach involves solving a convex minimization problem over an infinite-dimensional function space. For certain function classes defined by reproducing kernel Hilbert spaces (RKHS), however, they were able to reduce the computational load from solving infinite-dimensional problems to solving n -dimensional problems, where n denotes the sample size. When n is large, solving these convex problems can still be very demanding. They studied the convergence rate of the estimator, but did not derive exponential concentration bounds for the estimator.

Sricharan et al. (2010); Laurent (1996); Birge & Massart (1995) studied the estimation of non-linear functionals of density. They, however, did not study the Rényi divergence estimation and did not derive exponential concentration bounds either. Using ensemble estimators, Sricharan et al. (2012) derived fast rates for entropy estimation, but they did not investigate the divergence estimation problem.

Leonenko et al. (2008) and Gorja et al. (2005) considered Shannon and Rényi- α entropy estimation from a single sample.² In this work, we study divergence estimators using two independent samples. Recently, Pál et al. (2010) proposed a method for consistent Rényi information estimation, but this estimator also uses one sample only and cannot be used for estimating Rényi divergences.

Further information and useful reviews of several different divergences can be found, e.g., in Villmann & Haase (2010).

²The original presentations of these works contained some errors; Leonenko & Pronzato (2010) provide corrections for some of these theorems.

3. Problem and Assumptions

3.1. Notation

We use the notation of multi-indices common in multivariable calculus to index several expressions. For example, for analytic functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$,

$$f(\vec{y}) = \sum_{\vec{i} \in \mathbb{N}^d} \frac{D^{\vec{i}} f(\vec{x})}{\vec{i}!} (\vec{y} - \vec{x})^{\vec{i}},$$

where \mathbb{N}^d is the set of d -tuples of natural numbers,

$$\vec{i}! := \prod_{k=1}^d i_k!, \quad (\vec{y} - \vec{x})^{\vec{i}} := \prod_{k=1}^d (y_k - x_k)^{i_k}$$

and

$$D^{\vec{i}} f := \frac{\partial^{|\vec{i}|} f}{\partial^{i_1} x_1 \cdots \partial^{i_d} x_d}, \quad \text{for} \quad |\vec{i}| := \sum_{k=1}^d i_k.$$

We also use the multinomial theorem, which states that, $\forall k \in \mathbb{N}, \vec{x} \in \mathbb{R}^d$.

$$\left(\sum_{i=1}^d x_i \right)^k = \sum_{|\vec{i}|=k} \frac{k!}{\vec{i}!} \vec{x}^{\vec{i}}. \quad (1)$$

3.2. Problem

For a given $d \geq 1$, consider random d -dimensional real vectors X and Y in the unit cube $\mathcal{X} := [0, 1]^d$, distributed according to densities $p, q : \mathcal{X} \rightarrow \mathbb{R}$, respectively. For a given $\alpha \in (0, 1) \cup (1, \infty)$, we are interested in using a random sample of n i.i.d. points from p and n i.i.d. points from q to estimate the Rényi α -divergence

$$D_\alpha(p||q) = \frac{1}{\alpha - 1} \log \left(\int_{\mathcal{X}} p^\alpha(\vec{x}) q^{1-\alpha}(\vec{x}) d\vec{x} \right).$$

3.3. Assumptions

Density Assumptions:

We assume that p and q are in the bounded Hölder class $\Sigma_\kappa(\beta, L, r)$. That is, for some $\beta \in (0, \infty)$, if $\ell = \lfloor \beta \rfloor$ is the greatest integer with $\ell < \beta$, $\forall \vec{i} \in \mathbb{N}^d$ with $|\vec{i}| = \ell$, the densities p and q each satisfy a β -Hölder condition in the r -norm ($r \in [1, \infty)$):

$$\begin{aligned} |D^{\vec{i}} p(\vec{x} + \vec{v}) - D^{\vec{i}} p(\vec{x})| &\leq L \|\vec{v}\|_r^{\beta - \ell} \\ &= L \left(\sum_{i=1}^d |v_i|^r \right)^{(\beta - \ell)/r}, \end{aligned} \quad (2)$$

and, furthermore, there exist $\kappa = (\kappa_1, \kappa_2) \in (0, \infty)^2$ with $\kappa_1 \leq p, q \leq \kappa_2$. We could take p and q

to be in different Hölder classes $\Sigma_{\kappa_p}(\beta_p, L_p, r_p)$ and $\Sigma_{\kappa_q}(\beta_q, L_q, r_q)$, but the bounds we show depend, asymptotically, only on the weaker of the conditions on p and q (i.e., $\min\{\beta_p, \beta_q\}, \max\{L_p, L_q\}$, etc.).

It is worth commenting on the case that p (similarly, q) is γ times continuously differentiable for a positive integer γ . Since \mathcal{X} is compact, the γ -order derivatives of p are bounded. Hence, since \mathcal{X} is convex, the $(\gamma - 1)$ -order derivatives of p are Lipschitz, by the Mean Value Theorem. Consequently, any degree of continuous differentiability suffices for this assumption.

The existence of an upper bound κ_2 is trivial, since p, q are continuous and \mathcal{X} is compact. The existence of a positive lower bound κ_1 for q is natural, as otherwise Rényi- α divergence may be ∞ . The existence of κ_1 for p is a technical necessity due to certain singularities at 0 (see the Logarithm Bound in Section 6.1). However, in the important special case of Rényi- α entropy (i.e., q is the uniform distribution), the assumption of κ_1 for p can be dropped via an argument using Jensen's Inequality.

As explained later, we also desire p and q to be nearly constant near the boundary $\partial\mathcal{X} = \{\vec{x} \in \mathcal{X} : x_j \in \{0, 1\} \text{ for some } j \in [d]\}$ of \mathcal{X} . Thus, we assume that, for any sequence $\{\vec{x}_n\}_{n=1}^\infty \in \mathcal{X}$ with $\text{dist}(\vec{x}_n, \partial\mathcal{X}) \rightarrow 0$ as $n \rightarrow \infty$, $\forall \vec{i} \in \mathbb{N}^d$ with $1 \leq |\vec{i}| \leq \ell$,

$$\lim_{n \rightarrow \infty} D^{\vec{i}} q(\vec{x}_n) = \lim_{n \rightarrow \infty} D^{\vec{i}} p(\vec{x}_n) = 0.$$

Kernel Assumptions:

We assume the kernel $K : \mathbb{R} \rightarrow \mathbb{R}$ is non-negative, with bounded support $[-1, 1]$ and the following properties (with respect to the Lebesgue measure):

$$\int_{-1}^1 K(u) du = 1, \quad \int_{-1}^1 u^j K(u) du = 0,$$

for all $j \in \{1, \dots, \ell\}$, and K has finite 1-norm, i.e.,

$$\int_{-1}^1 |K(u)| du = \|K\|_1 < \infty. \quad (3)$$

4. Estimator

Let $[d] := \{1, 2, \dots, d\}$, and let

$$\begin{aligned} \mathcal{S} := \{ (S_1, S_2, S_3) : S_1 \cup S_2 \cup S_3 = [d], \\ S_i \cap S_j = \emptyset \text{ for } i \neq j \} \end{aligned}$$

denote the set of partitions of $[d]$ into 3 distinguishable parts. For a bandwidth $h \in (0, 1)$ (to be specified later), for each $S \in \mathcal{S}$, define the region

$$\begin{aligned} C_S = \{ x \in \mathcal{X} : \forall i \in S_1, 0 \leq x_i \leq h, \\ \forall j \in S_2, h < x_j < 1 - h, \\ \forall k \in S_3, 1 - h \leq x_k \leq 1 \} \end{aligned}$$

and the regional kernel $K_S : [-1, 2]^d \times \mathcal{X} \rightarrow \mathbb{R}$ by

$$K_S(\vec{x}, \vec{y}) := \prod_{j \in S_1} K\left(\frac{x_j + y_j}{h}\right) \cdot \prod_{j \in S_2} K\left(\frac{x_j - y_j}{h}\right) \cdot \prod_{j \in S_3} K\left(\frac{x_j - 2 + y_j}{h}\right).$$

Note that $\{C_S : S \in \mathcal{S}\}$ partitions \mathcal{X} (as illustrated in Figure 1), up to intersections of measure zero, and that K_S is supported only on $[-1, 2]^d \times C_S$. The term $K\left(\frac{x_j + y_j}{h}\right)$ corresponds to reflecting y across the hyperplane $x_j = 0$, whereas the term $K\left(\frac{x_j - 2 + y_j}{h}\right)$ reflects y across $x_j = 1$, so that $K_S(\vec{x}, y)$ is the product kernel (in x), with uniform bandwidth h , centered around a reflected copy of y .




Figure 1. Illustration of regions $C_{(S_1, S_2, S_3)}$ with $3 \in S_1$. The region labeled R corresponds to $S_1 = \{3\}, S_2 = \{1\}, S_3 = \{2\}$.

We define the “mirror image” kernel density estimator

$$\tilde{p}_h(\vec{x}) = \frac{1}{nh^d} \sum_{i=1}^n \sum_{S \in \mathcal{S}} K_S(\vec{x}, \vec{x}^i),$$

where \vec{x}^i denotes the i^{th} sample. Since the derivatives of p and q vanish near $\partial\mathcal{X}$, p and q are approximately constant near $\partial\mathcal{X}$, and so the mirror image estimator attempts to reduce boundary bias by mirroring data across $\partial\mathcal{X}$ before kernel-smoothing. We then clip the estimator at κ_1 and κ_2 :

$$\hat{p}_h(x) = \min(\kappa_2, \max(\kappa_1, \tilde{p}_h(x))).$$

Finally, we plug our clipped density estimate into the following plug-in estimator for Rényi α -divergence:

$$\begin{aligned} D_\alpha(p||q) &:= \frac{1}{\alpha - 1} \log \left(\int_{\mathcal{X}} p^\alpha(\vec{x}) q^{1-\alpha}(\vec{x}) d\vec{x} \right) \\ &= \frac{1}{\alpha - 1} \log \left(\int_{\mathcal{X}} f(p(\vec{x}), q(\vec{x})) d\vec{x} \right) \end{aligned} \quad (4)$$

for $f : [\kappa_1, \kappa_2]^2 \rightarrow \mathbb{R}$ defined by $f(x_1, x_2) := x_1^\alpha x_2^{1-\alpha}$. Our α -divergence estimate is then $D_\alpha(\hat{p}_h||\hat{q}_h)$.

5. Main Result

Rather than the usual decomposition of mean squared error into squared bias and variance, we decompose the error $|D_\alpha(\hat{p}_h||\hat{q}_h) - D_\alpha(p||q)|$ of our estimator into a bias term and a variance-like term via the triangle inequality:

$$|D_\alpha(\hat{p}_h||\hat{q}_h) - D_\alpha(p||q)| \leq \underbrace{|D_\alpha(\hat{p}_h||\hat{q}_h) - \mathbb{E}D_\alpha(\hat{p}_h||\hat{q}_h)|}_{\text{variance-like term}} + \underbrace{|\mathbb{E}D_\alpha(\hat{p}_h||\hat{q}_h) - D_\alpha(p||q)|}_{\text{bias term}}.$$

We will prove the “variance” bound

$$\mathbb{P}(|D_\alpha(\hat{p}_h, \hat{q}_h) - \mathbb{E}D_\alpha(\hat{p}_h, \hat{q}_h)| > \varepsilon) \leq 2 \exp\left(-\frac{k_1 \varepsilon^2 n}{\|K\|_1^{2d}}\right),$$

and the bias bound

$$|\mathbb{E}D_\alpha(\hat{p}_h||\hat{q}_h) - D_\alpha(p||q)| \leq k_2 \left(h^\beta + h^{2\beta} + \frac{1}{nh^d} \right),$$

where k_1, k_2 are constant in the sample size n and bandwidth h (see (15) and (16) for exact values of these constants). The variance bound does not depend on h , while the bias bound is minimized by $h \asymp n^{-\frac{1}{d+\beta}}$, giving the convergence rate

$$|\mathbb{E}D_\alpha(\hat{p}_h||\hat{q}_h) - D_\alpha(p||q)| \in O\left(n^{-\frac{\beta}{d+\beta}}\right).$$

Note that we can use this exponential concentration bound to obtain a bound on the variance of $D(\hat{p}_h||\hat{q}_h)$. If $F : [0, \infty) \rightarrow \mathbb{R}$ is the cumulative distribution of the squared deviation of $D_\alpha(\hat{p}_h||\hat{q}_h)$ from its mean,

$$\begin{aligned} 1 - F(\varepsilon) &= \mathbb{P}((D_\alpha(\hat{p}_h, \hat{q}_h) - \mathbb{E}D_\alpha(\hat{p}_h, \hat{q}_h))^2 > \varepsilon) \\ &\leq 2 \exp\left(-\frac{k_1 n}{\|K\|_1^{2d}}\right). \end{aligned}$$

Thus,

$$\begin{aligned} \mathbb{V}[D_\alpha(\hat{p}_h||\hat{q}_h)] &= \mathbb{E}[D_\alpha(\hat{p}_h, \hat{q}_h) - \mathbb{E}D_\alpha(\hat{p}_h, \hat{q}_h)]^2 \\ &= \int_0^\infty (1 - F(\varepsilon)) d\varepsilon \\ &\leq \int_0^\infty 2 \exp\left(-\frac{k_1 n \varepsilon}{\|K\|_1^{2d}}\right) d\varepsilon \\ &= 2 \frac{\|K\|_1^{2d}}{k_1} n^{-1}. \end{aligned} \quad (5)$$

We then have a mean squared-error of

$$\mathbb{E}\left[(D(\hat{p}_h||\hat{q}_h) - D(p||q))^2\right] \in O\left(n^{-1} + n^{-\frac{2\beta}{d+\beta}}\right).$$

which is in $O(n^{-1})$ if $\beta \geq d$ and in $O\left(n^{-\frac{2\beta}{d+\beta}}\right)$ otherwise. This asymptotic rate is consistent with previous bounds in density functional estimation (Birge & Massart, 1995; Sricharan et al., 2010).

6. Proof of Main Result

6.1. Lemmas

Bound on Derivatives of f : Let f be as in (4). Since f is analytic on the compact domain $[\kappa_1, \kappa_2]^2$, there is a constant $C_f \in \mathbb{R}$, depending only on κ , and α , such that, $\forall \xi \in (\kappa_1, \kappa_2)^2$,

$$\begin{aligned} & \left| \frac{\partial f}{\partial x_1}(\xi) \right|, \left| \frac{\partial f}{\partial x_2}(\xi) \right|, \left| \frac{\partial^2 f}{\partial x_1^2}(\xi) \right|, \\ & \left| \frac{\partial^2 f}{\partial x_2^2}(\xi) \right|, \left| \frac{\partial^2 f}{\partial x_1 \partial x_2}(\xi) \right| \leq C_f. \end{aligned} \quad (6)$$

C_f can be computed explicitly by differentiating f and observing that the derivatives of f are monotone in each argument. We will use this bound later in conjunction with the Mean Value and Taylor's theorems.

Logarithm Bound: If $g, \hat{g} : \mathcal{X} \rightarrow \mathbb{R}$ with $0 < c \leq g, \hat{g}$ for some $c \in \mathbb{R}$ depending only on κ and α , then, by the Mean Value Theorem, there exists C_L depending only on κ and α such that

$$\begin{aligned} & \left| \log \left(\int_{\mathcal{X}} \hat{g}(\vec{x}) d\vec{x} \right) - \log \left(\int_{\mathcal{X}} g(\vec{x}) d\vec{x} \right) \right| \\ & \leq C_L \int_{\mathcal{X}} |\hat{g}(\vec{x}) - g(\vec{x})| d\vec{x}. \end{aligned} \quad (7)$$

We will use this bound to eliminate logarithms from our calculations.

Bounds on Derivatives of p : Combining the assumption that the derivatives of p vanish on $\partial\mathcal{X}$ and the Hölder condition on p , we bound the derivatives of p near $\partial\mathcal{X}$. In particular, we show that, if $\vec{i} \in \mathbb{N}^d$ has $1 \leq |\vec{i}| \leq \ell$, then, $\forall \vec{x} \in \mathcal{B} := \{\vec{x} \in \mathcal{X} : \text{dist}(\vec{x}, \partial\mathcal{X}) \leq h\}$

$$|D^{\vec{i}} p(\vec{x})| \leq \frac{Lh^{\beta-|\vec{i}|}}{(\ell-|\vec{i}|)!}. \quad (8)$$

Proof: We proceed by induction on $|\vec{i}|$, as $|\vec{i}|$ decreases from ℓ to 0. The case $|\vec{i}| = \ell$ is precisely the Hölder assumption (2). Now suppose that we have the desired bound for derivatives of order $|\vec{i}| + 1$. Let $\vec{x} \in \partial\mathcal{X}$, $\vec{u} = (0, \dots, 0, \pm 1, 0, \dots, 0) \in \mathbb{R}^d$, where $u_j = \pm 1$. If $\vec{y} + h\vec{u} \in \mathcal{X}$ (any $\vec{x} \in \mathcal{B}$ is clearly of this form, for some $j \in [d]$), then

$$\begin{aligned} |D^{\vec{i}} p(\vec{y} + \vec{u})| & \leq \int_0^h \left| \frac{\partial}{\partial x_j} D^{\vec{i}} p(\vec{y} + t\vec{u}) \right| dt \\ & \leq \int_0^h \frac{Lt^{\beta-(|\vec{i}|+1)}}{(\ell-|\vec{i}|-1)!} dt \\ & = \frac{Lh^{\beta-|\vec{i}|}}{(\beta-|\vec{i}|)(\ell-|\vec{i}|-1)!} \leq \frac{Lh^{\beta-|\vec{i}|}}{(\ell-|\vec{i}|)!}. \end{aligned}$$

The desired result follows by induction on $|\vec{i}|$. ■

Bound on Integral of Mirrored Kernel: A key property of the mirrored kernel is that the mass of the kernel over \mathcal{X} is preserved, even near the boundary of \mathcal{X} , as the kernels about the reflected data points account exactly for the mass of the kernel about the original data point that is not in \mathcal{X} . In particular, $\forall \vec{y} \in \mathcal{X}$,

$$\sum_{S \in \mathcal{S}} \int_{\mathcal{X}} |K_S(\vec{x}, \vec{y})| d\vec{x} = h^d \|K\|_1^d. \quad (9)$$

Proof: For each $S \in \mathcal{S}$, the change of variables




Figure 2. A data point $x^1 \in C_{(\{1,2\}, \emptyset, \emptyset)} \subset [0, 1]^2$, along with its three reflected copies. The sum of the integrals over \mathcal{X} of the four kernels (shaded) is $\|K\|_1$.

$$\begin{aligned} u_j &= -x_j, & j \in S_1 \\ u_j &= x_j, & j \in S_2 \\ u_j &= 2 - x_j, & j \in S_3 \end{aligned}$$

returns the reflected data point created by K_S back onto its original data point. Applying this change of variables gives

$$\sum_{S \in \mathcal{S}} \int_{\mathcal{X}} |K_S(\vec{x}, \vec{y})| d\vec{x} = \int_{[-1, 2]^d} \left| K^d \left(\frac{u-y}{h} \right) \right| d\vec{x},$$

where $K^d(\vec{x}) := \prod_{i=1}^d K(x_i)$ denotes the product kernel.

Rescaling, translating, and applying Fubini's Theorem,

$$\begin{aligned} \sum_{S \in \mathcal{S}} \int_{\mathcal{X}} |K_S(\vec{x}, \vec{y})| d\vec{x} &= h^d \int_{[-1, 1]^d} |K^d(\vec{x})| d\vec{x} \\ &= h^d \left(\int_{-1}^1 |K(u)| du \right)^d = h^d \|K\|_1^d. \quad \blacksquare \end{aligned}$$

6.2. Bias Bound

The following lemma bounds the integrated square bias of \tilde{p}_h for an arbitrary $p \in \Sigma_{\kappa_1, \kappa_2}(\beta, L, r)$. We write the bias of \tilde{p}_h at $\vec{x} \in \mathcal{X}$ as $B_p(\vec{x}) = \mathbb{E}\tilde{p}_h(\vec{x}) - p(\vec{x})$.

Bias Lemma: There exists a constant $C > 0$ such that

$$\int_{\mathcal{X}} B_p^2(\vec{x}) d\vec{x} \leq Ch^{2\beta}. \quad (10)$$

We consider separately the interior $\mathcal{I} := (h, 1 - h)^d$ and boundary \mathcal{B} (noting $\mathcal{X} = \mathcal{I} \cup \mathcal{B}$). By a standard result for kernel density estimates of Hölder continuous functions (see, for example, Proposition 1.2 of [Tsybakov \(2008\)](#)),

$$\int_{\mathcal{I}} B_p^2(x) d\vec{x} \leq C_2 h^{2\beta}.$$

(In particular, this holds for the constant $C_2 := \frac{L}{\ell!} \|K\|_1^d$.)

We now show that $\int_{\mathcal{B}} B_p^2(\vec{x}) d\vec{x} \leq C_3^2 h^{2\beta}$.

Suppose $S = (S_1, S_2, S_3) \in \mathcal{S} \setminus \{(\emptyset, [d], \emptyset)\}$ (as $C_{(\emptyset, [d], \emptyset)} = \mathcal{I}$). We wish to bound $|B_p(\vec{x})|$ on C_S . To simplify notation, by geometric symmetry, we may assume $S_3 = \emptyset$. Let $\vec{u} \in [-1, 1]^d$, and define $\vec{y}_S \in \mathcal{X}$ by $(y_S)_i = hu_i - x_i, \forall i \in S_1$ and $(y_S)_i = x_i - hu_i, \forall i \in S_2$ (this choice arises from the change of variables we will use in (14)). By the Hölder condition (2) and the choice of y_S ,

$$\begin{aligned} \left| p(\vec{y}_S) - \sum_{|\vec{i}| \leq \ell} \frac{D^{\vec{i}} p(x)}{i!} (\vec{y}_S - \vec{x})^{\vec{i}} \right| &\leq L \|\vec{y}_S - \vec{x}\|_r^\beta \\ &= L \left(\sum_{j \in S_1} |2x_j + hu_j|^r + \sum_{j \in S_2} |hu_j|^r \right)^{\beta/r} \end{aligned}$$

Since each $|u_j| \leq 1$ and, for each $i \in S_1, 0 \leq x_j \leq h$,

$$\begin{aligned} \left| p(\vec{y}_S) - \sum_{|\alpha| \leq \ell} \frac{D^\alpha p(x)}{\alpha!} (\vec{y}_S - \vec{x})^\alpha \right| &= L \left(\sum_{j \in S_1} (3h)^r + \sum_{j \in S_2} h^r \right)^{\beta/r} \\ &\leq L (d(3h)^r)^{\beta/r} = L (3d^{1/r} h)^\beta. \end{aligned}$$

Rewriting this using the triangle inequality

$$\begin{aligned} |p(\vec{y}_S) - p(\vec{x})| &\leq L (3d^{1/r} h)^\beta + \left| \sum_{1 \leq |\alpha| \leq \ell} \frac{D^\alpha p(\vec{x})}{\alpha!} (\vec{y}_S - \vec{x})^\alpha \right|. \quad (11) \end{aligned}$$

Observing that $(\vec{y} - \vec{x})^{\vec{i}} \leq (3h)^{|\vec{i}|}$ and applying the bound

on the derivatives of p near $\partial\mathcal{X}$ (as computed in (8)),

$$\begin{aligned} &\left| \sum_{1 \leq |\vec{i}| \leq \ell} \frac{D^{\vec{i}} p(x)}{i!} (\vec{y}_S - \vec{x})^{\vec{i}} \right| \\ &\leq \left| \sum_{|\vec{i}| \leq \ell} \frac{Lh^{\beta - |\vec{i}|}}{(\ell - |\vec{i}|)! i!} (3h)^{|\vec{i}|} \right| \\ &= Lh^\beta \sum_{k=0}^{\ell} \sum_{|\vec{i}|=k} \frac{3^{|\vec{i}|}}{(\ell - k)! i!} \\ &\leq Lh^\beta \sum_{k=0}^{\ell} \frac{1}{k!(\ell - k)!} \sum_{|\vec{i}|=k} \frac{k! 3^{|\vec{i}|}}{i!} \end{aligned}$$

Then, applying the multinomial theorem (1) followed by the binomial theorem gives

$$\begin{aligned} \left| \sum_{1 \leq |\vec{i}| \leq \ell} \frac{D^{\vec{i}} p(x)}{i!} (\vec{y}_S - \vec{x})^{\vec{i}} \right| &\leq Lh^\beta \sum_{k=0}^{\ell} \frac{(3d)^k}{k!(\ell - k)!} \\ &= Lh^\beta \frac{1}{\ell!} \sum_{k=0}^{\ell} \frac{\ell!}{(\ell - k)! k!} (3d)^k \\ &= Lh^\beta \frac{(3d + 1)^\ell}{\ell!}. \end{aligned}$$

Combining this bound with (11) gives

$$|p(\vec{y}_S) - p(\vec{x})| = C_3 h^\beta, \quad (12)$$

where C_3 is the constant (in n and h)

$$C_3 := L \left((3d^{1/r})^\beta + \frac{(3d + 1)^\ell}{\ell!} \right). \quad (13)$$

For $x \in C_S$, we have $\tilde{p}_h(\vec{x}) = \frac{1}{nh^d} \sum_{i=1}^n K_S(\vec{x}, \vec{x}^i)$, and thus, by a change of variables, recalling that $K^d(\vec{x})$ denotes the product kernel,

$$\begin{aligned} \mathbb{E} \tilde{p}_h(\vec{x}) &= \frac{1}{h^d} \int_{\mathcal{X}} K_S(\vec{x}, \vec{u}) p(\vec{u}) d\vec{u} \\ &= \int_{[-1, 1]^d} K^d(\vec{v}) p(\vec{y}_S) d\vec{v}, \quad (14) \end{aligned}$$

Since $\int_{[-1, 1]^d} K^d(\vec{v}) d\vec{v} = 1$, by the bound in (12),

$$\begin{aligned} |B_p(x)| &= |\mathbb{E} \tilde{p}_h(x) - p(x)| \\ &= \left| \int_{[-1, 1]^d} K^d(\vec{v}) p(\vec{y}_S) d\vec{v} - \int_{[-1, 1]^d} K^d(\vec{v}) p(\vec{x}) d\vec{v} \right| \\ &\leq \int_{[-1, 1]^d} K^d(\vec{v}) |p(\vec{y}_S) - p(\vec{x})| d\vec{v} \\ &\leq \int_{[-1, 1]^d} K^d(\vec{v}) C_3 h^\beta d\vec{v} = C_3 h^\beta. \end{aligned}$$

Then, $\int_{\mathcal{B}} B_p^2(x) d\vec{x} \leq C_3^2 h^{2\beta}$ (since the measure of \mathcal{B} is less than 1), proving the Bias Lemma. ■

By Taylor's Theorem, $\forall \vec{x} \in \mathcal{X}$, for some $\xi : \mathcal{X} \rightarrow \mathbb{R}^2$ on the line segment between $(\hat{p}_h(\vec{x}), \hat{q}_h(\vec{x}))$ and $(p(\vec{x}), q(\vec{x}))$,

$$\begin{aligned} & |\mathbb{E}f(\hat{p}_h(\vec{x}), \hat{q}_h(\vec{x})) - f(p(\vec{x}), q(\vec{x}))| \\ &= \left| \mathbb{E} \frac{\partial f}{\partial x_1}(p(\vec{x}), q(\vec{x}))(\hat{p}_h(\vec{x}) - p(\vec{x})) \right. \\ &+ \frac{\partial f}{\partial x_2}(p(\vec{x}), q(\vec{x}))(\hat{q}_h(\vec{x}) - q(\vec{x})) \\ &+ \frac{1}{2} \left[\frac{\partial^2 f}{\partial x_1^2}(\xi)(\hat{p}_h(\vec{x}) - p(\vec{x}))^2 + \frac{\partial^2 f}{\partial x_2^2}(\xi)(\hat{q}_h(\vec{x}) - q(\vec{x}))^2 \right. \\ &\quad \left. + \frac{\partial^2 f}{\partial x_1 \partial x_2}(\xi)(\hat{p}_h(\vec{x}) - p(\vec{x}))(\hat{q}_h(\vec{x}) - q(\vec{x})) \right] \Big| \\ &\leq C_f (|B_p(\vec{x})| + |B_q(\vec{x})| + \mathbb{E}[\hat{p}_h(\vec{x}) - p(\vec{x})]^2 \\ &+ \mathbb{E}[\hat{q}_h(\vec{x}) - q(\vec{x})]^2 + |B_p(\vec{x})B_q(\vec{x})|), \end{aligned}$$

where the last line follows from the triangle inequality and (6). Thus, using (7),

$$\begin{aligned} & |\mathbb{E}D_\alpha(\hat{p}_h \parallel \hat{q}_h) - D_\alpha(p \parallel q)| \\ &= \left| \frac{1}{\alpha - 1} \left(\mathbb{E} \log \int_{\mathcal{X}} f(\hat{p}_h(\vec{x}), \hat{q}_h(\vec{x})) d\vec{x} \right. \right. \\ &\quad \left. \left. - \log \int_{\mathcal{X}} f(p(\vec{x}), q(\vec{x})) d\vec{x} \right) \right| \\ &\leq \frac{C_L}{|\alpha - 1|} \int_{\mathcal{X}} |\mathbb{E}f(\hat{p}_h(\vec{x}), \hat{q}_h(\vec{x})) - f(p(\vec{x}), q(\vec{x}))| d\vec{x} \\ &\leq \frac{C_f C_L}{|\alpha - 1|} \int_{\mathcal{X}} |B_p(\vec{x})| + |B_q(\vec{x})| + \mathbb{E}[\hat{p}_h(\vec{x}) - p(\vec{x})]^2 \\ &+ \mathbb{E}[\hat{q}_h(\vec{x}) - q(\vec{x})]^2 + |B_p(\vec{x})B_q(\vec{x})| d\vec{x}. \end{aligned}$$

By Hölder's Inequality, we then have

$$\begin{aligned} & |\mathbb{E}D_\alpha(\hat{p}_h \parallel \hat{q}_h) - D_\alpha(p \parallel q)| \\ &\leq \frac{C_f C_L}{|\alpha - 1| \kappa_1} \left(\sqrt{\int_{\mathcal{X}} B_p^2(\vec{x}) d\vec{x}} + \sqrt{\int_{\mathcal{X}} B_q^2(\vec{x}) d\vec{x}} \right. \\ &+ \int_{\mathcal{X}} \mathbb{E}[\hat{p}_h(\vec{x}) - p(\vec{x})]^2 + \mathbb{E}[\hat{q}_h(\vec{x}) - q(\vec{x})]^2 d\vec{x} \\ &\left. + \sqrt{\int_{\mathcal{X}} B_p^2(\vec{x}) \int_{\mathcal{X}} B_q^2(\vec{x}) d\vec{x}} \right). \end{aligned}$$

Applying Lemma 3.1 and a standard result in kernel density estimation (see, for example, Propositions 1.1 and 1.2 of

Tsybakov (2008)) gives

$$\begin{aligned} & |\mathbb{E}D_\alpha(\hat{p}_h \parallel \hat{q}_h) - D_\alpha(p \parallel q)| \\ &\leq (C_2 + C_3) h^\beta + C_2 h^{2\beta} + \kappa_2 \frac{\|K\|_1^d}{nh^d} \\ &\leq C \left(h^\beta + h^{2\beta} + \frac{1}{nh^d} \right), \end{aligned} \quad (15)$$

for some $C > 0$ not depending on n or h . ■

6.3. Variance Bound

Consider i.i.d. samples $\vec{x}^1, \dots, \vec{x}^n \sim p, \vec{y}^1, \dots, \vec{y}^n \sim q$. In anticipation of using McDiarmid's Inequality (McDiarmid, 1989), let $\tilde{p}'_h(\vec{x})$ denote our kernel density estimate with the sample \vec{x}^j replaced by $(\vec{x}^j)'$. By (7),

$$\begin{aligned} & |D_\alpha(\hat{p}_h \parallel \hat{q}_h) - D_\alpha(\tilde{p}'_h \parallel \tilde{q}'_h)| \\ &= \frac{1}{|\alpha - 1|} \left| \log \left(\int_{\mathcal{X}} f(\hat{p}_h(\vec{x}), \hat{q}_h(\vec{x})) d\vec{x} \right) \right. \\ &\quad \left. - \log \left(\int_{\mathcal{X}} f(\tilde{p}'_h(\vec{x}), \tilde{q}'_h(\vec{x})) d\vec{x} \right) \right| \\ &\leq \frac{C_L}{|\alpha - 1|} \int_{\mathcal{X}} |f(\hat{p}_h(\vec{x}), \hat{q}_h(\vec{x})) - f(\tilde{p}'_h(\vec{x}), \tilde{q}'_h(\vec{x}))| d\vec{x}. \end{aligned}$$

Then, applying the Mean Value Theorem followed by (6) gives, for some $\xi : \mathcal{X} \rightarrow \mathbb{R}^2$ on the line segment between (\hat{p}_h, \hat{q}_h) and (p, q) ,

$$\begin{aligned} & |D_\alpha(\hat{p}_h \parallel \hat{q}_h) - D_\alpha(\tilde{p}'_h \parallel \tilde{q}'_h)| \\ &\leq \frac{C_L}{|\alpha - 1|} \int_{\mathcal{X}} \left| \frac{\partial f}{\partial x_1}(\xi(\vec{x}))(\hat{p}_h(\vec{x}) - \tilde{p}'_h(\vec{x})) \right| dx \\ &\leq \frac{C_L C_f}{|\alpha - 1|} \int_{\mathcal{X}} |\hat{p}_h(\vec{x}) - \tilde{p}'_h(\vec{x})| d\vec{x}. \end{aligned}$$

Expanding \hat{p}_h as per its construction gives

$$\begin{aligned} & |D_\alpha(\hat{p}_h \parallel \hat{q}_h) - D_\alpha(\tilde{p}'_h \parallel \tilde{q}'_h)| \\ &\leq \frac{C_L C_f}{|\alpha - 1|} \int_{\mathcal{X}} |\tilde{p}_h(\vec{x}) - \tilde{p}'_h(\vec{x})| d\vec{x} \\ &\leq \frac{C_L C_f}{|\alpha - 1| nh^d} \sum_{S \in \mathcal{S}} \int_{\mathcal{X}} |K_S(x, x^j) - K_S(x, (x^j)')| d\vec{x} \\ &\leq \frac{2C_L C_f}{|\alpha - 1| nh^d} \sup_{y \in \mathcal{X}} \sum_{S \in \mathcal{S}} \int_{\mathcal{X}} |K_S(x, y)| d\vec{x} = \frac{2C_L C_f}{|\alpha - 1| n} \|K\|_1^d, \end{aligned}$$

where the last line follows from the triangle inequality and (9). An identical proof holds if we vary some \vec{y}^i rather than \vec{x}^i . Thus, since we have $2n$ independent samples, McDiarmid's Inequality gives the bound,

$$\mathbb{P}(|D_\alpha(\hat{p}_h, \hat{q}_h) - \mathbb{E}D_\alpha(\tilde{p}_h, \tilde{q}_h)| > \varepsilon) \leq 2 \exp\left(-\frac{C^2 \varepsilon^2 n}{\|K\|_1^{2d}}\right),$$

$$\text{where } C = \frac{|\alpha - 1|}{2C_L C_f} \quad (16)$$

depends only on κ and α . ■

7. Experiment

We used our estimator to estimate the Rényi α -divergence between two normal distributions in \mathbb{R}^3 restricted to the unit cube. In particular, for

$$\vec{\mu}_1 = \begin{bmatrix} 0.3 \\ 0.3 \\ 0.3 \end{bmatrix}, \vec{\mu}_2 = \begin{bmatrix} 0.7 \\ 0.7 \\ 0.7 \end{bmatrix}, \Sigma = \begin{bmatrix} 0.2 & 0 & 0 \\ 0 & 0.2 & 0 \\ 0 & 0 & 0.2 \end{bmatrix},$$

$p = \mathcal{N}(\vec{\mu}_1, \Sigma), q = \mathcal{N}(\vec{\mu}_2, \Sigma)$. For each $n \in \{1, 2, 5, 10, 50, 100, 500, 1000, 2000, 5000\}$, n data points were sampled according to each distribution and constrained (via rejection sampling) to lie within $[0, 1]^3$. Our estimator was computed from these samples, for $\alpha = 0.8$, using the Epanechnikov Kernel $K(u) = \frac{3}{4}(1 - u^2)$ on $[-1, 1]$, with constant bandwidth $h = 0.25$. The true α -divergence was computed directly according to its definition on the (renormalized) distributions on $[0, 1]^3$. The bias and variance of our estimator were then computed in the usual manner based on 100 trials. Figure 3 shows the error and variance of our estimator for each n .

We also compared our estimator's empirical error to our theoretical bound. Since the distributions used are infinitely differentiable, $\beta = \infty$, and so the estimator's mean squared error should converge as $O(n^{-1})$. An appropriate constant multiple was computed from (5), (13), and (15). The resulting bound is also shown in Figure 3.




Figure 3. Log-log plot of mean squared error (computed over 100 trials) of our estimator for various sample sizes n , alongside our theoretical bound. Error bars indicate standard deviation of estimator over 100 trials.

8. Conclusion

In this paper we derived a finite sample exponential concentration bound for a consistent, nonparametric, d -dimensional Rényi- α divergence estimator. To the best of our knowledge this is the first such exponential concentration bound for Rényi divergence.

References

- Adami, Christoph. Information theory in molecular biology. *Physics of Life Reviews*, 1:3–22, 2004.
- Aghagolzadeh, M., Soltanian-Zadeh, H., Araabi, B., and Aghagolzadeh, A. A hierarchical clustering based on mutual information maximization. In *in Proc. of IEEE International Conference on Image Processing*, pp. 277–280, 2007.
- Alemany, P. A. and Zanette, D. H. Fractal random walks from a variational formalism for Tsallis entropies. *Phys. Rev. E*, 49(2):R956–R958, Feb 1994. doi: 10.1103/PhysRevE.49.R956.
- Birge, L. and Massart, P. Estimation of integral functions of a density. *The Annals of Statistics*, 23:11–29, 1995.
- Borgwardt, Karsten M., Gretton, Arthur, Rasch, Malte J., Kriegel, Hans-Peter, Schlkopf, Bernhard, and Smola, Alex J. Integrating structured biological data by kernel maximum mean discrepancy. *Bioinformatics*, 22(14): e49–e57, 2006.
- Bregman, L. The relaxation method of finding the common points of convex sets and its application to the solution of problems in convex programming. *USSR Computational Mathematics and Mathematical Physics*, 7:79–86, 1967.
- Chai, Barry, Walther, Dirk B., Beck, Diane M., and Fei-Fei, Li. Exploring functional connectivity of the human brain using multivariate information analysis. In *NIPS*, 2009.
- Csiszár, I. Information-type measures of differences of probability distributions and indirect observations. *Studia Sci. Math. Hungarica*, 2:299–318, 1967.
- Goria, M. N., Leonenko, N. N., Mergel, V. V., and Inverardi, P. L. Novi. A new class of random vector entropy estimators and its applications in testing statistical hypotheses. *Journal of Nonparametric Statistics*, 17:277–297, 2005.
- Gupta, M. and Srivastava, S. Parametric bayesian estimation of differential entropy and relative entropy. *Entropy*, 12:818–843, 2010.
- Hero, A. O., Ma, B., Michel, O., and Gorman, J. Alpha-divergence for classification, indexing and retrieval, 2002a. Communications and Signal Processing Laboratory Technical Report CSPL-328.
- Hero, A. O., Ma, B., Michel, O. J. J., and Gorman, J. Applications of entropic spanning graphs. *IEEE Signal Processing Magazine*, 19(5):85–95, 2002b.

- Kozachenko, L. F. and Leonenko, N. N. A statistical estimate for the entropy of a random vector. *Problems of Information Transmission*, 23:9–16, 1987.
- Kullback, S. and Leibler, R.A. On information and sufficiency. *Annals of Mathematical Statistics*, 22:79–86, 1951.
- Kybic, J. Incremental updating of nearest neighbor-based high-dimensional entropy estimation. In *Proc. Acoustics, Speech and Signal Processing*, 2006.
- Laurent, B. Efficient estimation of integral functionals of a density. *The Annals of Statistics*, 24:659–681, 1996.
- Learned-Miller, Erik G. and Fisher, John W. ICA using spacings estimates of entropy. *Journal of Machine Learning Research*, 4:1271–1295, 2003.
- Leonenko, N. and Pronzato, L. Correction of ‘a class of Rényi information estimators for multidimensional densities’ *Ann. Statist.*, 36(2008) 2153–2182, 2010.
- Leonenko, Nikolai, Pronzato, Luc, and Savani, Vippal. A class of Rényi information estimators for multidimensional densities. *Annals of Statistics*, 36(5):2153–2182, 2008.
- Lewi, J., Butera, R., and Paninski, L. Real-time adaptive information-theoretic optimization of neurophysiology experiments. In *Advances in Neural Information Processing Systems*, volume 19, 2007.
- Liu, H., Lafferty, J., and Wasserman, L. Exponential concentration inequality for mutual information estimation. In *Neural Information Processing Systems (NIPS)*, 2012.
- McDiarmid, Colin. On the method of bounded differences. *Surveys in Combinatorics*, 141:148–188, 1989.
- Nguyen, X., Wainwright, M.J., and Jordan, M.I. Estimating divergence functionals and the likelihood ratio by convex risk minimization. *IEEE Transactions on Information Theory*, To appear, 2010.
- Oliva, J., Póczos, B., and Schneider, J. Distribution to distribution regression. In *International Conference on Machine Learning (ICML)*, 2013.
- Pál, D., Póczos, B., and Szepesvári, Cs. Estimation of Rényi entropy and mutual information based on generalized nearest-neighbor graphs. In *Proceedings of the Neural Information Processing Systems*, 2010.
- Peng, H. and Dind, C. Feature selection based on mutual information: Criteria of max-dependency, max-relevance, and min-redundancy. *IEEE Trans On Pattern Analysis and Machine Intelligence*, 27, 2005.
- Pérez-Cruz, F. Estimation of information theoretic measures for continuous random variables. In *Advances in Neural Information Processing Systems 21*, 2008.
- Póczos, B. and Schneider, J. On the estimation of alpha-divergences. In *International Conference on AI and Statistics (AISTATS)*, volume 15 of *JMLR Workshop and Conference Proceedings*, pp. 609–617, 2011.
- Póczos, B., Xiong, L., Sutherland, D., and Schneider, J. Nonparametric kernel estimators for image classification. In *25th IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- Rényi, A. On measures of entropy and information. In *Fourth Berkeley Symposium on Mathematical Statistics and Probability*. University of California Press, 1961.
- Rényi, A. *Probability Theory*. North-Holland Publishing Company, Amsterdam, 1970.
- Shan, Caifeng, Gong, Shaogang, and Mcowan, Peter W. Conditional mutual information based boosting for facial expression recognition. In *British Machine Vision Conference (BMVC)*, 2005.
- Sricharan, K., Raich, R., and Hero, A. Empirical estimation of entropy functionals with confidence. Technical Report, <http://arxiv.org/abs/1012.4188>, 2010.
- Sricharan, K., Wei, D., and Hero, A. Ensemble estimators for multivariate entropy estimation, 2012. <http://arxiv.org/abs/1203.5829>.
- Szabó, Z., Póczos, B., and Lőrincz, A. Undercomplete blind subspace deconvolution. *Journal of Machine Learning Research*, 8:1063–1095, 2007.
- Tsybakov, Alexandre B. *Introduction to Nonparametric Estimation*. Springer Publishing Company, Incorporated, 1st edition, 2008. ISBN 0387790519, 9780387790510.
- Villmann, T. and Haase, S. Mathematical aspects of divergence based vector quantization using Frechet-derivatives, 2010. University of Applied Sciences Mitweida.
- Wang, Qinq, Kulkarni, Sanjeev R., and Verdú, Sergio. Divergence estimation for multidimensional densities via k -nearest-neighbor distances. *IEEE Transactions on Information Theory*, 55(5), 2009.
- Wolsztynski, Eric, Thierry, Eric, and Pronzato, Luc. Minimum-entropy estimation in semi-parametric models. *Signal Process.*, 85(5):937–949, 2005. ISSN 0165-1684. doi: <http://dx.doi.org/10.1016/j.sigpro.2004.11.028>.