



# On the Estimation of Information Measures of Continuous Distributions

Pichler Georg, Günther Koliander, Pablo Piantanida

## ► To cite this version:

Pichler Georg, Günther Koliander, Pablo Piantanida. On the Estimation of Information Measures of Continuous Distributions. 2023. hal-04137020

**HAL Id: hal-04137020**

**<https://hal.science/hal-04137020>**

Preprint submitted on 22 Jun 2023

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# On the Estimation of Information Measures of Continuous Distributions

Georg Pichler<sup>1</sup>, Pablo Piantanida<sup>2</sup>, and Günther Koliander<sup>3</sup>

<sup>1</sup>*Institute of Telecommunications, TU Wien, Vienna, Austria*

<sup>2</sup>*Université Paris-Saclay, CNRS, CentraleSupélec, Laboratoire des Signaux et Systèmes, Gif-sur-Yvette, France and Montreal Institute for Learning Algorithms (Mila), Université de Montréal, QC, Canada*

<sup>3</sup>*Acoustics Research Institute, Austrian Academy of Sciences, Vienna, Austria*

## Abstract

The estimation of information measures of continuous distributions based on samples is a fundamental problem in statistics and machine learning. In this paper, we analyze estimates of differential entropy in  $K$ -dimensional Euclidean space, computed from a finite number of samples, when the probability density function belongs to a predetermined convex family  $\mathcal{P}$ . First, estimating differential entropy to any accuracy is shown to be infeasible if the differential entropy of densities in  $\mathcal{P}$  is unbounded, clearly showing the necessity of additional assumptions. Subsequently, we investigate sufficient conditions that enable confidence bounds for the estimation of differential entropy. In particular, we provide confidence bounds for simple histogram based estimation of differential entropy from a fixed number of samples, assuming that the probability density function is Lipschitz continuous with known Lipschitz constant and known, bounded support. Our focus is on differential entropy, but we provide examples that show that similar results hold for mutual information and relative entropy as well.

## 1 Introduction

Many learning tasks, especially in unsupervised/semi-supervised settings, use information theoretic quantities, such as relative entropy, mutual information, differential entropy, or other divergence functionals as target functions in numerical optimization problems [19, 20, 25, 38, 16, 8, 5]. Furthermore, estimators for information theoretical quantities are useful in other fields, such as neuroscience [26]. As these quantities typically cannot be computed directly, surrogate functions, either upper/lower bounds, or estimates are used in place. Here, we will investigate the problem of estimating differential entropy using a finite number of samples. Throughout, we will restrict our attention to differential entropy, but

similar results also hold for conditional differential entropy, mutual information and relative entropy (cf. Section 4).

## 1.1 Our Contribution

The contributions of this work can be summarized as follows:

- First, we explore the following basic but fundamental question: Fixing  $C, \delta > 0$  and given  $N \in \mathbb{N}$  samples from a probability density function (pdf)  $p \in \mathcal{P}$ , where  $\mathcal{P}$  is a family of pdfs on  $\mathbb{R}^K$ , is it possible to obtain an estimate  $\hat{h}$  of the differential entropy  $h(p) < \infty$  satisfying

$$\mathbb{P}\left\{|\hat{h} - h(p)| > C\right\} \leq \delta?$$

In Section 2 we show that the answer to this question is negative (Proposition 2) if  $\mathcal{P}$  is convex and the differential entropy of the pdfs in  $\mathcal{P}$  is unbounded.

- Subsequently, we investigate sufficient conditions for the class  $\mathcal{P}$  that enable estimation of differential entropy with such a confidence bound and in Section 3 (Theorem 3) we show that a known, bounded support together with an  $L$ -Lipschitz continuous pdf for fixed  $L > 0$ , suffices.<sup>1</sup> For a simple histogram based estimator we explicitly compute a relation between probability of correct estimation, accuracy, dimension  $K$ , sample size  $N$ , and Lipschitz constant  $L$ . It is shown that estimation becomes impossible if either assumption is removed.
- Finally, in Section 4 we obtain impossibility results, similar to Proposition 2 for the estimation of other information measures.

## 1.2 Previous Work

The problem of estimating information measures from a finite number of samples is as old as information theory itself. Shortly after his seminal paper [31], Shannon worked on estimating the entropy rate of English text [32]. There have been numerous works on the estimation of information measures, such as entropy, mutual information, and differential entropy, since. There are many different approaches for estimating information measures, including kernel based methods, nearest neighbor methods, methods based on sample distances as well as multiple variants of plug-in estimates. Many estimators have been shown to be consistent and/or asymptotically unbiased under various constraints, e.g., in [17, 11, 12, 21, 36]. An excellent overview can be found in [6].

In [36], rate-of-convergence results as well as a central limit theorem are provided for differential entropy and Rényi entropy. However, the confidence

---

<sup>1</sup>These assumptions assure that the differential entropy of the pdfs in  $\mathcal{P}$  is bounded. A known, bounded support bounds the differential entropy from above, and  $L$ -Lipschitz continuity bounds it from below.

bounds and the constants involved in the rate-of-convergence results depend on the underlying distribution which is typically unknown. Similarly, [18] obtains a rate-of-convergence result, assuming a Lipschitz ball smoothness assumption combined with known compact support, but the involved constants remain unspecified. In a similar spirit, [23] provides asymptotic results for the estimation of differential entropy in two dimensions, when certain smoothness conditions are satisfied and the pdf is bounded away from zero. The related task of estimating relative entropy is studied, e. g., in [39, 27] and partition-based estimation of mutual information is analyzed in [9]. While [39] only shows consistency, convergence rates are obtained in [27 Th. 2], but again the constants involved remain unspecified.

In contrast to our present work, the existing results for the estimation of differential entropy mentioned above fall short when addressing the practical problem of a finite sample size. However, some results are available in a more general context. In [35], a finite-sample analysis is conducted. Similar to our approach (cf. Section 3), the authors of [35] assume a fixed support  $[0, 1]^K$ , but instead of Lipschitz continuity,  $\beta$ -Hölder continuity,  $\beta \in (0, 2]$  is assumed. Additionally, strict positivity on the interior of the support is required and the constants bounding the approximation error depend on the underlying, unknown distribution. These additional complications are likely due to the extended scope, as [35] is not focused on differential entropy, but the expectation of arbitrary functionals of the probability density. The same authors also provide finite sample analysis for the estimation of Rényi divergence under similarly strong conditions in [34].

There are several negative results, which clearly show that information measures are hard to estimate from a finite number of samples. It was shown in [3 Th. 4] that rate-of-convergence results cannot be obtained for any consistent estimator of entropy on a countable alphabet and only when imposing various assumptions on the true distribution, rate-of-convergence results were obtained. More negative results on the estimation of entropy and mutual information can be found in [28]. In fact, obtaining confidence bounds for information measures from samples is inherently difficult and requires regularity assumptions about the involved distributions, which are not subject to empirical test. In the seminal work of [4] as well as subsequent works [13, 10, 29, 11] (and references therein) such necessary conditions for the estimation of statistical parameters with confidence bounds are discussed in great detail and generality. The results of [10, 29] can be applied to differential entropy estimation and yield a result very similar to Proposition 2 essentially showing that differential entropy cannot be bounded using a finite number of samples, unless additional assumptions on the distribution are made.

Especially in the context of unsupervised and semi-supervised machine learning it recently became popular to use variational bounds or estimates of information measures as part of the loss function for training neural networks [7, 30, 19]. Criticism to this approach, in particular the use of variational bounds, has been voiced [24]. The current paper has a more general scope, dealing with the estimation problem of information measures in general, not limited to specific

variational bounds or techniques.

The information flow in neural networks is also a recent topic of investigation. In [33], an argument for successive compression in the layers of a deep neural network is given, along the lines of the information bottleneck method [37]. While flaws in this argument were pointed out [2, 22], the authors of [15] found that a clustering phenomenon might elucidate the behavior of deep neural networks. These insights were obtained by estimating the differential entropy  $h(\mathbf{X} + \mathbf{G})$  of a sum of two random vectors, where  $\mathbf{X}$  is sub-Gaussian and  $\mathbf{G} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$  is an independent Gaussian vector. This is similar in spirit to the work conducted here, however, our assumption of compact support is replaced by assuming  $\mathbf{X}$  to be sub-Gaussian<sup>2</sup>. Note that the pdf of  $\mathbf{X} + \mathbf{G}$  is Lipschitz continuous with fixed Lipschitz constant  $L(\sigma^2)$ , so [15] is implicitly also using a Lipschitz assumption.

## 2 The Nonexistence of Confidence Sets

Let  $\mathcal{P}$  be a family of pdfs on  $\mathcal{X} := \mathbb{R}^K$  with finite differential entropy, i. e.,  $h(p) := -\int p(\mathbf{x}) \log p(\mathbf{x}) d\mathbf{x} \in \mathbb{R}$  for every  $p \in \mathcal{P}$ .

Suppose we observe  $N$  i.i.d. copies  $\mathcal{D} := (\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_N)$  of some random vector  $\mathbf{X} \sim p \in \mathcal{P}$  and want to obtain an estimate of differential entropy from these samples  $\mathcal{D}$ . Such an estimator is a function  $\hat{h}: \mathcal{X}^N \rightarrow \mathbb{R}$  that maps  $\mathcal{D}$  into  $\hat{h}(\mathcal{D})$ , approximating the differential entropy  $h := h(\mathbf{X}) := h(p) < \infty$ . Its accuracy can be measured by a confidence interval, a widely used tool in statistical practice for indicating the precision of point estimators. For a given error probability  $\delta > 0$ , we would like to have  $C > 0$  such that  $|h - \hat{h}(\mathcal{D})| \geq C$  with probability less than  $\delta$ , i. e., a confidence interval of size  $C$  with confidence  $1 - \delta$ . However, there is no free lunch when estimating differential entropy, as evidenced by the following result, a corollary of a more general result in [10], here specialized to a bound of differential entropy. It is based on the abstract notion of a dense graph condition (DGC)<sup>3</sup>.

**Theorem 1** ([10] Th. 2.1). *Assume that  $h$  satisfies the DGC over  $\mathcal{P}$  and define  $B := \sup\{h(p) : p \in \mathcal{P}\}$ , where, e. g.,  $B = +\infty$ . If for any  $C > 0$ ,  $\sup_{p \in \mathcal{P}} P_p\{\hat{h}(\mathcal{D}) + C \leq B\} = 1$ , then*

$$\inf_{p \in \mathcal{P}} P_p\{h(p) - \hat{h}(\mathcal{D}) \leq C\} = 0.$$

A similar result follows from [29, Prop. 3.1].

We will not work with the DGC, but make two practical assumptions:  $\mathcal{P}$  is a convex family and the differential entropy of the pdfs in  $\mathcal{P}$  is unbounded (either from above or from below). Under these assumptions we show that for any  $\delta, C > 0$  there is a pdf  $p \in \mathcal{P}$ , such that  $|h - \hat{h}(\mathcal{D})| \leq C$  with probability

<sup>2</sup>Similar to our assumption of an arbitrary but fixed compact support in Section 3 the constant  $K$  in the definition of the sub-Gaussian  $\mathbf{X}$  is assumed to be fixed in [14] eq. (1).

<sup>3</sup> $h$  satisfies the DGC over  $\mathcal{P}$  if the graph of  $h$  over  $\mathcal{P}$  is dense in its own epigraph [10] eq. (2.4)].

less than  $\delta$ , i. e.,  $\hat{h}(\mathcal{D})$  is far from  $h$  with high probability. Fundamentally, this follows from the fact that  $\mathcal{P}$  contains pdfs with a large difference in differential entropy, which cannot be accurately distinguished based on samples. Similar results hold true for mutual information and relative entropy and are given in Section 4

**Proposition 2.** *Let  $\mathcal{P}$  be a convex family of pdfs with unbounded differential entropy, i. e., for any  $\alpha \in [0, 1]$  and  $p, q \in \mathcal{P}$ , we have  $\alpha p + (1 - \alpha)q \in \mathcal{P}$  as well as  $\sup_{q \in \mathcal{P}} |h(q)| = \infty$ . Then, for any pair of constants  $C, \delta > 0$ , there exists a continuous random vector  $\mathbf{X} \sim p \in \mathcal{P}$ , satisfying*

$$P_p \left\{ |h(p) - \hat{h}(\mathcal{D})| \leq C \right\} \leq \delta.$$

*Remark 1.* Before proceeding with the proof of Proposition 2 we note that this result could be proved as a consequence of Theorem 1. However, this would necessitate to show that our conditions imply the DGC. Furthermore, the proof of [10] Th. 2.1] itself hinges on deep statistical results and thus we opted for providing a short, self-contained proof.

*Proof of Proposition 2.* The function  $\hat{h}$ , constants  $C, \delta > 0$ , and the sample size  $N \in \mathbb{N}$  are arbitrary, but fixed. Choose an arbitrary  $\bar{\mathbf{X}} \sim \bar{p} \in \mathcal{P}$  and let  $\bar{h} := h(\bar{p}) < \infty$ . Then fix  $b > 0$ , such that  $P \left\{ |\hat{h}(\bar{\mathbf{X}}_1, \bar{\mathbf{X}}_2, \dots, \bar{\mathbf{X}}_N) - \bar{h}| \leq b \right\} \geq 1 - \frac{\delta}{2}$ , where  $(\bar{\mathbf{X}}_n)_{n=1, \dots, N}$  are i.i.d. copies of  $\bar{\mathbf{X}}$ . Furthermore, let  $Q \sim \mathcal{B}(1 - \varepsilon)$  be a Bernoulli random variable with parameter  $1 - \varepsilon = P\{Q = 1\}$ , independent of  $\bar{\mathbf{X}}$ , where  $0 < \varepsilon \leq \frac{\delta}{2N}$ . Choose  $a > 0$  such that  $a\varepsilon > b + C + \log 2$ . By our assumption  $\sup_{q \in \mathcal{P}} |h(q)| = \infty$ , we can find  $\tilde{\mathbf{X}} \sim \tilde{p} \in \mathcal{P}$  with  $\tilde{h} := h(\tilde{p})$  such that  $|\tilde{h} - \bar{h}| \geq a$ .

Define  $\mathbf{X} := Q\bar{\mathbf{X}} - (1 - Q)\tilde{\mathbf{X}}$ , which yields  $h = h(\mathbf{X}) = I(\mathbf{X}; Q) + h(\mathbf{X}|Q) = I(\mathbf{X}; Q) + (1 - \varepsilon)\bar{h} + \varepsilon\tilde{h} \in (1 - \varepsilon)\bar{h} + \varepsilon\tilde{h} + [0, \log 2]$  where  $I(\cdot; \cdot)$  denotes mutual information. For convenience we use  $\hat{H} := \hat{h}(\mathcal{D})$ , where  $\mathcal{D} = (\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_N)$  and define the event  $\mathcal{E} = \{Q_1 = Q_2 = \dots = Q_N = 1\} = \{\mathcal{D} = (\bar{\mathbf{X}}_1, \bar{\mathbf{X}}_2, \dots, \bar{\mathbf{X}}_N)\}$ . By the union bound, we have  $P\{\mathcal{E}\} \geq 1 - N\varepsilon$  and obtain

$$\begin{aligned} & P\left\{ |\hat{H} - \bar{h}| \leq b \right\} \\ &= P\{\mathcal{E}\} P\left\{ |\hat{H} - \bar{h}| \leq b \middle| \mathcal{E} \right\} \\ &\quad + P\{\mathcal{E}^c\} P\left\{ |\hat{H} - \bar{h}| \leq b \middle| \mathcal{E}^c \right\} \\ &\geq (1 - N\varepsilon) P\left\{ |\hat{H} - \bar{h}| \leq b \middle| \mathcal{E} \right\} \\ &= (1 - N\varepsilon) P\left\{ |\hat{h}(\bar{\mathbf{X}}_1, \bar{\mathbf{X}}_2, \dots, \bar{\mathbf{X}}_N) - \bar{h}| \leq b \right\} \\ &\geq \left(1 - \frac{\delta}{2}\right)^2 \geq 1 - \delta. \end{aligned}$$

We thus found  $\mathbf{X}$  such that

$$\begin{aligned}
& \mathbb{P}\{|h - \hat{H}| \leq C\} \\
& \leq \mathbb{P}\{|\bar{h} - \hat{H}| \geq |\bar{h} - h| - C\} \\
& \leq \mathbb{P}\{|\bar{h} - \hat{H}| \geq \varepsilon|\bar{h} - \tilde{h}| - \log 2 - C\} \\
& \leq \mathbb{P}\{|\bar{h} - \hat{H}| \geq \varepsilon a - \log 2 - C\} \\
& \leq \mathbb{P}\{|\bar{h} - \hat{H}| > b\} \\
& = 1 - \mathbb{P}\{|\bar{h} - \hat{H}| \leq b\} \leq \delta. \quad \square
\end{aligned}$$

*Remark 2.* Proposition 2 shows that in order to obtain confidence bounds, one needs to make assumptions about the underlying distribution. However, as pointed out in [10] p. 1395], when making these assumptions, one uses information external to the samples.

*Remark 3.* Note that the family of all pdfs with support  $[0, 1]^K$  satisfies the requirements of Proposition 2. It also satisfies the DGC, but it is not strongly nonparametric, as defined in [10] p. 1395].

### 3 Lipschitz Density Assumption

One way to avoid the problems outlined in Section 2 is to impose additional assumptions on the underlying probability distribution, that bound the differential entropy from above and from below. We will showcase that the differential entropy of an  $L$ -Lipschitz continuous pdf with fixed, known  $L > 0$  on  $\mathbb{R}^K$  and known, compact support  $\mathcal{X}$  can be well approximated from samples. In the following, let  $\mathbf{X} \sim p$  be supported<sup>4</sup> on  $\mathcal{X} := [0, 1]^K$ , i. e.,  $\int_{\mathcal{X}} p \, d\lambda^K = \mathbb{P}\{\mathbf{X} \in \mathcal{X}\} = 1$ , where  $\lambda^K$  denotes the Lebesgue measure on  $\mathbb{R}^K$ . The pdf  $p: \mathbb{R}^K \rightarrow \mathbb{R}_+$  of  $\mathbf{X}$  is assumed to be  $L$ -Lipschitz continuous on  $\mathbb{R}^K$  with some fixed  $L > 0$ , where  $\mathbb{R}^K$  is equipped with the  $\ell_1$ -norm<sup>5</sup>  $\|\mathbf{x}\| := \|\mathbf{x}\|_1 = \sum_k |x_k|$ , hence,

$$\forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^K : |p(\mathbf{x}) - p(\mathbf{y})| \leq L\|\mathbf{x} - \mathbf{y}\|.$$

Given  $N$  i.i.d. copies  $\mathcal{D} = (\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_N)$  of  $\mathbf{X}$ , let  $\mathbf{Y}$  be distributed according to the empirical distribution of  $\mathcal{D}$ , i. e.,  $\mathbf{Y} = \mathbf{X}_U$ , where  $U \sim \mathcal{U}(\{1, 2, \dots, N\})$  is a uniform random variable on  $\{1, 2, \dots, N\}$ . Let the discrete random vector  $\tilde{\mathbf{Y}} = \Delta_M(\mathbf{Y})$  be the element-wise quantization of  $\mathbf{Y}$ , where  $\Delta_M(x) := \frac{\lfloor Mx \rfloor}{M}$  is the  $M$ -step discretization of  $[0, 1]$  for some  $M \in \mathbb{N}$ . Additionally define the continuous random vector  $\bar{\mathbf{Y}} = \tilde{\mathbf{Y}} + \mathcal{U}([0, \frac{1}{M}]^K)$ , i. e., independent uniform noise is

<sup>4</sup>Any known compact support suffices. An affine transformation then yields  $\mathcal{X} = [0, 1]^K$ , while possibly resulting in a different Lipschitz constant.

<sup>5</sup>The  $L_1$  norm is only chosen to facilitate subsequent computations. By the equivalence of norms on  $\mathbb{R}^K$ , any norm suffices.

added. Note also that  $h(\bar{\mathbf{Y}}) - H(\tilde{\mathbf{Y}}) = -K \log M$ , where  $H(\cdot)$  denotes Shannon entropy.

We will estimate differential entropy by  $\hat{h}(\mathcal{D}) = h(\bar{\mathbf{Y}}) = H(\tilde{\mathbf{Y}}) - K \log M$ , i.e., the Shannon entropy of the discretized and binned samples with a correction factor.

In the following, we shall also use the two constants

$$\eta(K, L) := \frac{1}{K} \left( \frac{2(K+1)!}{L} \right)^{\frac{1}{K+1}}, \text{ and}$$

$$\alpha := \frac{\sqrt{e^2 + 4} - e}{2e} \approx 0.12074.$$

**Theorem 3.** For  $M \geq \frac{1}{\alpha \eta(K, L)}$  and any  $\delta \in (0, 1)$ , we have with probability greater than  $1 - \delta$  that

$$|\hat{h}(\mathcal{D}) - h(\mathbf{X})| \leq \frac{LK}{2M} \log(M\eta(K, L)) \quad (1a)$$

$$+ \sqrt{\frac{2}{N} \log \frac{2}{\delta}} \log N \quad (1b)$$

$$+ \log \left( 1 + \frac{M^K - 1}{N} \right). \quad (1c)$$

The proof will be given in Appendix A

*Remark 4.* Of the three error terms (1a), (1b), and (1c) the terms (1a) and (1c) constitute the bias and (1b) is a variance-like error term. While the variance (1b) vanishes as  $N \rightarrow \infty$ , the term (1a) does not depend on the sample size  $N$  as it merely measures the error incurred due to the quantization  $\Delta_M$ , which is bounded by the Lipschitz constraint and approaches zero as  $M \rightarrow \infty$ . The final term (1c) results from ensuring that  $N$  samples suffice to suitably approximate the empirical distribution over  $M$  quantization steps. Thus, it ties the quantization to the sample size and approaches zero if  $\frac{M^K}{N} \rightarrow 0$ . In total, the RHS of (1) approaches zero for  $N, M \rightarrow \infty$  provided that  $\frac{M^K}{N} \rightarrow 0$ .

*Remark 5.* Theorem 3 should be regarded as a proof-of-concept rather than a practical tool for performing differential entropy estimation. While analytically tractable, the estimation strategy is crude and the bounds, especially the term (1b) while being a completely universal bound, is known to be loose, as pointed out in [28] p. 1200].

*Remark 6.* We want to note that requiring both a fixed Lipschitz constant  $L$  and a known bounded support, e.g.,  $\mathcal{X} = [0, 1]^K$ , is necessary. Consider for instance the set  $\mathcal{P}' = \{p : p \text{ supported on } [0, 1]^K \text{ and Lipschitz continuous}\}$  of pdfs with arbitrary Lipschitz constant or the set  $\mathcal{P}'' = \{p : p \text{ supported on a bounded set and } L\text{-Lipschitz continuous}\}$  with fixed Lipschitz constant, but arbitrary, bounded support. Both families satisfy the conditions of Proposition 2, i.e., they are convex and  $\sup_{p \in \mathcal{P}'} |h(p)| = \sup_{p \in \mathcal{P}''} |h(p)| = \infty$ .



In principle, Theorem 3 also allows for the approximation of mutual information with a confidence bound. Let  $(\mathbf{X}, \mathbf{Y}) \sim p_{\mathbf{XY}}$  be two random vectors, supported on  $[0, 1]^{K_1}$  and  $[0, 1]^{K_2}$ , respectively. Assuming that  $p_{\mathbf{XY}}$  is  $L$ -Lipschitz continuous on  $\mathbb{R}^{K_1+K_2}$ , it is clear that the marginals  $p_{\mathbf{X}}$  and  $p_{\mathbf{Y}}$  are  $L$ -Lipschitz continuous as well. Thus, Theorem 3 can be used to approximate all three terms in

$$I(\mathbf{X}; \mathbf{Y}) = h(\mathbf{X}) + h(\mathbf{Y}) - h((\mathbf{X}, \mathbf{Y})).$$

## 4 Estimation of other Measures

In this section, we showcase that similar statements as Proposition 2 also hold for mutual information and relative entropy. For simplicity we will not assume  $p_{\mathbf{XY}} \in \mathcal{P}$  for some family  $\mathcal{P}$  of probability density functions, but merely require  $I(\mathbf{X}; \mathbf{Y}), D(\mathbf{X} \parallel \mathbf{Y}) < \infty$ . Only proof sketches are provided as the examples provided in this section are similar to the proof of Proposition 2.

Here we show that in general, it is not possible to accurately estimate mutual information  $I(\mathbf{X}; \mathbf{Y})$  and relative entropy  $D(\mathbf{X} \parallel \mathbf{Y})$  from samples  $(\mathbf{X}, \mathbf{Y})$ .

### 4.1 Mutual Information

For any  $N$ , let  $\hat{i}: \mathbb{R}^N \times \mathbb{R}^N \rightarrow \mathbb{R}$  be a measurable function, which represents an estimate of the mutual information  $I := I(\mathbf{X}; \mathbf{Y}) < \infty$  from  $\mathbf{X}, \mathbf{Y}$ . For convenience we use  $\hat{I} := \hat{i}(\mathbf{X}, \mathbf{Y})$ . Let  $\mathbf{X}, \mathbf{Z} \sim \mathcal{U}([0, 1])$ ,  $W \sim \mathcal{U}([0, e^{-a}])$ , and  $\mathbf{Q} \sim \mathcal{B}(1 - \varepsilon)$  be independent random variables. Define

$$\mathbf{Y} := \mathbf{Q}\mathbf{Z} - (1 - \mathbf{Q})(\mathbf{X} + W).$$

We have

$$\begin{aligned} I(\mathbf{X}; \mathbf{Y}) &= h(\mathbf{Y}) - h(\mathbf{Y}|\mathbf{X}) \\ &\geq H_2(\varepsilon) - H_2(\varepsilon) + a\varepsilon \\ &= a\varepsilon. \end{aligned}$$

The random vectors  $(\mathbf{X}, \mathbf{Y})$   $[(\mathbf{X}, \mathbf{Z})]$  are  $N$  i.i.d. realizations of  $(X, Y)$   $[(X, Z)]$ . For any  $\delta > 0$ , we can find  $b \in \mathbb{R}$  such that  $P\{\hat{i}(\mathbf{X}, \mathbf{Z}) \leq b\} \geq 1 - \frac{\delta}{2}$ . Letting  $\mathcal{E} = \{\mathbf{Y} = \mathbf{Z}\} = \{\mathbf{Q} = \mathbf{1}\}$ , we have  $P\{\mathcal{E}\} \geq 1 - N\varepsilon$ . Thus, when choosing  $\varepsilon \leq \frac{\delta}{2N}$ ,

$$\begin{aligned} P\{\hat{I} \leq b\} &= P\{\mathcal{E}\}P\{\hat{I} \leq b | \mathcal{E}\} + P\{\mathcal{E}^c\}P\{\hat{I} \leq b | \mathcal{E}^c\} \\ &\geq (1 - N\varepsilon)P\{\hat{I} \leq b | \mathcal{E}\} \\ &\geq (1 - N\varepsilon)\left(1 - \frac{\delta}{2}\right) \\ &\geq 1 - \delta. \end{aligned}$$

We may choose  $a \geq \frac{b+C}{\varepsilon}$ . Then, for arbitrary  $C, \delta > 0$  and  $n \in \mathbb{N}$ , we found  $\mathbf{X}$ ,  $\mathbf{Y}$  and  $b \in \mathbb{R}$ , such that  $\mathbb{P}\{\hat{I} > b\} \leq \delta$ , yet  $I \geq b + C$ .

*Remark 7.* Note that [7] Th. 3] claims a confidence bound for mutual information, that together with the approximation result [7] Lem. 1] seemingly contradicts our result. However, the confidence bound proved in [7] Th. 3] requires strong conditions on the functions<sup>6</sup>  $T_\theta$  and [7] Lem. 1] does not necessarily hold under these conditions. Moreover, both approximation results [7] Lem. 1 and Lem. 2] do not hold uniformly for a family of distributions, but implicitly assume a fixed, underlying distribution. This is especially evident in [7] Lem. 2], which also seemingly contradicts our result, when assuming that the optimal function  $T^*$  is in the family  $T_\theta$ . However, this apparent contradiction is resolved by noting that the chosen  $N \in \mathbb{N}$  depends on the underlying, true distribution.

#### 4.1.1 One Discrete Random Variable

In the following we show that a similar result holds if  $\mathbf{Y}$  has a fixed finite alphabet, say  $\mathcal{Y} = \{1, 2, \dots, K\}$ . Again, for every  $N \in \mathbb{N}$ , let  $\hat{i}_N: \mathbb{R}^N \times \mathcal{Y}^N \rightarrow \mathbb{R}$  be an estimator that estimates  $I := I(\mathbf{X}; \mathbf{Y}) \leq \log K$  from  $\mathbf{X}, \mathbf{Y}$ . Note that the result for continuous  $\mathbf{Y}$  cannot carry over unchanged as we have  $I \in [0, \log K]$ . We shall assume that  $\hat{i}_N$  is consistent in the sense that  $\hat{I}_N \rightarrow I$  in probability as  $N \rightarrow \infty$ , where we use  $\hat{I}_N := \hat{i}_N(\mathbf{X}, \mathbf{Y})$ .

Let  $\mathbf{X} \sim \mathcal{U}([0, 1])$  and  $\mathbf{W} \sim \mathcal{U}(\mathcal{Y})$  be independent random variables. Fix  $\delta > 0$  and by consistency find  $N_0$  such that  $\mathbb{P}\{\hat{i}_N(\mathbf{X}, \mathbf{W}) \geq \delta\} \leq \frac{\delta}{2}$  for all  $N \geq N_0$ . In the following consider  $N \geq N_0$  fixed. Fix  $M \in \mathbb{N}$  and  $\mathbf{v} \in \mathcal{Y}^M$ , and define the quantization  $\mathbf{Z} := 1 + \lfloor M\mathbf{X} \rfloor$ . The random variable  $\mathbf{Y}$  is simply  $\mathbf{Y} = \mathbf{v}_\mathbf{Z}$ . We use the notation  $\hat{I}_N^\mathbf{v}$  to highlight that  $\hat{I}_N$  depends on the particular choice of  $\mathbf{v}$  and wish to show that  $\mathbb{P}\{\hat{I}_N^\mathbf{v} \geq \delta\} \leq \delta$  for at least one  $\mathbf{v}$ .

Assume to the contrary, that  $\mathbb{P}\{\hat{I}_N^\mathbf{v} \geq \delta\} > \delta$  for all  $\mathbf{v} \in \mathcal{Y}^M$ . Let  $\mathcal{E} = \{\exists i, j \in \{1, 2, \dots, N\} : i \neq j, Z_i = Z_j\}$  be the event that two elements of  $\mathbf{X}$  fall in the same “bin.” Note that  $\mathbf{Z}$  is the quantization of  $\mathbf{X}$ . For  $M$  large enough, we obtain

$$\begin{aligned} \mathbb{P}\{\mathcal{E}\} &= 1 - \frac{M!}{M^N(M-N)!} \\ &\leq 1 - \left(\frac{M-N+1}{M}\right)^N \\ &\leq \varepsilon. \end{aligned}$$

Defining  $\mathbf{V} \sim \mathcal{U}(\mathcal{Y}^M)$ , independent of  $\mathbf{X}$ , we obtain for  $\varepsilon$  small enough

$$\delta < K^{-M} \sum_{\mathbf{w} \in \mathcal{Y}^M} \mathbb{P}\{\hat{I}_N^\mathbf{v} \geq \delta\}$$

---

<sup>6</sup>Here we use the notation of [7].

$$\begin{aligned}
&\leq \varepsilon + K^{-M} \sum_{\mathbf{v} \in \mathcal{Y}^M} \mathbb{P}\left\{\hat{\mathbf{I}}_N^{\mathbf{v}} \geq \delta \middle| \mathcal{E}^c\right\} \\
&= \varepsilon + \sum_{\mathbf{v} \in \mathcal{Y}^M} \mathbb{P}\left\{\hat{\mathbf{I}}_N^{\mathbf{v}} \geq \delta \middle| \mathcal{E}^c, \mathbf{V} = \mathbf{v}\right\} \mathbb{P}\{\mathbf{V} = \mathbf{v}\} \\
&= \varepsilon + \mathbb{P}\left\{\hat{\mathbf{I}}_N^{\mathbf{V}} \geq \delta \middle| \mathcal{E}^c\right\} \\
&= \varepsilon + \mathbb{P}\left\{\hat{i}_N(\mathbf{X}, \mathbf{V}_{\mathbf{Z}}) \geq \delta \middle| \mathcal{E}^c\right\} \\
&= \varepsilon + \mathbb{P}\left\{\hat{i}_N(\mathbf{X}, \mathbf{W}) \geq \delta \middle| \mathcal{E}^c\right\} \\
&\leq \varepsilon + \frac{\delta}{2(1-\varepsilon)} \leq \delta,
\end{aligned}$$

leading to a contradiction.

To summarize, for arbitrary  $\delta$  and  $N$  large enough, there exists  $\mathbf{w}$  such that  $\mathbb{P}\left\{\hat{\mathbf{I}}_N^{\mathbf{w}} \geq \delta\right\} \leq \delta$ , but clearly  $\mathbf{Y}$  is a deterministic function of  $\mathbf{X}$  and hence  $I = \log K$ .

## 4.2 Relative Entropy

Let  $p$  and  $q$  be two continuous pdfs (w.r.t.  $\lambda$ ) and  $\mathbf{X}, \mathbf{Y}$  be  $N$  i.i.d. random variables distributed according to  $p$  and  $q$ , respectively. For any  $N$ , let  $\hat{d}_N: \mathbb{R}^N \times \mathbb{R}^N \rightarrow \mathbb{R}$  be an estimator that estimates  $D := D(p\|q) < \infty$  from  $\mathbf{X}, \mathbf{Y}$ . For convenience we use  $\hat{D}_N := \hat{d}_N(\mathbf{X}, \mathbf{Y})$ . Let  $\mathbf{Z}_1, \mathbf{Z}_2$  be two independent i.i.d.  $N$  vectors with components uniformly distributed on  $[-1, 0]$ . For an arbitrary  $\delta > 0$  we can find  $c \in \mathbb{R}$  such that  $\mathbb{P}\left\{\hat{d}_N(\mathbf{Z}_1, \mathbf{Z}_2) \leq c\right\} \geq 1 - \frac{\delta}{2}$ . Consider  $C, \delta > 0$  and an arbitrary  $N \in \mathbb{N}$ .

Define the pdfs

$$\begin{aligned}
p(x) &= e^{-a} \mathbb{1}_{[0,1)}(x) + (1 - e^{-a}) \mathbb{1}_{[-1,0)}(x), \text{ and} \\
q(x) &= e^{-a-b} \mathbb{1}_{[0,1)}(x) + (1 - e^{-a-b}) \mathbb{1}_{[-1,0)}(x)
\end{aligned}$$

for  $a, b \in \mathbb{R}_+$ .

With  $b = ke^a$ , where  $k \in \mathbb{R}_+$ , we have  $D = D(a, k)$ , with the function

$$\begin{aligned}
D(a, k) &= e^{-a}b + (1 - e^{-a}) \log \frac{1 - e^{-a}}{1 - e^{-a-b}} \\
&= k + (1 - e^{-a}) \log \frac{1 - e^{-a}}{1 - e^{-a-ke^a}} \\
&\geq k - e^{-1}.
\end{aligned}$$

Let  $\mathcal{E}_1 = \{\mathbf{X} < \mathbf{0}\}$  and  $\mathcal{E}_2 = \{\mathbf{Y} < \mathbf{0}\}$  be the events that every component of  $\mathbf{X}$  and  $\mathbf{Y}$ , respectively, is negative. Then,

$$\mathbb{P}\{\mathcal{E}_1^c \cup \mathcal{E}_2^c\} \leq 2Ne^{-a},$$

$$\mathbb{P}\{\mathcal{E}_1 \cap \mathcal{E}_2\} \geq 1 - 2Ne^{-a}.$$

Choose  $a \geq \log \frac{4N}{\delta}$  and  $k \geq c + C + e^{-1}$  such that  $2Ne^{-a} \leq \frac{\delta}{2}$  and  $D \geq c + C$ . We can now bound the probability

$$\begin{aligned} \mathbb{P}\{\hat{D}_N \leq c\} &= \mathbb{P}\{\hat{D}_N \leq c \mid \mathcal{E}_1 \cap \mathcal{E}_2\} \mathbb{P}\{\mathcal{E}_1 \cap \mathcal{E}_2\} \\ &\quad + \mathbb{P}\{\hat{D}_N \leq c \mid \mathcal{E}_1^c \cup \mathcal{E}_2^c\} \mathbb{P}\{\mathcal{E}_1^c \cup \mathcal{E}_2^c\} \\ &\geq \mathbb{P}\{\hat{D}_N \leq c \mid \mathcal{E}_1 \cap \mathcal{E}_2\} (1 - \frac{\delta}{2}) \\ &= \mathbb{P}\{\hat{d}_N(\mathbf{Z}_1, \mathbf{Z}_2) \leq c\} (1 - \frac{\delta}{2}) \\ &\geq 1 - \delta. \end{aligned}$$

In summary, for an estimator  $\hat{d}_N$  and any  $\delta, C > 0$  and  $N \in \mathbb{N}$ , we can find distributions  $p, q$  and  $c \in \mathbb{R}$ , such that  $\mathbb{P}\{\hat{D}_N > c\} \leq \delta$ , even though  $D \geq c + C$ .

## 5 Discussion and Perspectives

We showed that under mild assumptions on the family of allowed distributions, differential entropy cannot be reliably estimated solely based on samples, no matter how many samples are available. In particular, as first noted in [10] no non-trivial bound or estimate of an information measure can be obtained based only on samples. External information about the regularity of the underlying probability distribution needs to be taken into account. However, such regularity assumptions are not subject to empirical verification and thus, the existence of statistical guarantees for an empirical estimate cannot be empirically tested. This shows that researchers should take great care when approximating or bounding information measures, and specifically explore the necessary assumptions for the underlying distribution.

Regarding the use of information measures in machine learning, we note that our results apply to all estimators of information measures. In particular, empirical versions of variational bounds cannot provide estimates of information measures with high reliability in general.

It would be interesting to investigate the type of assumptions on the underlying distributions that may hold in typical machine learning setups. However, as pointed out previously, these properties cannot be deduced from data, but must result from the model under consideration. In a related note, it might be interesting if the confidence bounds for differential entropy estimation under bounded support and Lipschitz condition from Section 3 carry over to empirical versions of variational bounds. Extensions of these results to other information measures, e. g., Rényi entropy, Rényi divergences, or  $f$ -divergences, could also be of particular interest for future work.

## Acknowledgments

The authors are very grateful to Prof. Elisabeth Gassiat for pointing out the connection between the present work and the reference [10]. This project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 797805.

## A Proof of Theorem 3

We shall first introduce auxiliary random variables which are depicted in Figure 1. Let  $\tilde{\mathbf{X}}$  be the element-wise discretization  $\tilde{\mathbf{X}} = \Delta_M(\mathbf{X})$ . The continuous random vector  $\bar{\mathbf{X}} = \tilde{\mathbf{X}} + \mathcal{U}([0, \frac{1}{M}]^K)$  is obtained by adding independent uniform random noise. Let  $q$  be the pdf of  $\bar{\mathbf{X}}$ . It is straightforward to see that  $\tilde{\mathbf{Y}}$  is distributed according to the empirical distribution of  $N$  i.i.d. copies of  $\tilde{\mathbf{X}}$ . Also note that  $H(\tilde{\mathbf{X}}) = h(\bar{\mathbf{X}}) + K \log M$ .




Figure 1: Connections between all involved random variables.

In order to prove Theorem 3 we use the triangle inequality twice to obtain

$$\begin{aligned}
& |\hat{h}(\mathcal{D}) - h(\mathbf{X})| \\
&= |h(\bar{\mathbf{Y}}) - h(\mathbf{X})| \\
&\leq |h(\bar{\mathbf{Y}}) - h(\bar{\mathbf{X}})| + |h(\bar{\mathbf{X}}) - h(\mathbf{X})| \\
&= |H(\tilde{\mathbf{Y}}) - H(\tilde{\mathbf{X}})| + |h(\bar{\mathbf{X}}) - h(\mathbf{X})| \\
&\leq |H(\tilde{\mathbf{Y}}) - \mathbb{E}[H(\tilde{\mathbf{Y}})]| + |\mathbb{E}[H(\tilde{\mathbf{Y}})] - H(\tilde{\mathbf{X}})| \\
&\quad + |h(\bar{\mathbf{X}}) - h(\mathbf{X})|, \tag{2}
\end{aligned}$$

noting that  $h(\bar{\mathbf{Y}}) - H(\tilde{\mathbf{Y}}) = h(\bar{\mathbf{X}}) - H(\tilde{\mathbf{X}}) = -K \log M$ . Note that  $H(\tilde{\mathbf{Y}})$  is a random quantity that depends on  $\mathcal{D}$ . We thus split the bound in three terms, where the first term in (2) is variance-like and the second and third terms constitute the bias. In the remainder of this appendix, we will complete the proof by showing that all three terms in (2) can be bounded as follows. First,

we have

$$|\mathbb{E}[\mathbf{H}(\tilde{\mathbf{Y}})] - \mathbf{H}(\tilde{\mathbf{X}})| \leq \log\left(1 + \frac{M^K - 1}{N}\right), \text{ and} \quad (3)$$

$$|\mathbf{h}(\tilde{\mathbf{X}}) - \mathbf{h}(\mathbf{X})| \leq \frac{LK}{2M} \log(M\eta(K, L)). \quad (4)$$

And with probability greater than  $1 - \delta$  we also have

$$|\mathbf{H}(\tilde{\mathbf{Y}}) - \mathbb{E}[\mathbf{H}(\tilde{\mathbf{Y}})]| \leq \sqrt{\frac{2}{N} \log \frac{2}{\delta} \log N}. \quad (5)$$

As  $\tilde{\mathbf{Y}}$  is distributed according to the empirical distribution of  $N$  i.i.d. copies of  $\tilde{\mathbf{X}}$ , on an alphabet of size  $M^K$ , the inequalities (3) and (5) follow directly from the following well-known lemma, concerning the estimation of (discrete) Shannon entropy.

**Lemma 4** ([28] eq. (3.4), and Prop. 1] and [3] Remark iii, p.168]). *Let  $Z$  be a random variable on  $\{1, 2, \dots, M\}$  and  $\hat{Z}$  distributed according to the empirical measure of  $N$  i.i.d. copies of  $Z$ . We then have  $|\mathbf{H}(\hat{Z}) - \mathbf{H}(Z)| \leq |\mathbf{H}(Z) - \mathbb{E}[\mathbf{H}(\hat{Z})]| + |\mathbf{H}(\hat{Z}) - \mathbb{E}[\mathbf{H}(\hat{Z})]|$ , where*

$$|\mathbf{H}(Z) - \mathbb{E}[\mathbf{H}(\hat{Z})]| \leq \log\left(1 + \frac{M - 1}{N}\right),$$

and for any  $\delta \in (0, 1]$ , with probability greater than  $1 - \delta$ ,

$$|\mathbf{H}(\hat{Z}) - \mathbb{E}[\mathbf{H}(\hat{Z})]| \leq \sqrt{\frac{2}{N} \log \frac{2}{\delta} \log N}.$$

In order to show (4) we will first obtain some preliminary results and then conclude the proof in Lemma 10. We start by bounding the difference between  $p$  and the approximation  $q$  using the following auxiliary results.

**Lemma 5.** *Let  $f: [\mathbf{0}, \mathbf{A}] \rightarrow \mathbb{R}$  be an arbitrary  $L$ -Lipschitz continuous function and assume  $f(\mathbf{y}) = 0$  for some  $\tilde{\mathbf{y}} \in [\mathbf{0}, \mathbf{A}]$ , then*

$$\int_{[\mathbf{0}, \mathbf{A}]} |f(\mathbf{x})| \lambda^K(d\mathbf{x}) \leq \frac{L}{2} \left( \prod_{k=1}^K A_k \right) \left( \sum_{k=1}^K A_k \right).$$

*In particular, for  $A_k = \varepsilon$ ,  $k \in \{1, 2, \dots, K\}$ , we have  $\int_{[0, \varepsilon]^K} |f(\mathbf{x})| \lambda^K(d\mathbf{x}) \leq \varepsilon^{K+1} \frac{LK}{2}$ .*

---

<sup>7</sup>We use the notation  $[\mathbf{0}, \mathbf{A}] = \{\mathbf{x} \in \mathbb{R}^K : 0 \leq x_k \leq A_k, k \in \{1, 2, \dots, K\}\}$ .

*Proof.* For  $\mathbf{x} \in [0, \mathbf{A}]$  we have  $|f(\mathbf{x})| = |f(\mathbf{x}) - f(\mathbf{y})| \leq L\|\mathbf{x} - \mathbf{y}\|$  and hence

$$\begin{aligned}
& \int_{[0, \mathbf{A}]} |f(\mathbf{x})| \lambda^K(d\mathbf{x}) \\
& \leq L \int_{[0, \mathbf{A}]} \|\mathbf{x} - \mathbf{y}\| \lambda^K(d\mathbf{x}) \\
& = L \sum_{k=1}^K \int_{[0, \mathbf{A}]} |x_k - y_k| \lambda^K(d\mathbf{x}) \\
& = \frac{L}{2} \left( \prod_{k=1}^K A_k \right) \left( \sum_{k=1}^K \frac{A_k^2 - 2y_k(A_k - y_k)}{A_k} \right) \\
& \leq \frac{L}{2} \left( \prod_{k=1}^K A_k \right) \left( \sum_{k=1}^K A_k \right). \quad \square
\end{aligned}$$

**Lemma 6.** For an  $L$ -Lipschitz continuous pdf  $p$  on  $\mathcal{X}$  and  $q$  the pdf of  $\bar{\mathbf{X}}$ , as defined above, we have  $|p(\mathbf{x}) - q(\mathbf{x})| \leq \frac{LK}{2M}$  for every  $\mathbf{x} \in \mathcal{X}$ .

*Proof.* Let  $\mathbf{x} \in \mathcal{X}$  and  $\tilde{\mathbf{x}} = \Delta_M(\mathbf{x})$ . The function  $q$  is constant on  $\Delta_M^{-1}(\tilde{\mathbf{x}})$  and given by  $q(\mathbf{x}') = \lambda^K(\Delta_M^{-1}(\tilde{\mathbf{x}}))^{-1} \int_{\Delta_M^{-1}(\tilde{\mathbf{x}})} p d\lambda^K$  for all  $\mathbf{x}' \in \Delta_M^{-1}(\tilde{\mathbf{x}})$ , where  $\lambda^K(\Delta_M^{-1}(\tilde{\mathbf{x}})) = M^{-K}$ . Thus, since  $\mathbf{x} \in \Delta_M^{-1}(\tilde{\mathbf{x}})$ , we obtain

$$\begin{aligned}
M^{-K}|q(\mathbf{x}) - p(\mathbf{x})| &= \left| \int_{\Delta_M^{-1}(\tilde{\mathbf{x}})} p(\mathbf{x}') - p(\mathbf{x}) \lambda^K(d\mathbf{x}') \right| \\
&\leq \int_{\Delta_M^{-1}(\tilde{\mathbf{x}})} |p(\mathbf{x}') - p(\mathbf{x})| \lambda^K(d\mathbf{x}') \\
&\leq M^{-K-1} \frac{LK}{2}, \tag{6}
\end{aligned}$$

where we applied Lemma 5 to  $\mathbf{x}' \mapsto p(\mathbf{x}') - p(\mathbf{x})$  in (6).  $\square$

**Lemma 7.** If  $p$  is an  $L$ -Lipschitz continuous pdf on  $\mathbb{R}^K$ , then  $p \leq \left( \frac{L^K(K+1)!}{2^K} \right)^{\frac{1}{K+1}}$ .

*Proof.* Let  $\mathbf{x} \in \mathbb{R}^K$  and define the ball  $B_{\mathbf{x}}(\frac{p(\mathbf{x})}{L}) = \{\mathbf{x}' \in \mathbb{R}^K : \|\mathbf{x}' - \mathbf{x}\| \leq \frac{p(\mathbf{x})}{L}\}$  with radius  $\frac{p(\mathbf{x})}{L}$ , centered at  $\mathbf{x}$ . We then have  $\lambda^K(B_{\mathbf{x}}(\frac{p(\mathbf{x})}{L})) = \frac{2^K p(\mathbf{x})^K}{L^K K!}$  and hence,

$$\begin{aligned}
1 &\geq \int_{B_{\mathbf{x}}(\frac{p(\mathbf{x})}{L})} p(\mathbf{x}) \lambda^K(d\mathbf{x}) \\
&= \frac{2^K p(\mathbf{x})^{K+1}}{L^K K!} - \int_{B_{\mathbf{x}}(\frac{p(\mathbf{x})}{L})} p(\mathbf{x}) - p(\mathbf{x}') \lambda^K(d\mathbf{x}') \\
&\geq \frac{2^K p(\mathbf{x})^{K+1}}{L^K K!} - L \int_{B_{\mathbf{x}}(\frac{p(\mathbf{x})}{L})} \|\mathbf{x} - \mathbf{x}'\| \lambda^K(d\mathbf{x}') \tag{7}
\end{aligned}$$

$$= \frac{2^K p(\mathbf{x})^{K+1}}{L^K} \frac{1}{(K+1)!},$$

where the fact that  $p(\mathbf{x}) - p(\mathbf{x}') \leq L\|\mathbf{x} - \mathbf{x}'\|$  for all  $\mathbf{x}' \in \mathbb{R}^K$  is used in (7)  $\square$

Using the previous lemmas to bound the distance between  $p$  and  $q$ , the following two results will allow us to bound the difference of the differential entropies.

**Lemma 8.** For  $x \in [0, 1]$ ,  $y \geq 0$ , and  $a := |x - y| \leq \alpha \approx 0.121$  we have

$$|x \log x - y \log y| \leq -a \log a. \quad (8)$$

*Proof.* In the following, we assume w.l.o.g. that  $x \leq y = x + a$ . If  $x \leq y \leq e^{-1}$ , then  $|x \log x - (x + a) \log(x + a)| = x \log x - (x + a) \log(x + a)$  is monotonically decreasing in  $x$  and thus maximal at  $x = 0$  and hence, (8) follows. If, on the other hand,  $y \geq e^{-1}$ , then necessarily  $\alpha \leq e^{-1} - \alpha \leq x \leq y \leq 1 + \alpha$ . Define the function  $f(x) := -x \log x$ ,  $x > 0$  and  $f(0) := 0$ . Note that by the mean value theorem there are  $a_0 \in (0, a)$  and  $x_0 \in (x, y)$  such that  $|f(0) - f(a)| = f(a) = af'(a_0)$  and  $|f(x) - f(y)| = a|f'(x_0)|$ . Inequality (8) then follows by observing that  $|f'(x_0)| \leq f'(a_0)$  whenever  $a_0 \in (0, \alpha)$  and  $x_0 \in (\alpha, 1 + \alpha)$ .  $\square$

**Lemma 9.** Let  $p$  and  $q$  be two pdfs supported on  $\mathcal{X}$  with finite differential entropies. Assume that for all  $\mathbf{x} \in \mathcal{X}$  we have  $|p(\mathbf{x}) - q(\mathbf{x})| \leq \varepsilon$  and  $0 \leq p(\mathbf{x}) \leq A$ , and that  $\frac{\varepsilon}{A} \leq \alpha$  holds. Then,

$$|h(p) - h(q)| \leq \varepsilon \log \frac{A}{\varepsilon}. \quad (9)$$

*Proof.* Define  $p'(\mathbf{x}) := A^{-1}p\left(A^{-\frac{1}{K}}\mathbf{x}\right)$  and  $q'(\mathbf{x}) := A^{-1}q\left(A^{-\frac{1}{K}}\mathbf{x}\right)$  for  $\mathbf{x} \in [0, A^{\frac{1}{K}}]^K$ . We have  $|p'(\mathbf{x}) - q'(\mathbf{x})| \leq A^{-1}\varepsilon$ ,  $0 \leq p'(\mathbf{x}) \leq 1$  as well as

$$\begin{aligned} |h(p) - h(q)| &= |h(p') - h(q')| \\ &= \left| \int p' \log p' - q' \log q' d\lambda^K \right| \\ &\leq \int |p' \log p' - q' \log q'| d\lambda^K \\ &\leq -\lambda^K([0, A^{\frac{1}{K}}]^K) \cdot A^{-1}\varepsilon \log(A^{-1}\varepsilon) \\ &= \varepsilon \log \frac{A}{\varepsilon}, \end{aligned} \quad (10)$$

where Lemma 8 was applied in (10)  $\square$

We can now finish the proof of Theorem 3 by showing (4)



**Lemma 10.** *If  $M \geq \frac{1}{\alpha\eta(K,L)}$  we have*

$$|\mathbf{h}(\overline{\mathbf{X}}) - \mathbf{h}(\mathbf{X})| \leq \frac{LK}{2M} \log(M\eta(K, L)).$$

*Proof.* By Lemma 6  $|\mathbf{p}(\mathbf{x}) - \mathbf{q}(\mathbf{x})| \leq \frac{LK}{2M}$  and, by Lemma 7  $\mathbf{p} \leq \left(\frac{L^K(K+1)!}{2^K}\right)^{\frac{1}{K+1}}$ . We can thus apply Lemma 9 with  $\varepsilon = \frac{LK}{2M}$  and  $A = \left(\frac{L^K(K+1)!}{2^K}\right)^{\frac{1}{K+1}}$  provided that  $\frac{\varepsilon}{A} \leq \alpha$ , which is equivalent to  $M \geq \frac{1}{\alpha\eta(K,L)}$ . Inserting  $\varepsilon$  and  $A$  in (9) proves the result.  $\square$

## References

- [1] I. Ahmad and P.-E. Lin. A nonparametric estimation of the entropy for absolutely continuous distributions. *IEEE Trans. Inf. Theory*, 22(3):372–375, May 1976.
- [2] R. A. Amjad and B. C. Geiger. Learning representations for neural network-based classification using the information bottleneck principle. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. to appear.
- [3] A. Antos and I. Kontoyiannis. Convergence properties of functional estimates for discrete distributions. *Random Structures & Algorithms*, 19(3-4):163–193, Nov. 2001.
- [4] R. R. Bahadur and L. J. Savage. The nonexistence of certain statistical procedures in nonparametric problems. *Ann. Math. Statist.*, 27(4):1115–1122, 1956.
- [5] D. Barber and F. Agakov. The IM algorithm: A variational approach to information maximization. In *NIPS’03*, volume 16 of *Advances in neural information processing systems*, pages 201–208, Cambridge, MA, USA, Dec. 2003.
- [6] J. Beirlant, E. J. Dudewicz, L. Györfi, and E. C. Van der Meulen. Nonparametric entropy estimation: An overview. *International Journal of Mathematical and Statistical Sciences*, 6(1):17–39, 1997.
- [7] M. I. Belghazi, A. Baratin, S. Rajeshwar, S. Ozair, Y. Bengio, A. Courville, and D. Hjelm. Mutual information neural estimation. In *ICML’18*, volume 80 of *PMLR*, pages 531–540, Stockholm, Sweden, July 2018.
- [8] X. Chen, Y. Duan, R. Houthoofd, J. Schulman, I. Sutskever, and P. Abbeel. InfoGAN: Interpretable representation learning by information maximizing generative adversarial nets. In *NIPS’16*, volume 29 of *Advances in neural information processing systems*, pages 2172–2180, Barcelona, Spain, 2016.
- [9] G. A. Darbellay and I. Vajda. Estimation of the information by an adaptive partitioning of the observation space. *IEEE Trans. Inf. Theory*, 45(4):1315–1321, May 1999.
- [10] D. L. Donoho. One-sided inference about functionals of a density. *Ann. Statist.*, 16(4):1390–1420, 1988.
- [11] D. L. Donoho and R. C. Liu. Geometrizing rates of convergence, II. *Ann. Statist.*, 19(2):633–667, 1991.
- [12] W. Gao, S. Kannan, S. Oh, and P. Viswanath. Estimating mutual information for discrete-continuous mixtures. In *NIPS’17*, volume 30 of *Advances in Neural Information Processing Systems*, pages 5986–5997, Long Beach, CA, USA, 2017.

- [13] L. J. Gleser and J. T. Hwang. The nonexistence of 100  $(1-\alpha)\%$  confidence sets of finite expected diameter in errors-in-variables and related models. *Ann. Statist.*, pages 1351–1362, 1987.
- [14] Z. Goldfeld, K. Greenewald, Y. Polyanskiy, and J. Weed. Convergence of smoothed empirical measures with applications to entropy estimation. *arXiv preprint*, 2019.
- [15] Z. Goldfeld, E. Van Den Berg, K. Greenewald, I. Melnyk, N. Nguyen, B. Kingsbury, and Y. Polyanskiy. Estimating information flow in deep neural networks. In *ICML’19*, volume 97 of *PMLR*, pages 2299–2308, Long Beach, CA, USA, 2019.
- [16] S. Gordon, H. Greenspan, and J. Goldberger. Applying the information bottleneck principle to unsupervised clustering of discrete and continuous image representations. In *Proc. Ninth IEEE Int. Conf. Comput. Vision*, pages 370–377, Nice, France, Oct. 2003.
- [17] L. Györfi and E. C. Van der Meulen. Density-free convergence properties of various estimators of entropy. *Comput. Stat. Data Anal.*, 5(4):425–436, Sept. 1987.
- [18] Y. Han, J. Jiao, T. Weissman, and Y. Wu. Optimal rates of entropy estimation over Lipschitz balls. *arXiv preprint*, 2019.
- [19] R. D. Hjelm, A. Fedorov, S. Lavoie-Marchildon, K. Grewal, P. Bachman, A. Trischler, and Y. Bengio. Learning deep representations by mutual information estimation and maximization. In *ICLR*, New Orleans, LA, USA, 2019.
- [20] W. Hu, T. Miyato, S. Tokui, E. Matsumoto, and M. Sugiyama. Learning discrete representations via information maximizing self-augmented training. In *ICML’17*, volume 70 of *PMLR*, pages 1558–1567, Sydney, Australia, 2017.
- [21] K. Kandasamy, A. Krishnamurthy, B. Poczos, L. Wasserman, and J. M. Robins. Nonparametric von Mises estimators for entropies, divergences and mutual informations. In *NIPS’15*, volume 28 of *Advances in Neural Information Processing Systems*, pages 397–405, Montréal, Canada, 2015.
- [22] A. Kolchinsky, B. D. Tracey, and S. V. Kuyk. Caveats for information bottleneck in deterministic scenarios. In *ICLR*, New Orleans, LA, USA, 2019.
- [23] H. Liu, L. Wasserman, and J. D. Lafferty. Exponential concentration for mutual information estimation with application to forests. In *NIPS’12*, volume 26 of *Advances in Neural Information Processing Systems*, pages 2537–2545, Lake Tahoe, NV, USA, 2012.

- [24] D. McAllester and K. Stratos. Formal limitations on the measurement of mutual information. In *ICLR*, New Orleans, LA, USA, 2019.
- [25] T. Miyato, S.-i. Maeda, M. Koyama, K. Nakae, and S. Ishii. Distributional smoothing with virtual adversarial training. In *ICLR*, San Juan, Puerto Rico, 2016.
- [26] I. Nemenman, W. Bialek, and R. D. R. Van Steveninck. Entropy and information in neural spike trains: Progress on the sampling problem. *Phys. Rev. E*, 69(5):056111, 2004.
- [27] X. Nguyen, M. J. Wainwright, and M. I. Jordan. Estimating divergence functionals and the likelihood ratio by convex risk minimization. *IEEE Trans. Inf. Theory*, 56(11):5847–5861, Nov. 2010.
- [28] L. Paninski. Estimation of entropy and mutual information. *Neural Comput.*, 15(6):1191–1253, 2003.
- [29] J. Pfanzagl. The nonexistence of confidence sets for discontinuous functionals. *J. Stat. Plan. Inference*, 75(1):9–20, 1998.
- [30] B. Poole, S. Ozair, A. van den Oord, A. A. Alemi, and G. Tucker. On variational bounds of mutual information. In *ICML’19*, volume 97 of *PMLR*, pages 5171–5180, Long Beach, CA, USA, 2019.
- [31] C. E. Shannon. A mathematical theory of communication. *Bell Syst. Tech. J.*, 27(3):379–423, July 1948.
- [32] C. E. Shannon. Prediction and entropy of printed English. *Bell Syst. Tech. J.*, 30(1):50–64, 1951.
- [33] R. Shwartz-Ziv and N. Tishby. Opening the black box of deep neural networks via information. *arXiv preprint*, 2017.
- [34] S. Singh and B. Poczos. Generalized exponential concentration inequality for renyi divergence estimation. In *ICML’14*, volume 32 of *PMLR*, pages 333–341, Beijing, China, 2014.
- [35] S. Singh and B. Poczos. Finite-sample analysis of fixed-k nearest neighbor density functional estimators. In *NIPS’16*, volume 29 of *Advances in Neural Information Processing Systems*, pages 1217–1225, Barcelona, Spain, 2016.
- [36] K. Sricharan, R. Raich, and A. O. Hero. k-nearest neighbor estimation of entropies with confidence. In *Proc. IEEE Int. Symp. Inf. Theory (ISIT 2011)*, pages 1205–1209, Saint Petersburg, Russia, July 2011.
- [37] N. Tishby, F. C. Pereira, and W. Bialek. The information bottleneck method. In *Annu. Allerton Conf. Commun., Control, and Comput.*, pages 368–377, Monticello, IL, Sept. 1999.

- [38] A. van den Oord, Y. Li, and O. Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint*, 2018.
- [39] Q. Wang, S. R. Kulkarni, and S. Verdú. Divergence estimation of continuous distributions based on data-dependent partitions. *IEEE Trans. Inf. Theory*, 51(9):3064–3074, Sept. 2005.