

Article

Estimating Mixture Entropy with Pairwise Distances

Artemy Kolchinsky ^{1,*} and Brendan D. Tracey ^{1,2}

¹ Santa Fe Institute, Santa Fe, NM 87501, USA; tracey.brendan@gmail.com

² Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

* Correspondence: artemy@santafe.edu; Tel.: +1-(505)-984-8800

Received: 8 June 2017; Accepted: 12 July 2017; Published: 14 July 2017

Abstract: Mixture distributions arise in many parametric and non-parametric settings—for example, in Gaussian mixture models and in non-parametric estimation. It is often necessary to compute the entropy of a mixture, but, in most cases, this quantity has no closed-form expression, making some form of approximation necessary. We propose a family of estimators based on a pairwise distance function between mixture components, and show that this estimator class has many attractive properties. For many distributions of interest, the proposed estimators are efficient to compute, differentiable in the mixture parameters, and become exact when the mixture components are clustered. We prove this family includes lower and upper bounds on the mixture entropy. The Chernoff α -divergence gives a lower bound when chosen as the distance function, with the Bhattacharyaa distance providing the tightest lower bound for components that are symmetric and members of a location family. The Kullback–Leibler divergence gives an upper bound when used as the distance function. We provide closed-form expressions of these bounds for mixtures of Gaussians, and discuss their applications to the estimation of mutual information. We then demonstrate that our bounds are significantly tighter than well-known existing bounds using numeric simulations. This estimator class is very useful in optimization problems involving maximization/minimization of entropy and mutual information, such as MaxEnt and rate distortion problems.

Keywords: mixture distribution; mixture of Gaussians; entropy estimation; MaxEnt; rate distortion

1. Introduction

A mixture distribution is a probability distribution whose density function is a weighted sum of individual densities. Mixture distributions are a common choice for modeling probability distributions, in both parametric settings, for example, learning a mixture of Gaussians statistical model [1], and non-parametric settings, such as kernel density estimation.

It is often necessary to compute the differential entropy [2] of a random variable with a mixture distribution, which is a measure of the inherent uncertainty in the outcome of the random variable. Entropy estimation arises in image retrieval tasks [3], image alignment and error correction [4], speech recognition [5,6], analysis of debris spread in rocket launch failures [7], and many other settings. Entropy also arises in optimization contexts [4,8–10], where it is minimized or maximized under some constraints (e.g., MaxEnt problems). Finally, entropy also plays a central role in minimization or maximization of mutual information, such as in problems related to rate distortion [11].

Unfortunately, in most cases, the entropy of a mixture distribution has no known closed-form expression [12]. This is true even when the entropy of each component distribution does have a known closed-form expression. For instance, the entropy of a Gaussian has a well-known form, while the entropy of a mixture of Gaussians does not [13]. As a result, the problem of finding a tractable and accurate estimate for mixture entropy has been described as “a problem of considerable current interest and practical significance” [14].

One way to approximate mixture entropy is with Monte Carlo (MC) sampling. MC sampling provides an unbiased estimate of the entropy, and this estimate can become arbitrarily accurate by increasing the number of MC samples. Unfortunately, MC sampling is very computationally intensive, as, for each sample, the (log) probability of the sample location must be computed under every component in the mixture. MC sampling typically requires a large number of samples to estimate entropy, especially in high-dimensions. Sampling is thus typically impractical, especially for optimization problems where, for every parameter change, a new entropy estimate is required. Alternatively, it is possible to approximate entropy using numerical integration, but this is also computationally expensive and limited to low-dimensional applications [15,16].

Instead of Monte Carlo sampling or numerical integration, one may use an analytic estimator of mixture entropy. Analytic estimators have estimation bias but are much more computationally efficient. There are several existing analytic estimators of entropy, discussed in-depth below. To summarize, however, commonly-used estimators have significant drawbacks: they have large bias relative to the true entropy, and/or they are invariant to the amount of “overlap” between mixture components. For example, many estimators do not depend on the locations of the means in a Gaussian mixture model.

In this paper, we introduce a novel family of estimators for the mixture entropy. Each member of this family is defined via a pairwise-distance function between component densities. The estimators in this family have several attractive properties. They are computationally efficient, as long as the pairwise-distance function and the entropy of each component distribution are easy to compute. The estimation bias of any member of this family is bounded by a constant. The estimator is continuous and smooth and is therefore useful for optimization problems. In addition, we show that when the Chernoff α -divergence (i.e., a scaled Rényi divergence) is used as a pairwise-distance function, the corresponding estimator is a lower-bound on the mixture entropy. Furthermore, among all the Chernoff α -divergences, the Bhattacharyya distance ($\alpha = 0.5$) provides the tightest lower bound when the mixture components are symmetric and belong to a location family (such as a mixture of Gaussians with equal covariances). We also show that when the Kullback–Leibler [KL] divergence is used as a pairwise-distance function, the corresponding estimator is an upper-bound on the mixture entropy. Finally, our family of estimators can compute the exact mixture entropy when the component distributions are grouped into well-separated clusters, a property not shared by other analytic estimators of entropy. In particular, the bounds mentioned above converge to the same value for well-separated clusters.

The paper is laid out as follows. We first review mixture distributions and entropy estimation in Section 2. We then present the class of pairwise distance estimators in Section 3, prove bounds on the error of any estimator in this class, and show distance functions that bound the entropy as discussed above. In Section 4, we consider the special case of mixtures of Gaussians, and give explicit expressions for lower and upper bounds on the mixture entropy. When all the Gaussian components have the same covariance matrix, we show that these bounds have particularly simple expressions. In Section 5, we consider the closely related problem of estimating the mutual information between two random variables, and show that our estimators can be directly used to estimate and bound the mutual information. For the Gaussian case, these can be used to bound the mutual information across a type of Additive White Noise Gaussian channel. Finally, in Section 6, we run numerical experiments and compare the performance of our lower and upper bounds relative to existing estimators. We consider both mixtures of Gaussians and mixtures of uniform distributions.

2. Background and Definitions

We consider the differential entropy of a continuous random variable X , defined as

$$H(X) := - \int p_X(x) \ln p_X(x) dx,$$

where p_X is a mixture distribution,

$$p_X(x) = \sum_{i=1}^N c_i p_i(x),$$

and where c_i indicates the weight of component i ($c_i \geq 0$, $\sum_i c_i = 1$) and p_i the probability density of component i .

We can treat the set of component weights as the probabilities of outcomes $1 \dots N$ of a discrete random variable C , where $\Pr(C = i) = c_i$. Consider the mixed joint distribution of the discrete random variable C and the continuous random variable X ,

$$p_{X,C}(x, i) = p_i(x)c_i,$$

and note the following identities for conditional and joint entropy [17],

$$H(X, C) = H(X|C) + H(C) = H(C|X) + H(X),$$

where we use H for discrete and differential entropy interchangeably. Here, the conditional entropies are defined as

$$H(X|C) = \sum_i c_i H(p_i) \quad \text{and} \quad H(C|X) = \int p_{X,C}(x, i) \log \frac{p_{X,C}(x, i)}{p_X(x)} dx.$$

Using elementary results from information theory [2], $H(X)$ can be bounded from below by

$$H(X) \geq H(X|C), \quad (1)$$

since conditioning can only decrease entropy. Similarly, $H(X)$ can be bounded from above by

$$H(X) \leq H(X, C) = H(X|C) + H(C), \quad (2)$$

following from $H(X) = H(X, C) - H(C|X)$ and the non-negativity of the conditional discrete entropy $H(C|X)$. This upper bound on the mixture entropy was previously proposed by Huber et al. [18].

It is easy to see that the bound in Equation (1) is tight when all the components have the same distribution, since then $H(p_X) = H(p_i)$ for all i . The bound in Equation (2) becomes tight when $H(C|X) = 0$, i.e., when any sample from p_X uniquely determines the component identity C . This occurs when the different mixture components have non-overlapping supports, $p_i(x) > 0 \implies p_j(x) = 0$ for all x and $i \neq j$. More generally, the bound of Equation (2) becomes increasingly tight as the mixture distributions move farther apart from one another.

In the case where the entropy of each component density, $H(p_i)$ for $i = 1 \dots N$, has a simple closed form expression, the bounds in Equations (1) and (2) can be easily computed. However, neither bound depends on the ‘‘overlap’’ between components. For instance, in a Gaussian mixture model, these bounds are invariant to changes in the component means. The bounds are thus unsuitable for many problems; for instance, in optimization, one typically tunes parameters to adjust component means, but the above entropy bounds remain the same regardless of mean location.

There are two other estimators of the mixture entropy that should be mentioned. The first estimator is based on kernel density estimation [16,19]. It estimates the entropy using the mixture probability of the component means, μ_i ,

$$\hat{H}_{\text{KDE}}(X) := - \sum_i c_i \ln \sum_j c_j p_j(\mu_i). \quad (3)$$

The second estimator is a lower bound that is derived using Jensen’s inequality [2], $-E[\ln f(X)] \geq -\ln[E(f(X))]$, giving

$$H(X) := - \int \sum_i c_i p_i(x) \ln \sum_j c_j p_j(x) dx \geq - \sum_i c_i \ln \sum_j c_j \left(\int p_i(x) p_j(x) dx \right) =: \hat{H}_{\text{ELK}}. \quad (4)$$

In the literature, the term $\int p_i(x)p_j(x) dx$ has been referred to as the “Cross Information Potential” [20,21] and the “Expected Likelihood Kernel” [22,23] (ELK, we use this second acronym to label this estimator). When the component distributions are Gaussian, $p_i := \mathcal{N}(\mu_i, \Sigma_i)$, the ELK has a simple closed-form expression,

$$\hat{H}_{\text{ELK}}(X) = - \sum_i c_i \ln \sum_j c_j q_{j,i}(\mu_i), \quad (5)$$

where each $q_{j,i}$ is a Gaussian defined as $q_{j,i} := \mathcal{N}(\mu_j, \Sigma_i + \Sigma_j)$. This lower bound was previously proposed for Gaussian mixtures in [18] and in a more general context in [12].

Both \hat{H}_{KDE} , Equation (3), and \hat{H}_{ELK} , Equation (5), are computationally efficient, continuous and differentiable, and depend on component overlap, making them suitable for optimization. However, as will be shown via numerical experiments (Section 6), they exhibit significant underestimation bias. At the same time, we will show that for Gaussian mixtures with equal covariance, \hat{H}_{KDE} is only an additive constant away from an estimator in our proposed class.

3. Estimators Based on Pairwise-Distances

3.1. Overview

Let $D(p_i \| p_j)$ be some (generalized) distance function between probability densities p_i and p_j . Formally, we assume that D is a *premetric*, meaning that it is non-negative and $D(p_i \| p_j) = 0$ if $p_i = p_j$. We do not assume that D is symmetric, nor that it obeys the triangle inequality, nor that it is strictly greater than 0 when $p_i \neq p_j$.

For any allowable distance function D , we propose the following entropy estimator:

$$\hat{H}_D(X) := H(X|C) - \sum_i c_i \ln \sum_j c_j \exp(-D(p_i \| p_j)). \quad (6)$$

This estimator can be efficiently computed if the entropy of each component and $D(p_i \| p_j)$ for all i, j have simple closed-form expressions. There are many distribution-distance function pairs that satisfy these conditions (e.g., Kullback–Leibler divergence, Renyi divergences, Bregman divergences, f-divergences, etc., for Gaussian, uniform, exponential, etc.) [24–28].

It is straightforward to show that for any D , \hat{H}_D falls between the bounds of Equations (1) and (2),

$$H(X|C) \leq \hat{H}_D(X) \leq H(X, C). \quad (7)$$

To do so, consider the “smallest” and “largest” allowable distance functions,

$$D_{\min}(p_i \| p_j) = 0 \quad \text{and} \quad D_{\max}(p_i \| p_j) = \begin{cases} 0, & \text{if } p_i = p_j, \\ \infty, & \text{otherwise.} \end{cases} \quad (8)$$

For any D and p_i, p_j , $D_{\min}(p_i \| p_j) \leq D(p_i \| p_j) \leq D_{\max}(p_i \| p_j)$, thus

$$\hat{H}_{D_{\min}}(X) \leq \hat{H}_D(X) \leq \hat{H}_{D_{\max}}(X).$$

Plugging D_{\min} into Equation (6) (noting that $\sum_j c_j = 1$) gives $\hat{H}_{D_{\min}}(X) = H(X|C)$, while plugging D_{\max} into Equation (6) gives

$$\hat{H}_{D_{\max}} = H(X|C) - \sum_i c_i \ln \left(c_i + \sum_{j \neq i} c_j e^{-D_{\max}(p_i \| p_j)} \right) \leq H(X|C) - \sum_i c_i \ln c_i = H(X, C).$$

These two inequalities yield Equation (7). The true entropy, as shown in Section 2, also obeys $H(X|C) \leq H(X) \leq H(X, C)$. The magnitude of the bias of \hat{H}_D is thus bounded by

$$|\hat{H}_D(X) - H(X)| \leq H(X, C) - H(X|C) = H(C).$$

In the next two subsections, we improve upon the bounds suggested in Equations (1) and (2), by examining bounds induced by particular distance functions.

3.2. Lower Bound

The ‘‘Chernoff α -divergence’’ [28,29] for some real-valued α is defined as

$$C_\alpha(p||q) := -\ln \int p^\alpha(x)q^{1-\alpha}(x) dx. \tag{9}$$

Note that $C_\alpha(p||q) = (1 - \alpha)R_\alpha(p||q)$, where R_α is Rényi divergence of order α [30].

We show that for any $\alpha \in [0, 1]$, $\hat{H}_{C_\alpha}(X)$ is a lower bound on the entropy (for $\alpha \notin [0, 1]$, C_α is not a valid distance function (see Appendix A)). To do so, we make use of a derivation from [31] and write,

$$\begin{aligned} H(X) &= H(X|C) - \int \sum_i c_i p_i(x) \ln \frac{p_X(x)}{p_i(x)} dx \\ &= H(X|C) - \int \sum_i c_i p_i(x) \ln \frac{p_X(x)}{p_i(x)^\alpha \sum_j c_j p_j(x)^{1-\alpha}} dx - \int \sum_i c_i p_i(x) \ln \left(\frac{\sum_j c_j p_j(x)^{1-\alpha}}{p_i(x)^{1-\alpha}} \right) dx \\ &\stackrel{(a)}{\geq} H(X|C) - \int \sum_i c_i p_i(x) \ln \left(\frac{\sum_j c_j p_j(x)^{1-\alpha}}{p_i(x)^{1-\alpha}} \right) dx \\ &\stackrel{(b)}{\geq} H(X|C) - \sum_i c_i \ln \int p_i(x)^\alpha \sum_j c_j p_j(x)^{1-\alpha} dx \\ &= H(X|C) - \sum_i c_i \ln \sum_j c_j e^{-C_\alpha(p_i||p_j)} := \hat{H}_{C_\alpha}(X). \end{aligned} \tag{10}$$

The inequalities (a) and (b) follow from Jensen’s inequality. This inequality is used directly in (b), while in (a) it follows from

$$\begin{aligned} - \int \sum_i c_i p_i(x) \ln \frac{p_X(x)}{p_i(x)^\alpha \sum_j c_j p_j(x)^{1-\alpha}} dx &\geq -\ln \int \sum_i c_i p_i(x) \frac{p_X(x)}{p_i(x)^\alpha \sum_j c_j p_j(x)^{1-\alpha}} dx \\ &= -\ln \int \sum_i c_i p_i(x)^{1-\alpha} \frac{p_X(x)}{\sum_j c_j p_j(x)^{1-\alpha}} dx \\ &= -\ln \sum_i c_i \int p_X(x) dx = 0. \end{aligned}$$

Note that Jensen’s inequality is used in the derivations of both this lower bound as well as the lower bound \hat{H}_{ELK} in Equation (4). However, the inequality is applied differently in the two cases, and, as will be demonstrated in Section 6, the estimators have different performance.

We have shown that using C_α as a distance function gives a lower bound on the mixture entropy for any $\alpha \in [0, 1]$. For a general mixture distribution, one could optimize over the value of α to find the tightest lower bound. However, we can show that the tightest bound is achieved for $\alpha = 0.5$ in the special case when all of the mixture components p_i are symmetric and come from a location family,

$$p_i(x) = b(x - \mu_i) = b(\mu_i - x).$$

Examples of this situation include mixtures of Gaussians with the same covariance (‘‘homoscedastic’’ mixtures), multivariate t -distributions with the same covariance, location-shifted bounded uniform distributions, most kernels used in kernel density estimation, etc. It does not apply to skewed distributions, such as the skew-normal distribution [12].

To show that $\alpha = 0.5$ is optimal, first define the Chernoff α -coefficient as

$$c_\alpha(p_i||p_j) := \int p_i(x)^\alpha p_j(x)^{1-\alpha} dx.$$

We show that for any pair p_i, p_j of symmetric distributions from a location family, $c_\alpha(p_i||p_j)$ is minimized by $\alpha = 0.5$. This means that all pairwise distances $C_\alpha(p_j||p_i) \equiv -\ln c_\alpha(p_j||p_i)$ are maximized by $\alpha = 0.5$, and, therefore, the entropy estimator H_{C_α} (Equation (6)) is maximized by $\alpha = 0.5$.

First, define a change of variables

$$y := \mu_i + \mu_j - x,$$

which gives $x - \mu_i = \mu_j - y$ and $x - \mu_j = \mu_i - y$. This allows us to write the Chernoff α -coefficient as

$$\begin{aligned} c_\alpha(p_i||p_j) &= \int p_i(x)^\alpha p_j(x)^{1-\alpha} dx \\ &= \int b(x - \mu_i)^\alpha b(x - \mu_j)^{1-\alpha} dx \\ &\stackrel{(a)}{=} \int b(\mu_j - y)^\alpha b(\mu_i - y)^{1-\alpha} dy \\ &\stackrel{(b)}{=} \int b(y - \mu_j)^\alpha b(y - \mu_i)^{1-\alpha} dy \\ &= c_{1-\alpha}(p_i||p_j), \end{aligned}$$

where, in (a), we have substituted variables, and in (b) we used the assumption that $b(x) = b(-x)$. Since we have shown that $c_\alpha(p_i||p_j) = c_{1-\alpha}(p_i||p_j)$, c_α is symmetric in α about $\alpha = 0.5$. In Appendix A, we show that $c_\alpha(p||q)$ is everywhere convex in α . Together, this means that $c_\alpha(p_i||p_j)$ must achieve a minimum value at $\alpha = 0.5$.

The Chernoff α -coefficient for $\alpha = 0.5$ is known as the *Bhattacharyya coefficient*, with the corresponding *Bhattacharyya distance* [32] defined as

$$BD(p||q) := -\ln \int \sqrt{p(x)q(x)} dx = C_{0.5}(p||q).$$

Since any Chernoff α -divergence is a lower bound for the entropy, we write the particular case of Bhattacharyya-distance lower bound as

$$H(X) \geq \hat{H}_{BD}(X) := H(X|C) - \sum_i c_i \ln \sum_j c_j e^{-BD(p_i||p_j)}. \tag{11}$$

3.3. Upper Bound

The *Kullback–Leibler [KL] divergence* [2] is defined as

$$KL(p||q) := \int p(x) \ln \frac{p(x)}{q(x)} dx.$$

Using KL divergence as the pairwise distance provides an upper bound on the mixture entropy. We show this as follows:

$$\begin{aligned} H(X) &= -\sum_i c_i E_{p_i} \left[\ln \sum_j c_j p_j(X) \right] \\ &\stackrel{(a)}{\leq} -\sum_i c_i \ln \sum_j c_j e^{E_{p_i}[\ln p_j(X)]} \end{aligned}$$

$$\begin{aligned}
 &= - \sum_i c_i \ln \sum_j c_j e^{-H(p_i \| p_j)} \\
 &= \sum_i c_i H(p_i) - \sum_i c_i \ln \sum_j c_j e^{-\text{KL}(p_i \| p_j)},
 \end{aligned}$$

where E_{p_i} indicates expectation when X is distributed according to p_i , $H(\cdot \| \cdot)$ indicates the cross-entropy function, and we employ the identity $H(p_i \| p_j) = H(p_i) + \text{KL}(p_i \| p_j)$. The inequality in step (a) uses a variational lower bound on the expectation of a log-sum [5,33],

$$E \left[\ln \sum_j Z_j \right] \geq \ln \sum_j \exp (E [\ln Z_j]) .$$

Combining yields the upper bound

$$H(X) \leq \hat{H}_{\text{KL}} := H(X|C) - \sum_i c_i \ln \sum_j c_j e^{-\text{KL}(p_i \| p_j)} . \tag{12}$$

3.4. Exact Estimation in the “Clustered” Case

In the previous sections, we derived lower and upper bounds on the mixture entropy, using estimators based on Chernoff α -divergence and KL divergence, respectively.

There are situations in which the lower and upper bounds become similar. Consider a pair of component distributions, p_i and p_j . By applying Jensen’s inequality to Equation (9), we can derive the inequality $C_\alpha(p_i \| p_j) \leq \alpha \text{KL}(p_i \| p_j)$. There are two cases in which a pair of components contributes similarly to the lower and upper bounds. The first case is when $C_\alpha(p_i \| p_j)$ is very large, meaning that the KL is also very large. By Equation (6), distances enter into our estimators as $\exp(-D(p_i \| p_j))$, and, in this case, $\exp(-\text{KL}(p_i \| p_j)) \approx \exp(-C_\alpha(p_i \| p_j)) \approx 0$. In the second case, $\text{KL}(p_i \| p_j) \approx 0$, meaning that $C_\alpha(p_i \| p_j)$ must also be near zero, and, in this case, $\exp(-\text{KL}(p_i \| p_j)) \approx \exp(-C_\alpha(p_i \| p_j)) \approx 1$. Thus, the lower and upper bounds become similar when all pairs of components are either very close together or very far apart.

In this section, we analyze this special case. Specifically, we consider the situation when mixture components are “clustered”, meaning that there is a grouping of component distributions such that distributions in the same group are approximately the same and distributions assigned to different groups are very different from one another. We show that in this case our lower and upper bounds become equal and our pairwise-distance estimate of the entropy is tight. Though this situation may seem like an edge case, clustered distributions do arise in mixture estimation, e.g., when there are repeated data points, or as solutions to information-theoretic optimization problems [11]. Note that the number of groups is arbitrary, and therefore this situation includes the extreme cases of a single group (all component distributions are nearly the same) as well as N different groups (all component distributions are very different).

Formally, let the function $g(i)$ indicate the group of component i . We define that the components are “clustered” with respect to grouping g iff $\text{KL}(p_i \| p_j) \leq \kappa$ whenever $g(i) = g(j)$ for some small κ , and $\text{BD}(p_i \| p_j) \geq \beta$ whenever $g(i) \neq g(j)$ some large β . We use the notation $p_G(k) = \sum_i \delta_{g(i),k} c_i$ to indicate the sum of the weights of the components in group k , where δ_{ij} indicates the Kronecker delta function. For technical reasons, below we only consider C_α where α is strictly greater than 0.

We show that when κ is small and β is large, both \hat{H}_{C_α} for $\alpha \in (0, 1]$ and \hat{H}_{KL} approach

$$H(X|C) - \sum_k p_G(k) \ln p_G(k) .$$

Since one is a lower bound and one is an upper bound on the true entropy, the estimators become exact as they converge in value.

Recall that $BD(p||q) = C_{0.5}(p||q)$. For $\alpha \in (0, 1]$, $\alpha^{-1}C_\alpha(p||q)$ is a monotonically decreasing function for $\alpha \in (0, 1]$ [34], meaning that $C_\alpha \geq 2\alpha BD(p||q)$ for $\alpha \in (0, 0.5]$. In addition, $(1 - \alpha)^{-1}C_\alpha(p||q)$ is a monotonically increasing function for $\alpha > 0$ [34], thus $C_\alpha \geq 2(1 - \alpha)BD(p||q)$ for $\alpha \in [0.5, 1]$. Using the assumption that $BD(p_i||p_j) \geq \beta$ and combining gives the bound

$$C_\alpha(p||q) \geq (1 - |1 - 2\alpha|)\beta$$

for $\alpha \in (0, 1]$, leading to

$$\begin{aligned} \hat{H}_{C_\alpha}(X) &:= H(X|C) - \sum_i c_i \ln \sum_j c_j \exp(-C_\alpha(p_i||p_j)) \\ &\geq H(X|C) - \sum_i c_i \ln \left[\sum_j \delta_{g(i),g(j)} c_j + \sum_j (1 - \delta_{g(i),g(j)}) c_j \exp(-C_\alpha(p_i||p_j)) \right] \\ &\geq H(X|C) - \sum_k p_G(k) \ln [p_G(k) + (1 - p_G(k)) \exp(-(1 - |1 - 2\alpha|)\beta)] . \end{aligned}$$

In the second line, for the summation over i, j in the same group, we've used the non-negativity of $C_\alpha(p_i||p_j) \geq 0$.

For the upper bound \hat{H}_{KL} , we use that $KL(p_i||p_j) \leq \kappa$ for i and j in the same group, and otherwise $\exp(-KL(p_i||p_j)) \geq 0$. This gives the bound

$$\begin{aligned} \hat{H}_{KL}(X) &:= H(X|C) - \sum_i c_i \ln \sum_j c_j \exp(-KL(p_i||p_j)) \\ &\leq H(X|C) - \sum_i c_i \ln \left[\sum_j \delta_{g(i),g(j)} c_j e^{-\kappa} + \sum_j (1 - \delta_{g(i),g(j)}) c_j \exp(-KL(p_i||p_j)) \right] \\ &\leq H(X|C) - \sum_k p_G(k) \ln p_G(k) e^{-\kappa} . \end{aligned}$$

The difference between the bounds is bounded by

$$\begin{aligned} \hat{H}_{KL}(X) - \hat{H}_{C_\alpha}(X) &\leq \kappa + \sum_k p_G(k) \ln \left[1 + \frac{(1 - p_G(k)) \exp(-(1 - |1 - 2\alpha|)\beta)}{p_G(k)} \right] \\ &\leq \kappa + \sum_k p_G(k) \frac{(1 - p_G(k)) \exp(-(1 - |1 - 2\alpha|)\beta)}{p_G(k)} \\ &= \kappa + (|G| - 1) \exp(-(1 - |1 - 2\alpha|)\beta) , \end{aligned}$$

where $|G|$ is the number of groups. Thus, the difference decreases at least linearly in κ and exponentially in β . This shows that, in the clustered case, when $\kappa \approx 0$ and β is very large, our lower and upper bounds become exact.

It also shows that any distance measure bounded between BD and KL also gives an exact estimate of entropy in the clustered case. Furthermore, the idea behind this proof can be extended to estimators induced by other bounding distances, beyond BD and KL, so as to show that a particular estimator converges to an exact entropy estimate in the clustered case. Note, however, that, for some distribution-distance pairs, the components will never be considered as "clustered"; e.g., the α -Chernoff distance for $\alpha = 0$ between any two Gaussians is 0, and so a Gaussian mixture distribution will never be considered clustered according to this distance.

Finally, in the perfectly clustered case, we can show that our lower bound, \hat{H}_{BD} , is at least as good as the Expected Likelihood Kernel lower bound, \hat{H}_{ELK} , as defined in Equation (4). See Appendix B for details.

4. Gaussian Mixtures

Gaussians are very frequently used as components in mixture distributions. Our family of estimators is well-suited to estimating the entropies of Gaussian mixtures, since the entropy of a d -dimensional Gaussian $p_i = \mathcal{N}(\mu_i, \Sigma_i)$ has a simple closed-form expression,

$$H(p_i) = \frac{1}{2} [\ln |\Sigma_i| + d \ln 2\pi + d], \quad (13)$$

and because there are many distance functions between Gaussians with closed-form expressions (KL divergence, the Chernoff α -divergences [35], 2-Wasserstein distance [36,37], etc.). In this section, we consider Gaussian mixtures and state explicit expression for the lower and upper bounds on the mixture entropy derived in the previous section. We also consider these bounds in the special case where all Gaussian components have the same covariance matrix (homoscedastic mixtures).

We first consider the lower bound, \hat{H}_{C_α} , based on the Chernoff α -divergence distance function. For two multivariate Gaussians $p_1 \sim \mathcal{N}(\mu_1, \Sigma_1)$ and $p_2 \sim \mathcal{N}(\mu_2, \Sigma_2)$, this distance is defined as [35]:

$$C_\alpha(p_1 \| p_2) = \frac{(1-\alpha)\alpha}{2} (\mu_1 - \mu_2)^T ((1-\alpha)\Sigma_1 + \alpha\Sigma_2)^{-1} (\mu_1 - \mu_2) + \frac{1}{2} \ln \left(\frac{|(1-\alpha)\Sigma_1 + \alpha\Sigma_2|}{|\Sigma_1|^{1-\alpha} |\Sigma_2|^\alpha} \right). \quad (14)$$

(As a warning, note that most sources show erroneous expressions for the Chernoff and/or Rényi α -divergence between two multivariate Gaussians, including [27,29,38–40], and even a late draft of this manuscript.)

For the upper bound \hat{H}_{KL} , the KL divergence between two multivariate Gaussians $p_1 \sim \mathcal{N}(\mu_1, \Sigma_1)$ and $p_2 \sim \mathcal{N}(\mu_2, \Sigma_2)$ is

$$\text{KL}(p_1 \| p_2) = \frac{1}{2} [\ln |\Sigma_2| - \ln |\Sigma_1| + (\mu_1 - \mu_2)^T \Sigma_2^{-1} (\mu_1 - \mu_2) + \text{tr}(\Sigma_2^{-1} \Sigma_1) - d]. \quad (15)$$

The appropriate lower and upper bounds are found by plugging in Equations (14) and (15) into Equation (6).

These bounds have simple forms when all of the mixture components have equal covariance matrices; i.e., $\Sigma_i = \Sigma$ for all i . First, define a transformation in which each Gaussian component p_j is mapped to a different Gaussian $\tilde{p}_{j,\alpha}$, which has the same mean but where the covariance matrix is rescaled by $\frac{1}{\alpha(1-\alpha)}$,

$$p_j := \mathcal{N}(\mu_j, \Sigma) \quad \mapsto \quad \tilde{p}_{j,\alpha} := \mathcal{N}\left(\mu_j, \frac{1}{\alpha(1-\alpha)} \Sigma\right).$$

Then, the lower bound of Equation (10) can be written as

$$\hat{H}_{C_\alpha} = \frac{d}{2} + \frac{d}{2} \ln(\alpha(1-\alpha)) - \sum_i c_i \ln \sum_j c_j \tilde{p}_{j,\alpha}(\mu_i).$$

This is derived by combining the expressions for C_α , Equation (14), the entropy of a Gaussian, Equation (13), and the Gaussian density function. For a homoscedastic mixture, the tightest lower bound among the Chernoff α -divergences is given by $\alpha = 0.5$, corresponding to the Bhattacharyya distance,

$$\hat{H}_{BD} = \frac{d}{2} + \frac{d}{2} \ln \frac{1}{4} - \sum_i c_i \ln \sum_j c_j \tilde{p}_{j,0.5}(\mu_i).$$

(This is derived above in Section 3.2.)

For the upper bound, when all Gaussians have the same covariance matrix, we again combine the expressions for KL, Equation (15), the entropy of a Gaussian, Equation (13), and the Gaussian density function to give

$$\hat{H}_{KL}(X) = \frac{d}{2} - \sum_i c_i \ln \sum_j c_j p_j(\mu_i).$$

Note that this is exactly the expression for the kernel density estimator \hat{H}_{KDE} (Equation (3)), plus a dimensional correction. Thus, surprisingly \hat{H}_{KDE} is a reasonable entropy estimator for homoscedastic Gaussian mixtures, since it is only an additive constant away from KL-distance based estimator \hat{H}_{KL} (which has various beneficial properties, as described above). This may explain why \hat{H}_{KDE} has been used effectively in optimization contexts [4,8–10], where the additive constant is often irrelevant, despite lacking a principled justification in terms of being a bound on entropy.

5. Estimating Mutual Information

It is often of interest, for example in rate distortion and related problems [11], to calculate the mutual information across a communication channel,

$$MI(X;U) = H(X) - H(X|U),$$

where U is the distribution of signals sent across the channel, and X is the distribution of messages received on the other end of the channel. As with mixture distributions, it is often easy to compute $H(X|U)$, the entropy of the received signal given the sent signal (i.e., the distribution of noise on the channel). The marginal entropy of the received signals, $H(X)$, on the other hand, is often difficult to compute.

In some cases, the distribution of U may be well approximated by a mixture model. In this case, we can estimate the entropy of the received signals, $H(X)$, using our pairwise distance estimators, as discussed in Section 3. In particular, we have the lower bound

$$MI(X;U) \geq -\sum_i c_i \ln \sum_j c_j \exp(-C_\alpha(p_i||p_j)),$$

and the upper bound,

$$MI(X;U) \leq -\sum_i c_i \ln \sum_j c_j \exp(-\text{KL}(p_i||p_j)),$$

where p_i is the density of component i , and noting that the $H(X|U)$ terms cancel in the expression.

This also illuminates that the actual pairwise portion of the estimator, $-\sum_i c_i \ln \sum_j c_j \exp(-D(p_i||p_j))$ is a measure of the mutual information between the random variable specifying the component identity and the random variable distributed as the mixture of the component densities. If the components are identical, this mutual information is zero, since knowing the component identity tells one nothing about the outcome of X . On the other hand, when all of the components are very different from one another, knowing the component that generated X is very informative, giving the maximum amount of information, $H(C)$.

As a practical example, consider a scenario in which U is a random variable representing the outside temperature on any particular day. This temperature is measured with a thermometer with Gaussian measurement noise (the “Additive White Noise Gaussian channel”). This gives our measurement distribution

$$X = U + \mathcal{N}(0, \Sigma').$$

If the actual temperature U is (approximately or exactly) distributed as a mixture of M Gaussians, each one having mixture weight c_i , mean μ_i , and covariance matrix Σ_i , then X will also be distributed as a mixture of M Gaussians, each with weight c_i , mean μ_i , and covariance matrix $\Sigma'_j := \Sigma_i + \Sigma'$. We can then use our estimators to estimate the mutual information between the actual temperature, U , and thermometer measurements, X .

6. Numerical Results

In this section, we run numerical experiments and compare estimators of mixture entropy under a variety of conditions. We consider two different types of mixtures, mixtures of Gaussians and mixtures of uniform distributions, for a variety of parameter values. We evaluate the following estimators:

1. The true entropy, $H(X)$, as estimated by a Monte Carlo sampling of the mixture model. Two thousand samples were used for each MC estimate for the mixtures of Gaussians, and 5000 samples were used for the mixtures of uniform distributions.
2. Our proposed upper-bound, based on the KL divergence, \hat{H}_{KL} (Equation (12))
3. Our proposed lower-bound, based on the Bhattacharyya distance, \hat{H}_{BD} (Equation (11))
4. The kernel density estimate based on the component means, \hat{H}_{KDE} (Equation (3))
5. The lower bound based on the “Expected Likelihood Kernel”, \hat{H}_{ELK} (Equation (4))
6. The lower bound based on the conditional entropy, $H(X|C)$ (Equation (1))
7. The upper bound based on the joint entropy, $H(X, C)$ (Equation (2)).

We show the values of the estimators 1–5 as line plots, while the region between the conditional (6) and joint entropy (7) is shown in shaded green. The code for these figures can be found at [41], and uses the Gonum numeric library [42].

6.1. Mixture of Gaussians

In the first experiment, we evaluate the estimators on a mixture of randomly placed Gaussians, and look at their behavior as the distance between the means of the Gaussians increases. The mixture is composed of 100 10-dimensional Gaussians, each Gaussian distributed as $p_i = \mathcal{N}(\mu_i, \mathbf{I}_{(10)})$, where $\mathbf{I}_{(d)}$ indicates the $d \times d$ identity matrix. Means are sampled from $\mu_i \sim \mathcal{N}(0, \sigma \mathbf{I}_{(10)})$. Figure 1A depicts the change in estimated entropy as the means grow farther apart, in particular a function of $\ln(\sigma)$. We see that our proposed bounds are closer to the true entropy than the other estimators over the whole range of σ values, and in the extremes, our bounds approach the exact value of the true entropy. This is as expected, since as $\sigma \rightarrow 0$ all of the Gaussian mixture components become identical, and as $\sigma \rightarrow \infty$ all of the Gaussian components grow very far apart, approaching the case where each Gaussian is in its own “cluster”. The ELK lower bound is a strictly worse estimate than \hat{H}_{BD} , in this experiment. As expected, the KDE estimator differs by exactly $d/2$ from the KL estimator.

In the second experiment, we evaluate the entropy estimators as the covariance matrices change from less to more similar. We again generate 100 10-dimensional Gaussians. Each Gaussian is distributed as $p_i = \mathcal{N}(\mu_i, \Sigma_i)$, where now $\mu_i \sim \mathcal{N}(0, \mathbf{I}_{(10)})$ and $\Sigma_i \sim \mathcal{W}(\frac{1}{10+n} \mathbf{I}_{(10)}, n)$, where $\mathcal{W}(\mathbf{V}, n)$ is a Wishart distribution with scale-matrix \mathbf{V} and n degrees of freedom. Figure 1B compares the the estimators with the true entropy as a function of $\ln(n)$. When n is small, the Wishart distribution is broad and the covariance matrices differ significantly from one another, while as $n \rightarrow \infty$, all the covariance matrices become close to the identity $\mathbf{I}_{(10)}$. Thus, for small n , we essentially recover a “clustered” case, in which every component is in its own cluster and our lower and upper bounds give highly accurate estimates. For large n , we converge to the $\sigma = 1$ case of the first experiment.

In the third experiment, we again generate a mixture of 100 10-dimensional Gaussians. Now, however, the Gaussians are grouped into five “clusters”, with each Gaussian component randomly assigned to one of the clusters. We use $g(i) \in \{1 \dots 5\}$ to indicate the group of each Gaussian’s component $i \in \{1 \dots 100\}$, and each of the 100 Gaussians is distributed as $p_i = \mathcal{N}(\tilde{\mu}_{g(i)}, \mathbf{I}_{(10)})$. The cluster centers $\tilde{\mu}_k$ for $k \in \{1 \dots 5\}$ are drawn from $\mathcal{N}(0, \sigma \mathbf{I}_{(10)})$. The results are depicted in Figure 1C as a function of $\ln(\sigma)$. In the first experiment, we saw that the joint entropy $H(X, C)$ became an increasingly better estimator as the Gaussians grew increasingly far apart. Here, however, we see that there is a significant difference between $H(X, C)$ and the true entropy, even as the groups become increasingly separated. Our proposed bounds, on the other hand, provide accurate estimates of the entropy across the entire parameter sweep. As expected, they become exact in the limit when all clusters are at the same location, as well as when all clusters are very far apart from each other.

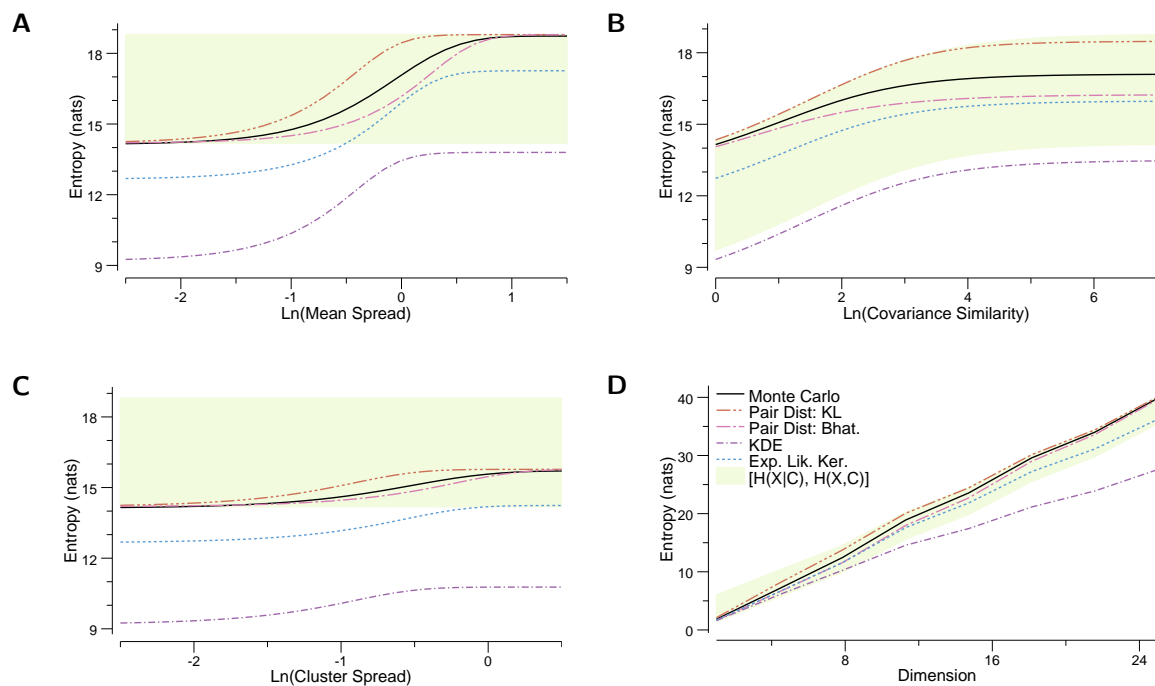


Figure 1. Entropy estimates for a mixture of 100 Gaussians. In each plot, the vertical axis shows the entropy of the distribution, and the horizontal axis changes a feature of the components: **(A)** the distance between means is increased; **(B)** the component covariances become more similar (at the right side of the plot, all Gaussians have covariance matrices approximately equal to the identity matrix); **(C)** the components are grouped into five “clusters”, and the distance between the locations of the clusters is increased; **(D)** the dimension is increased.

Finally, we evaluate the entropy estimators while changing the dimension of the Gaussian components. We again generate 100 Gaussian components, each distributed as $p_i = \mathcal{N}(\mu_i, \mathbf{I}_{(d)})$, with $\mu_i \sim \mathcal{N}(0, \sigma \mathbf{I}_{(d)})$. We vary the dimensionality d from 1 to 60. The results are shown in Figure 1D. First, we see that when $d = 1$, the KDE estimator and the KL-divergence based estimator give a very similar prediction (differing only by 0.5), but as the dimension increases, the two estimates diverge at a rate of $d/2$. Similarly, \hat{H}_{ELK} grows increasingly less accurate as the dimension increases. Our proposed lower and upper bounds provide good estimates of the mixture entropy across the whole sweep across dimensions.

As previously mentioned, our lower and upper bounds tend to perform best at the “extremes” and worse in the intermediate regimes. In particular, in Figure 1A,C,D, the distances between component means increase from left to right. On the left hand side of these figures, all of the component means are close and the component distributions overlap, as evidenced by the fact that the mixture entropy is $\approx H(X|C)$, i.e., $I(X;C) \approx 0$. In this regime, when there is essentially a single “cluster”, and our bounds become tight (see Section 3.4). On the right hand side of these figures, the components’ means are all far apart from each other, and the mixture entropy $\approx H(X,C)$, i.e., $I(X;C) \approx H(C)$ (in Figure 1C, it is the five clusters that become far apart, and the mixture entropy $\approx H(X|C) + \ln 5$). In this regime where there are many well-separated clusters, our bounds again become tight. In between these two extremes, however, there is no clear clustering of the mixture components, and the entropy bounds are not as tight.

As noted in the previous paragraph, the extremes in three out of the four subfigures approach the perfectly clustered case. In this situation, we show in Appendix B that the BD-based estimator is a better bound on the true entropy than the Expected Likelihood Kernel estimator. We see confirmation of this in the experimental results, where \hat{H}_{ELK} performs worse than the pairwise-distance based estimators.

6.2. Mixture of Uniforms

In the second set of experiments, we consider a mixture of uniform distributions. Unlike Gaussians, uniform distributions are bounded within a hyper-rectangle and do not have full support over the domain. In particular, a uniform distribution $p = \mathcal{U}(a, b)$ over d dimensions is defined as

$$p(x) \propto \begin{cases} 1, & \text{if } x_i \in [a_i, b_i] \forall i = 1 \dots d, \\ 0, & \text{otherwise,} \end{cases}$$

where x , a , and b are d -dimensional vectors, and the subscript x_i refers to value of x on dimension i . Note that when p_X is a mixture of uniforms, there can be significant regions where $p_X(x) > 0$, but $p_i(x) = 0$ for some i .

Here, we list the formulae for pairwise distance measure between uniform distributions. In the following, we use $V_i := \int_x 1\{p_i(x) > 0\}dx$ to indicate the “volume” of distribution p_i . Uniform components have a constant $p(x)$ over their support, and so $p_i(x) = 1/V_i$ for all x where $p_i(x) > 0$. Similarly, we use $V_{i \cap j}$ as the “volume of overlap” between p_i and p_j , i.e., the volume of the intersection of the support of p_i and p_j , $V_{i \cap j} := \int_x 1\{p_i(x) > 0\}1\{p_j(x) > 0\}dx$. The distance measures between uniforms are then

$$H(p_i) = \ln V_i,$$

$$\text{KL}(p_i \| p_j) = \begin{cases} \ln(V_j/V_i), & \text{supp } p_i \subseteq \text{supp } p_j, \\ \infty, & \text{otherwise,} \end{cases}$$

$$\text{BD}(p_i \| p_j) = 0.5 \ln V_i + 0.5 \ln V_j - \ln V_{i \cap j}, \quad (16)$$

$$\hat{H}_{\text{ELK}} = - \sum_i c_i \ln \sum_j c_j \frac{V_{i \cap j}}{V_i V_j}. \quad (17)$$

Like the Gaussian case, we run four different computational experiments and compare the mixture entropy estimates to the true entropy, as determined by Monte Carlo sampling.

In the first experiment, the mixture consists of 100 10-dimensional uniform components, with $p_i = \mathcal{U}(\mu_i - \mathbf{1}_{(10)}, \mu_i + \mathbf{1}_{(10)})$, and $\mu_i \sim \mathcal{N}(0, \sigma \mathbf{I}_{(10)})$, where $\mathbf{1}_{(d)}$ refers to a d -dimensional vector of 1s. Figure 2A depicts the change in entropy as a function of $\ln(\sigma)$. For very small σ , the distributions are almost entirely overlapping, while for large σ they tend very far apart. As expected, the entropy increases with σ . Here, we see that the prediction of \hat{H}_{KL} is identical to $H(X, C)$, which arises because $\text{KL}(p_i \| p_j)$ is infinite whenever the support of p_i is not entirely contained in the support of p_j . Uniform components with equal size and non-equal means must have some region of non-overlap, and so the KL is infinite between all pairs of components, thus KL is effectively D_{max} (Equation (8)). In contrast, we see that \hat{H}_{BD} estimates the true entropy quite well. This example demonstrates that getting an accurate estimate of mixture entropy may require selecting a distance function that works well with the component distributions. Finally, it turns out that, for uniform components of equal size, $\hat{H}_{\text{ELK}} = \hat{H}_{\text{BD}}$. This can be seen by combining Equations (6) and (16), and comparing to Equation (17) (note that $V_i = V_j$ when the components have equal size).

In the second experiment, we adjust the variance in the size of the uniform components. We again use 100 10-dimensional components, $p_i = \mathcal{U}(\mu_i - \gamma_i \mathbf{1}_{(10)}, \mu_i + \gamma_i \mathbf{1}_{(10)})$, where $\mu_i \sim \mathcal{N}(0, \mathbf{I}_{(10)})$, and $\gamma_i \sim \Gamma(1 + \sigma, 1 + \sigma)$, where $\Gamma(\alpha, \beta)$ is the Gamma distribution with shape parameter α and rate parameter β . Figure 2B shows the change in entropy estimates as a function of $\ln(\sigma)$. When σ is small, the sizes have significant spread, while as σ grows the distributions become close to equally sized. We again see that \hat{H}_{BD} is a good estimator of entropy, outperforming all of the other estimators. Generally, not all supports will be non-overlapping, so \hat{H}_{KL} will not necessarily be equal to $H(X, C)$, though we find the two to be numerically quite close. In this experiment, we find that the lower and upper bounds specified by \hat{H}_{BD} and \hat{H}_{KL} provide a tight estimate of the true entropy.

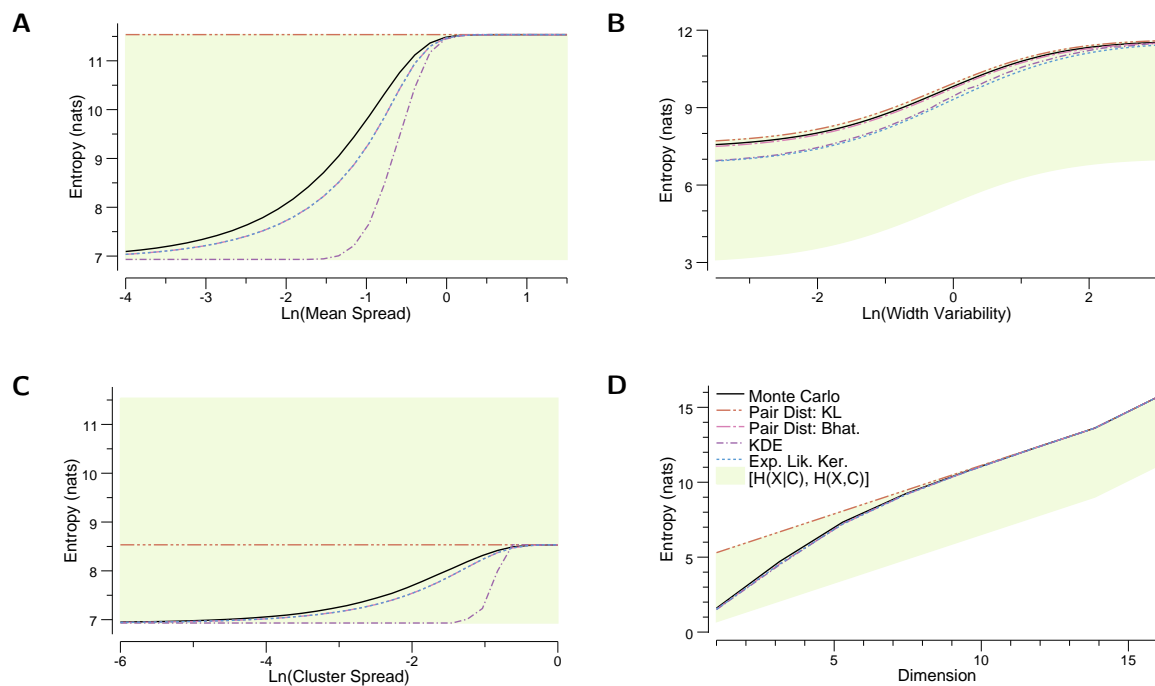


Figure 2. Entropy estimates for a mixture of 100 uniform components. In each plot, the vertical axis shows the entropy of the distribution, and the horizontal axis changes a feature of the components: (A) the distance between means is increased; (B) the component sizes become more similar (at the right side of the plot, all components have approximately the same size); (C) the components are grouped into five “clusters”, and the distance between these clusters is increased; (D) the dimension is increased.

In the third experiment, we again consider a clustered mixture, and evaluate the entropy estimators as these clusters grow apart. Here, there are 100 components with $p_i = \mathcal{U}(\tilde{\mu}_{g(i)} - \mathbf{1}_{(10)}, \tilde{\mu}_{g(i)} + \mathbf{1}_{(10)})$, where $g(i) \in \{1 \dots 5\}$ is the randomly assigned cluster identity of component i . The cluster centers $\tilde{\mu}_k$ for $k \in \{1 \dots 5\}$ are generated according to $\mathcal{N}(0, \sigma \mathbf{I}_{(10)})$. Figure 2C shows the change in entropy as the clusters locations move apart. Note that, in this case, the upper bound \hat{H}_{KL} significantly outperforms $H(X, C)$, unlike in the first and second experiment, because in this experiment, components in the same cluster have perfect overlap. We again see that \hat{H}_{BD} provides a relatively accurate lower bound for the true entropy.

In the final experiment, the dimension of the components is varied. There are again 100 components, with $p_i = \mathcal{U}(\mu_i - \mathbf{1}_{(d)}, \mu_i + \mathbf{1}_{(d)})$, and $\mu_i \sim \mathcal{N}(0, \sigma \mathbf{I}_{(d)})$. Figure 2D shows the change in entropy as the dimension increases from $d = 1$ to $d = 16$. Interestingly, in the low-dimensional case, $H(X|C)$ is a very close estimate for the true entropy, while in the high-dimensional case, the entropy becomes very close to $H(X, C)$. This is because in higher dimensions, there is more ‘space’ for the components to be far from each other. As in the first experiment, \hat{H}_{KL} is equal to $H(X, C)$. We again observe that \hat{H}_{BD} provides a tight lower bound on the mixture entropy, regardless of dimension.

7. Discussion

We have presented a new class of estimators for the entropy of a mixture distribution. We have shown that any estimator in this class has a bounded estimation bias, and that this class includes useful lower and upper bounds on the entropy of a mixture. Finally, we show that these bounds become exact when mixture components are grouped into well-separated clusters.

Our derivation of the bounds make use of some existing results [5,31]. However, to our knowledge, these results have not been previously used to estimate mixture entropies. Furthermore, they have not been compared numerically or analytically to better-known bounds.

We evaluated these estimators using numerical simulations of mixtures of Gaussians as well as mixtures of bounded (hypercube) uniform distributions. Our results demonstrate that our estimators perform much better than existing well-known estimators.

This estimator class can be especially useful for optimization problems that involve minimization of entropy or mutual information. If the distance function used in the pairwise estimator class is continuous and smooth in the parameters of the mixture components, then the entropy estimate is also continuous and smooth. This permits our estimators to be used within gradient-based optimization techniques, for example gradient descent, as often done in machine learning problems.

In fact, we have used our upper bound to implement a non-parametric, nonlinear version of the “Information Bottleneck” [43]. Specifically, we minimized an upper bound on the mutual information between input and hidden layer in a neural networks [11]. We found that the optimal distributions were often clustered (Section 3.4). That work demonstrated practically the value of having an accurate, differentiable upper bound on mixture entropy that performs well in the clustered regime.

Note that we have not proved that the bounds derived here are the best possible. Identifying better bounds, or proving that our results are optimal within some class of bounds, remains for future work.

Acknowledgments: We thank David H. Wolpert for useful discussions. We would also like to thank the Santa Fe Institute for helping to support this research. This work was made possible through the support of AFOSR MURI on multi-information sources of multi-physics systems under Award Number FA9550-15-1-0038.

Author Contributions: Artemy Kolchinsky and Brendan D. Tracey designed the method and experiments; Brendan D. Tracey performed the experiments; Artemy Kolchinsky and Brendan D. Tracey wrote the paper.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A. Chernoff α -Divergence Is Not a Distance Function for $\alpha \notin [0, 1]$

For any pair of densities p and q , consider the Chernoff α -divergence

$$C_\alpha(p||q) := -\ln \int p^\alpha(x)q^{1-\alpha}(x) dx = -\ln c_\alpha(p||q),$$

where the quantity c_α is called the Chernoff α -coefficient [29]. Taking the second derivative of c_α with respect to α gives

$$\frac{d^2}{d\alpha^2}c_\alpha(p||q) = \int p^\alpha(x)q^{1-\alpha}(x) \left(\ln \frac{p(x)}{q(x)} \right)^2 dx.$$

Observe that this quantity is everywhere positive, meaning that $c_\alpha(p||q)$ is convex everywhere. For simplicity, consider the case $p \neq q$, in which case this function is *strictly* convex. In addition, observe that for any p and q , $c_\alpha(p||q) = 1$ when $\alpha = 0$ and $\alpha = 1$. If $c_\alpha(p||q)$ is strictly convex in α , this must mean that $c_\alpha(p||q) > 1$ for $\alpha \notin [0, 1]$. This in turn implies that the Chernoff α -divergence C_α is strictly negative for $\alpha \notin [0, 1]$. Thus, C_α is not a valid distance function for $\alpha \notin [0, 1]$, as defined in Section 3.

Appendix B. For Clustered Mixtures, $\hat{H}_{BD} \geq \hat{H}_{ELK}$

Assume a mixture with perfect clustering (Section 3.4). Specifically, we assume that if $g(i) = g(j)$, then $p_i(x) = p_j(x)$, and if $g(i) \neq g(j)$ then both $\exp(-C_\alpha(p_i||p_j)) \approx 0$ and $\int p_i(x)p_j(x)dx \approx 0$.

In this case, our lower bound \hat{H}_{BD} is at least as good as \hat{H}_{ELK} . Specifically, \hat{H}_{ELK} becomes

$$\begin{aligned} \hat{H}_{ELK} &= -\sum_i c_i \ln \sum_j c_j \int p_i(x)p_j(x)dx \\ &= -\sum_i c_i \ln \sum_j c_j \delta_{g(i),g(j)} \int p_i(x)^2 dx \\ &= -\sum_k p_G(k) \ln \sum_k p_G(k) \int p_k(x)^2 dx, \end{aligned}$$

where $p_k(x)$ is shorthand for the density of any component in cluster k (remember that all components in the same cluster have equal density). \hat{H}_{C_α} becomes

$$\begin{aligned}\hat{H}_{C_\alpha} &= H(X|C) - \sum_i c_i \ln \sum_j c_j \exp(-C_\alpha(p_i||p_j)) \\ &= - \sum_i c_i \int p_i(x) \ln p_i(x) dx - \sum_i c_i \ln \sum_j c_j \delta_{g(i),g(j)} \\ &= - \sum_k p_G(k) \int p_k(x) \ln p_k(x) dx - \sum_k p_G(k) \ln p_G(k) \\ &\stackrel{(a)}{\geq} - \sum_k p_G(k) \ln \int p_k(x)^2 dx - \sum_k p_G(k) \ln p_G(k) \\ &= - \sum_k p_G(k) \ln p_G(k) \int p_k(x)^2 dx = \hat{H}_{ELK},\end{aligned}$$

where (a) uses Jensen's inequality.

References

1. McLachlan, G.; Peel, D. *Finite Mixture Models*; John Wiley & Sons: Hoboken, NJ, USA, 2004.
2. Cover, T.M.; Thomas, J.A. *Elements of Information Theory*; John Wiley & Sons: Hoboken, NJ, USA, 2012.
3. Goldberger, J.; Gordon, S.; Greenspan, H. An Efficient Image Similarity Measure Based on Approximations of KL-Divergence Between Two Gaussian Mixtures. In Proceedings of the 9th International Conference on Computer Vision, Nice, France, 13–16 October 2003; Volume 3, pp. 487–493.
4. Viola, P.; Schraudolph, N.N.; Sejnowski, T.J. Empirical Entropy Manipulation for Real-World Problems. In *Advances in Neural Information Processing Systems*; The MIT Press: Cambridge, MA, USA, 1996; pp. 851–857.
5. Hershey, J.R.; Olsen, P.A. Approximating the Kullback Leibler divergence between Gaussian mixture models. In Proceedings of the 2007 IEEE International Conference on Acoustics, Speech and Signal Processing, Honolulu, HI, USA, 15–20 April 2007; Volume 4, pp. 317–320.
6. Chen, J.Y.; Hershey, J.R.; Olsen, P.A.; Yashchin, E. Accelerated monte carlo for kullback-leibler divergence between gaussian mixture models. In Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing, Las Vegas, NV, USA, 31 March–4 April 2008; pp. 4553–4556.
7. Capristan, F.M.; Alonso, J.J. Range Safety Assessment Tool (RSAT): An analysis environment for safety assessment of launch and reentry vehicles. In Proceedings of the 52nd Aerospace Sciences Meeting, National Harbor, MD, USA, 13–17 January 2014; p. 304.
8. Schraudolph, N.N. Optimization of Entropy with Neural Networks. Ph.D. Thesis, University of California, San Diego, CA, USA, 1995.
9. Schraudolph, N.N. Gradient-based manipulation of nonparametric entropy estimates. *IEEE Trans. Neural Netw.* **2004**, *15*, 828–837.
10. Shwartz, S.; Zibulevsky, M.; Schechner, Y.Y. Fast kernel entropy estimation and optimization. *Signal Process.* **2005**, *85*, 1045–1058.
11. Kolchinsky, A.; Tracey, B.D.; Wolpert, D.H. Nonlinear Information Bottleneck. *arXiv* **2017**, arXiv:1705.02436.
12. Contreras-Reyes, J.E.; Cortés, D.D. Bounds on Rényi and Shannon Entropies for Finite Mixtures of Multivariate Skew-Normal Distributions: Application to Swordfish (*Xiphias gladius* Linnaeus). *Entropy* **2016**, *18*, 382.
13. Carreira-Perpinan, M.A. Mode-finding for mixtures of Gaussian distributions. *IEEE Trans. Pattern Anal. Mach. Intell.* **2000**, *22*, 1318–1323.
14. Zobay, O. Variational Bayesian inference with Gaussian-mixture approximations. *Electron. J. Stat.* **2014**, *8*, 355–389.
15. Beirlant, J.; Dudewicz, E.J.; Györfi, L.; van der Meulen, E.C. Nonparametric entropy estimation: An overview. *Int. J. Math. Stat. Sci.* **1997**, *6*, 17–39.
16. Joe, H. Estimation of entropy and other functionals of a multivariate density. *Ann. Inst. Stat. Math.* **1989**, *41*, 683–697.

17. Nair, C.; Prabhakar, B.; Shah, D. On entropy for mixtures of discrete and continuous variables. *arXiv* **2006**, arXiv:cs/0607075.
18. Huber, M.F.; Bailey, T.; Durrant-Whyte, H.; Hanebeck, U.D. On entropy approximation for Gaussian mixture random vectors. In Proceedings of the IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems, Seoul, Korea, 20–22 August 2008; pp. 181–188.
19. Hall, P.; Morton, S.C. On the estimation of entropy. *Ann. Inst. Stat. Math.* **1993**, *45*, 69–88.
20. Principe, J.C.; Xu, D.; Fisher, J. Information theoretic learning. *Unsuperv. Adapt. Filter* **2000**, *1*, 265–319.
21. Xu, J.W.; Paiva, A.R.; Park, I.; Principe, J.C. A reproducing kernel Hilbert space framework for information-theoretic learning. *IEEE Trans. Signal Process.* **2008**, *56*, 5891–5902.
22. Jebara, T.; Kondor, R. Bhattacharyya and expected likelihood kernels. In *Learning Theory and Kernel Machines*; Springer: Berlin, Germany, 2003; pp. 57–71.
23. Jebara, T.; Kondor, R.; Howard, A. Probability product kernels. *J. Mach. Learn. Res.* **2004**, *5*, 819–844.
24. Banerjee, A.; Merugu, S.; Dhillon, I.S.; Ghosh, J. Clustering with Bregman divergences. *J. Mach. Learn. Res.* **2005**, *6*, 1705–1749.
25. Cichocki, A.; Zdunek, R.; Phan, A.H.; Amari, S.I. Similarity Measures and Generalized Divergences. In *Nonnegative Matrix and Tensor Factorizations: Applications to Exploratory Multi-Way Data Analysis and Blind Source Separation*; John Wiley & Sons: Hoboken, NJ, USA, 2009; pp. 81–129.
26. Cichocki, A.; Amari, S.I. Families of alpha-beta-and gamma-divergences: Flexible and robust measures of similarities. *Entropy* **2010**, *12*, 1532–1568.
27. Gil, M.; Alajaji, F.; Linder, T. Rényi divergence measures for commonly used univariate continuous distributions. *Inf. Sci.* **2013**, *249*, 124–131.
28. Crooks, G.E. On Measures of Entropy and Information. Available online: http://threeplusone.com/on_information.pdf (accessed on 12 July 2017).
29. Nielsen, F. Chernoff information of exponential families. *arXiv* **2011**, arXiv:1102.2684.
30. Van Erven, T.; Harremoës, P. Rényi divergence and Kullback-Leibler divergence. *IEEE Trans. Inf. Theory* **2014**, *60*, 3797–3820.
31. Haussler, D.; Opper, M. Mutual information, metric entropy and cumulative relative entropy risk. *Ann. Stat.* **1997**, *25*, 2451–2492.
32. Fukunaga, K. *Introduction to Statistical Pattern Recognition*, 2nd ed.; Academic Press: Boston, MA, USA, 1990.
33. Paisley, J. *Two Useful Bounds for Variational Inference*; Princeton University: Princeton, NJ, USA, 2010.
34. Sason, I.; Verdú, S. *f-Divergence Inequalities*. *IEEE Trans. Inf. Theory* **2016**, *62*, 5973–6006.
35. Hero, A.O.; Ma, B.; Michel, O.; Gorman, J. Alpha-Divergence for Classification, Indexing and Retrieval. Available online: <https://pdfs.semanticscholar.org/6d51/fbf90c59c2bb8cbf0cb609a224f53d1b68fb.pdf> (accessed on 14 July 2017).
36. Dowson, D.; Landau, B. The Fréchet distance between multivariate normal distributions. *J. Multivar. Anal.* **1982**, *12*, 450–455.
37. Olkin, I.; Pukelsheim, F. The distance between two random vectors with given dispersion matrices. *Linear Algebra Appl.* **1982**, *48*, 257–263.
38. Pardo, L. *Statistical Inference Based on Divergence Measures*; CRC Press: Boca Raton, FL, USA, 2005.
39. Hobza, T.; Morales, D.; Pardo, L. Rényi statistics for testing equality of autocorrelation coefficients. *Stat. Methodol.* **2009**, *6*, 424–436.
40. Nielsen, F. Generalized Bhattacharyya and Chernoff upper bounds on Bayes error using quasi-arithmetic means. *Pattern Recognit. Lett.* **2014**, *42*, 25–34.
41. GitHub. Available online: <https://www.github.com/btracey/mixent> (accessed on 14 July 2017).
42. Gonum Numeric Library. Available online: <https://www.gonum.org> (accessed on 14 July 2017).
43. Tishby, N.; Pereira, F.; Bialek, W. The information bottleneck method. In Proceedings of the 37th Annual Allerton Conference on Communication, Control, and Computing, Monticello, IL, USA, 22–24 September 1999.

