

# Local Intrinsic Dimensional Entropy

Rohan Ghosh<sup>1</sup>, Mehul Motani<sup>1,2</sup>

<sup>1</sup> College of Design and Engineering, Department of Electrical and Computer Engineering, National University of Singapore

<sup>2</sup> N.1 Institute for Health, Institute for Digital Medicine (WisDM), Institute of Data Science, National University of Singapore  
rghosh92@gmail.com, motani@nus.edu.sg

## Abstract

Most entropy measures depend on the spread of the probability distribution over the sample space  $\mathcal{X}$ , and the maximum entropy achievable scales proportionately with the sample space cardinality  $|\mathcal{X}|$ . For a finite  $|\mathcal{X}|$ , this yields robust entropy measures which satisfy many important properties, such as invariance to bijections, while the same is not true for continuous spaces (where  $|\mathcal{X}| = \infty$ ). Furthermore, since  $\mathbb{R}$  and  $\mathbb{R}^d$  ( $d \in \mathbb{Z}^+$ ) have the same cardinality (from Cantor’s correspondence argument), cardinality-dependent entropy measures cannot encode the data dimensionality. In this work, we question the role of cardinality and distribution spread in defining entropy measures for continuous spaces, which can undergo multiple rounds of transformations and distortions, e.g., in neural networks. We find that the average value of the local intrinsic dimension of a distribution, denoted as ID-Entropy, can serve as a robust entropy measure for continuous spaces, while capturing the data dimensionality. We find that ID-Entropy satisfies many desirable properties and can be extended to conditional entropy, joint entropy and mutual-information variants. ID-Entropy also yields new information bottleneck principles and also links to causality. In the context of deep learning, for feedforward architectures, we show, theoretically and empirically, that the ID-Entropy of a hidden layer directly controls the generalization gap for both classifiers and auto-encoders, when the target function is Lipschitz continuous. Our work primarily shows that, for continuous spaces, taking a structural rather than a statistical approach yields entropy measures which preserve intrinsic data dimensionality, while being relevant for studying various architectures.

## Introduction and Motivation

In this paper, we consider an alternative interpretation of what outlines a good measure of information, particularly in the context of sample spaces which can repeatedly undergo continuous transformations of any kind. This applies to domains where data can undergo multiple rounds of processing and distortion, for example, in machine learning with deep neural networks. Our objective here is to find whether a robust information measure can be defined in this context, which ideally satisfies the various properties exhibited by discrete entropy, which is defined for discrete spaces.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Entropy for discrete sample spaces  $\mathcal{X}$  ( $|\mathcal{X}| < \infty$ ) is essentially a fundamental property of the space. If one were to map the elements of a finite  $\mathcal{X}$  to another space  $\mathcal{X}'$  using a bijection,  $\mathcal{X}'$  will have the same entropy as  $\mathcal{X}$ . Therefore, bijective transformations such as shift and scaling will yield the same discrete entropy. Furthermore, the notion of *surprise*, which is often associated with entropy, aptly applies to the case of finite  $\mathcal{X}$ , as a larger spread of the distribution essentially points to less certainty about the outcome of any random draw from  $\mathcal{X}$ . It is also worth noting that any deterministic function, which maps a finite  $\mathcal{X}$  to another finite space  $\mathcal{X}'$ , can only reduce discrete entropy, as deterministic maps should not increase the amount of information.

It is interesting to note then, that the counterpart of entropy measures for continuous spaces, such as differential entropy and other variants (Renyi 1960; Shannon 1948), don’t satisfy the same desirable properties as discrete entropy. First, we note that differential entropy can change significantly in response to trivial bijective maps, such as simply scaling the axes by a factor of  $\alpha$ . In fact, for a continuous random variable (RV)  $X \in \mathbb{R}^d$ , we have that  $h(\alpha X) = h(X) + \log |\alpha|$ , where  $h(\cdot)$  denotes the differential entropy. Second, we reason that the notion of distribution spread cannot be used as a measure of surprise for continuous distributions. This is because if one were to simply scale the axes by  $\alpha$ , we can directly tune the spread of the distribution. By doing so, one would be able to yield arbitrarily higher measures of surprise by simply processing the RV  $X$  with a bijective map such as scaling. Note that we cannot do this for discrete entropy, as bijections such as scaling and translation do not alter the discrete entropy of the space. Third, similar to our previous arguments, we see that simply scaling the variable to have a much higher spread can increase differential entropy arbitrarily, and thus processing can increase differential entropy, contrary to how discrete entropy behaves. Note that similar conclusions would follow even if one were to estimate the differential entropy after normalizing the variance of  $X$ , by replacing the scaling transformation with any other scale-preserving distortion.

**Towards Local Measures of Surprise:** We next consider how to quantify surprise for continuous spaces. We start by noting that maximum entropy for a discrete set  $\mathcal{X}$  is  $\log |\mathcal{X}|$ , and thus depends on the cardinality of the set. In contrast, for continuous spaces of  $d$  dimensions  $\mathbb{R}^d$ , the cardinality is ac-

tually invariant to the dimension  $d$ . As observed first by Cantor in his 1874 work (Cantor 1874), there exists a bijective transformation between  $\mathbb{R}^{d_1}$  and  $\mathbb{R}^{d_2}$  for any  $d_1, d_2 \in \mathbb{Z}^+$ , indicating that  $\mathbb{R}^{d_1}$  and  $\mathbb{R}^{d_2}$  have the same cardinality. These observations point to a potentially different interpretation of what an entropy measure could signify in the context of continuous spaces. We see that both cardinality and spread cannot be used as a basis for computing any measure of *surprise*. Furthermore, we observe that the notion of *surprise* itself cannot be immediately defined in this case, as the true surprise  $\lim_{dx \rightarrow 0} \log \frac{1}{P(x)dx}$ , for any element  $x \in \mathcal{X}$  would be infinite, as  $dx \rightarrow 0$  implies  $P(x)dx \rightarrow 0$  as well (except for Dirac-delta distributions).

This leads us to question whether there are alternate ways to quantify surprise for continuous spaces. To do so, we define the *covering* of a distribution  $P(X)$ , as the minimal number of spheres  $B_0, B_1, \dots, B_K$  of fixed radius  $\epsilon$ , which together cover the entire support of  $P(X)$ . By doing so, we can effectively treat  $P(X)$  as a discrete distribution of cardinality  $K$ , and we can define surprise in the conventional *log-probability* sense. Next, for any point  $X_0$  which lies in the support of  $P(X)$ , we define a *local-neighborhood*  $\mathcal{N}(X_0)$  centered at  $X_0$ , which is a sphere of radius  $\delta$ , s.t.  $\delta = \tau\epsilon$  for some constant  $\tau > 1$ . Let the number of spheres among  $B_0, B_1, \dots, B_K$  that lie inside  $\mathcal{N}(X_0)$ , be denoted by  $n$ . As  $\epsilon \rightarrow 0$ , we can easily show that  $n \rightarrow (\delta/\epsilon)^{d(X_0)}$ , where  $d(X_0)$  is the dimension of the distribution manifold around  $X_0$ , also called the local intrinsic dimension (Bac and Zinovyev 2020). As such, the *local surprise* converges to  $\log n = \log(\delta/\epsilon)^{d(X_0)} = d(X_0) \log(\delta/\epsilon) = d(X_0) \log(\tau)$ . Furthermore, the expected value of local surprise is  $\mathbb{E}_{X \sim P} [d(X)] \log(\tau)$ . The quantity  $\mathbb{E}_{X \sim P} [d(X)]$  is what we propose to analyze in this work. We term this quantity as the *local intrinsic dimensional entropy* (ID-Entropy) of the distribution, and find that it satisfies all of the earlier discussed properties of discrete entropy. Furthermore, we also find that ID-entropy can be relevant in the context of analyzing the information in deep neural network layers, and predicting their generalization behaviour, which is discussed next.

**Information Bottlenecks and Generalization:** The *Information Bottleneck* (IB) principle, originally proposed in (Tishby and Zaslavsky 2015), discusses the relevance of quantities such as the mutual information (MI) between the input  $X$  and features  $T$ , denoted by  $I(X; T)$ , and between the features and labels  $Y$ , denoted by  $I(T; Y)$  to generalization performance. Broadly, the IB principle hypothesizes that neural networks undergo two phases during learning: (a) a *fitting* phase where both  $I(X; T)$  and  $I(Y; T)$  increases, and (b) a *compression* phase where  $I(X; T)$  decreases. However, the true MI  $I(X; T)$  is actually infinity, as given a fixed network configuration,  $T$  is a deterministic function of  $X$ , and  $I(X; f(X)) = \infty$  for deterministic  $f$  (Goldfeld and Polyanskiy 2020). Furthermore, these two phases have not always been observed in networks (Saxe et al. 2018). In this work, we show that ID-Entropy of the feature layer  $T$  of a neural network, denoted by  $\mathfrak{h}(T)$ , can be analyzed similarly, leading to analogous bottleneck crite-

ria, and new generalization error bounds. Furthermore, unlike the vacuous nature of the estimate of  $I(X; T)$ ,  $\mathfrak{h}(T)$  is non-zero and finite, and has *opposite trends* during training compared to its IB counterpart. From ID-Entropy’s perspective, we find that the network undergoes compression first, as  $\mathfrak{h}(T)$  reduces initially and increases slowly thereafter. Furthermore, both in theory and practice, we find that if the network converges to a higher  $\mathfrak{h}(T)$  after training, it shows worse generalization performance.

## Contributions

The following are the specific contributions of our work.

- 1. Definition and Properties:** A novel measure of entropy for continuous domains called *ID-Entropy* is proposed. Unlike differential entropy and other counterparts, *ID-Entropy* is motivated from the observation that the intrinsic dimensionality of the data relates to a local measure of information gain, yielding a new notion of surprise for continuous spaces. We find that *ID-Entropy* satisfies many properties of discrete Shannon entropy itself.
- 2. New Information Bottleneck Criteria:** *ID-Entropy* measures leads to a new notion of information bottlenecks for auto-encoder and classifier architectures. Unlike conventional bottlenecks where the notion of MI and differential entropy can be vacuous, *ID-Entropy* avoids this problem, while yielding new perspectives on the problem.
- 3. Connection to Generalization:** We find that the *ID-Entropy* of input and hidden layers for feedforward architectures in large dataset scenarios controls the generalization gap for both auto-encoders and classifier architectures.
- 4. Connection to Causality:** *ID-Entropy* encodes the underlying causal factors responsible for generating any given dataset, when considering a class of causal models where causal effects are generated by continuous functions. This pertains to the cases where, if  $X$  is the only cause of  $Y$ ,  $Y$  is a continuous function of  $X$  (but not necessarily vice-versa).
- 5. Experimental Validation:** We empirically compute the *ID-Entropy* of feature layers for feed-forward architectures such as Convolutional Neural Networks (CNNs) and Convolutional Auto-Encoders trained on MNIST and CIFAR-10. For both auto-encoders and classifiers, we find that the estimated ID-Entropy correlates with the generalization gap.

## *ID-Entropy: A New Perspective on Entropy*

### Assumptions and Definitions

In this section, we present some useful definitions and assumptions, building up to *ID-Entropy*.

Let  $X$  be a random variable (RV) and  $P(X)$  be the distribution of  $X$ . We first enumerate some useful notions. **(A1)** For a RV  $X$ , the support of  $P(X)$  can be expressed as a union of a finite number of *locally Euclidean* (Lee 2010) manifolds, where locally Euclidean means that every point in the support of  $P(X)$  has a local neighborhood which is homeomorphic to  $\mathbb{R}^k$  for some non-negative integer  $k$ . **(A2)** For a set of RVs  $\mathbb{X} = \{X_1, X_2, \dots, X_n\}$  and for any  $S \subseteq \mathbb{X}$ , the geodesic distance between any two connected points in  $P(S)$  is finite. **(A3)** For a RV  $X$ , the number of connected components in the support of  $P(X)$  is countable. Unless

otherwise mentioned, all RVs in this work are assumed to satisfy A1, A2, and A3. These assumptions are not very stringent, and we expect them to be satisfied for most cases. We provide a more detailed discussion regarding the same in the Appendix, available in (Ghosh and Motani 2023).

Next, we define a novel relation between two RVs called *f-relatedness*. Note that f-relatedness is used in the context of vector fields with a different meaning (Wikipedia 2022).

**Definition 1. (*f-relatedness*)** We say that a RV  $Y$  is *f-related* to another RV  $X$ , if there exists a continuous function  $f$  such that  $f(X) = Y$ . This is denoted by  $X \rightsquigarrow Y$ . We also refer to this as an *f-relation*. When  $Y$  is not *f-related* to  $X$ , we denote this as  $X \not\rightsquigarrow Y$ .

Note that two RVs  $X$  and  $Y$  are either *f-related* or they are not. Furthermore, *f-relatedness* is asymmetric, i.e.,  $X \rightsquigarrow Y$  does not imply  $Y \rightsquigarrow X$ . Next, we define *symmetric f-relations*.

**Definition 2. (*Symmetric f-relation*)** We say that RV's  $X$  and  $Y$  have a *symmetric f-relation* when  $X \rightsquigarrow Y$  and  $Y \rightsquigarrow X$ . A *symmetric f-relation* is denoted as  $X \leftrightarrow Y$ .

Note that both *f-relations* and *symmetric f-relations* can be extended to the multiple RVs scenario. One can similarly define  $(X_1, X_2, \dots, X_k) \rightsquigarrow Y$  and  $X \rightsquigarrow (Y_1, Y_2, \dots, Y_k)$  for RVs  $X_1, X_2, \dots, X_k$  and  $Y_1, Y_2, \dots, Y_k$ . We note that  $(X_1, X_2, \dots, X_k) \rightsquigarrow Y$  does not imply  $X_i \rightsquigarrow Y$  for any  $i$ .

Next, we define the  $\epsilon$ -*Neighborhood Intrinsic Dimension* of a RV  $X$ , at any point  $\rho \in \mathbb{R}^d$ , as follows. First, for some  $0 \leq \epsilon \leq \infty$ , let us consider another RV  $X_\epsilon^\rho$ , which follows a distribution  $P_\epsilon^\rho$  given by:

$$P_\epsilon^\rho(X) = \begin{cases} \frac{P(X)}{c}, & \text{if } \|X - \rho\| \leq \epsilon \\ 0, & \text{otherwise} \end{cases}$$

Here,  $c$  is a constant which scales the distribution to ensure  $\int P_\epsilon^\rho(X) dX = 1$ .

Given this, we define the  $\epsilon$ -*Neighborhood Intrinsic Dimension*  $d_\epsilon(\rho)$  as follows:

$$d_\epsilon(\rho) = \min n \text{ s.t. } (v_1, v_2, \dots, v_n) \leftrightarrow X_\epsilon^\rho \quad (1)$$

where  $v_1, v_2, \dots, v_N$  represent either continuous RVs such that  $0 \leq v_i \leq 1$ , or binary RVs such that  $v_i \in \{0, 1\}$ . Thus, in other words,  $d_\epsilon(\rho)$  represents the minimum number of total continuous and/or discrete RVs of any arbitrary dense distribution  $Q$ , which can be associated with  $X_\epsilon^\rho$  via a homeomorphism, i.e., infinitesimal changes in  $v_1, v_2, \dots, v_N$  always yields infinitesimal changes in  $X_\epsilon^\rho$  and vice versa. Note that  $d_\epsilon(\rho)$  will be finite for small  $\epsilon$  due to the locally Euclidean assumption. Thus, note that cycling through all possible values of the RVs and remapping them to  $\mathbb{R}^d$  using the appropriate continuous map should precisely yield the set of all possible values of the RV  $X_\epsilon^\rho$ .

With this, we can define  $\epsilon$ -*ID-Entropy*  $\mathfrak{h}^\epsilon(X)$  of a RV  $X$  with an underlying distribution  $P(X)$  as follows:

**Definition 3. ( $\epsilon$ -*ID-Entropy*)** We define the  $\epsilon$ -*ID-Entropy*  $\mathfrak{h}^\epsilon(X)$  of a RV  $X$  as  $\mathfrak{h}^\epsilon(X) = \mathbb{E}_{\rho \sim P(X)} [d_\epsilon(\rho)]$ , where  $P(X)$  is the underlying distribution of  $X$ .

---

#### Algorithm 1: Estimation of ID-Entropy

---

**Input:**  $S = \{X_1, \dots, X_m\}$  (i.i.d samples of RV  $X$ ), a global-ID estimator  $f_{ID}(S')$  of points in  $S'$ , & parameters  $(k, n)$ .

**Output:**  $ID_X$  (Estimate of ID-Entropy of  $X$ )

- 1:  $ID_{sum} = 0$
  - 2: **for**  $j = 1, \lfloor 2m/n \rfloor, \lfloor 3m/n \rfloor, \dots, m$  **do**
  - 3:     Let  $S' = k$ -nearest neighbors of  $X_j$  in  $S$
  - 4:      $ID_{sum} = ID_{sum} + f_{ID}(S')$
  - 5:  $ID_X = ID_{sum}/n$
- 

Note that the above formulation represents the *ID-Entropy* for a finite fixed  $\epsilon > 0$ , and, as such, one can similarly consider all possible values of  $0 < \epsilon < \infty$ . Next, with this, we can define the *ID-Entropy*  $\mathfrak{h}(X)$  of  $X$  as follows.

**Definition 4. (*ID-Entropy*)** We define the *ID-Entropy*  $\mathfrak{h}(X)$  of a RV  $X$  as

$$\mathfrak{h}(X) = \lim_{\epsilon \rightarrow 0^+} \mathfrak{h}^\epsilon(X). \quad (2)$$

Note that the above definition can be extended to the case of multiple RVs, just by considering their concatenated form as the new RV. Thus, the joint *ID-Entropy* of RV  $X \in \mathbb{R}^d$  and  $Y \in \mathbb{R}^l$  can be defined as follows:

$$\mathfrak{h}(X, Y) = \mathfrak{h}(V), \quad (3)$$

where  $V \in \mathbb{R}^{d+l}$  is simply a concatenation of the two RVs  $X$  and  $Y$ . Note that one can similarly define *mutual ID-information* and *conditional ID-Entropy* from the definition of *ID-Entropy*, in the same way as their traditional counterparts. They are defined as follows.

**Definition 5. (*Mutual ID-information*)** Given RVs  $X$  and  $Y$ ,  $\mathcal{I}(X; Y) = \mathfrak{h}(X) + \mathfrak{h}(Y) - \mathfrak{h}(X, Y)$ . Note that  $\mathcal{I}(X; Y) \geq 0$ .

**Definition 6. (*Conditional ID-Entropy*)** Given RVs  $X$  and  $Y$ ,  $\mathfrak{h}(Y|X) = \mathfrak{h}(X, Y) - \mathfrak{h}(X)$ . Note that  $\mathfrak{h}(Y|X) \geq 0$ , and  $\mathfrak{h}(Y|X) = 0$  iff  $X \rightsquigarrow Y$ .

Algorithm 1 shows how to estimate *ID-Entropy*, making use of a global-ID estimator, denoted by  $f_{ID}$ .

#### Properties of *ID-Entropy*

We outline the properties of the proposed *ID-Entropy* measure and its extensions. The proofs of all the properties, and all subsequent theoretical results in this paper, are given in the Appendix in (Ghosh and Motani 2023). In what follows,  $\dim(X)$  represents the dimensionality of  $X$ 's co-domain.

**P1**  $0 \leq \mathfrak{h}(X) \leq \dim(X)$  (Finite and Bounded)

**P2**  $\mathfrak{h}(X) \leq \mathfrak{h}^\epsilon(X)$  for any  $\epsilon > 0$ .

**P3**  $\mathfrak{h}(\alpha X) = \mathfrak{h}(X)$ , for any real  $\alpha$  (Scale-invariance).

**P4**  $\mathfrak{h}(X, Y) = \mathfrak{h}(Y, X)$  (Symmetric).

**P5**  $\mathfrak{h}(X, Y) \geq \mathfrak{h}(X)$  and  $\mathfrak{h}(X, Y) \geq \mathfrak{h}(Y)$ .

**P6**  $\mathfrak{h}(X, Y) \leq \mathfrak{h}(X) + \mathfrak{h}(Y)$ .

**P7** If  $X \rightsquigarrow Y$ ,  $\mathfrak{h}(X, Y) = \mathfrak{h}(X) \geq \mathfrak{h}(Y)$ . (Cannot increase with continuous processing)

**P8** If  $X \rightsquigarrow Y \rightsquigarrow Z$ , then we have that  $\mathcal{I}(X; Y) \geq \mathcal{I}(X; Z)$ . (DPI for Continuous Maps).

**P9** If  $X \rightarrow Y \rightarrow Z$  represents a general Markov Chain, then we have that  $\mathcal{I}(X; Y) \geq \mathcal{I}(X; Z)$ . (DPI for General Markov Chains).

**P10** If  $X \leftrightarrow Y$ ,  $\mathfrak{h}(X) = \mathfrak{h}(Y)$ .

**P11** (Sampling Invariance) Let  $X \in \mathbb{R}^d$  be represented as  $d$  RVs  $(X_1, X_2, \dots, X_d)$ . Consider functions  $f_1, f_2, \dots, f_d$  all from  $\mathbb{R} \rightarrow \mathbb{R}$ , such that they are order preserving and continuous ( $X < Y$  implies  $f_i(X) < f_i(Y)$ ). Then we have that,  $\mathfrak{h}(X) = \mathfrak{h}(X_1, X_2, \dots, X_d) = \mathfrak{h}(f_1(X_1), f_2(X_2), \dots, f_d(X_d))$ . We denote this property of *ID-Entropy* as *sampling-invariance*, where  $f_1, f_2, \dots, f_n$  represent the re-sampling functions.

**Remark 1.** We note that *ID-Entropy* shares many of the desirable properties of discrete Shannon entropy. This includes P4, P5, P6, P8 and P9. In the Appendix, available in (Ghosh and Motani 2023), we also define a tighter variant of  $\epsilon$ -*ID-Entropy*, called  $\epsilon$ -*log-ID-Entropy*, which more smoothly connects to discrete Shannon entropy.

**Remark 2.** We see that *ID-Entropy* is sampling-invariant (P11), thus  $\mathfrak{h}(X)$ , does not change when the individual 1- $d$  RVs within  $X$  undergo separate order-preserving sampling functions. This is a desirable when  $X$  represents real data sampled through sensors, as this ensures that  $\mathfrak{h}(X)$  does not change in response to different choices of sensors with slightly different sensitivities, as long as they are order preserving. For instance, for vision sensors, one can potentially have many possible ways by which the pixels respond to brightness changes. As long as lower brightness gets mapped to a lower response, the *ID-Entropy* of the sampled signal stays the same.

**Remark 3.** We refer to our initial discussion regarding the observation that  $\mathbb{R}^d$  and  $\mathbb{R}$  are equivalent in terms of cardinality, thus  $\mathbb{R}^d$  and  $\mathbb{R}$  cannot be differentiated cardinality-wise. However, as  $\mathbb{R}$  is not homeomorphic to  $\mathbb{R}^d$ , only via their topological properties we can differentiate  $\mathbb{R}^d$  and  $\mathbb{R}$ . If we consider distributions  $P(X)$  dense in  $\mathbb{R}^d$ , then it will have an *ID-Entropy* of  $d$ , as any dense distribution is homeomorphic to  $\mathbb{R}^d$ . Applying this to the case of  $d = 1$ , we see that the *ID-Entropies* of all 1- $D$  distributions with dense support is 1, contrary to differential entropy, which heavily depends on the shape and statistics of the distribution.

### Connection to Differential Entropy

We show that there exists scenarios where *ID-Entropy* is related to the traditional differential entropy (Cover 1999).

**Definition 7. (Differential Entropy)** The differential entropy  $H(X)$  of a RV  $X$  with an underlying probability density function of  $P(X)$  is defined as  $H(X) = \mathbb{E}_X [\log(P(X))]$ .

The result, shown in Proposition 1, finds that only when one re-imagines the notion of *probability density*, the two metrics can be shown to be of similar form.

**Proposition 1.** Let us consider distributions  $P(X)$  such that for all points  $X_0 \in \mathbb{R}^d$  which lie in the support of  $P(X)$ , and

for  $\epsilon \rightarrow 0^+$  the following holds.

$$\int_{\|X - X_0\| \leq \epsilon} P(X) dX = F(\epsilon) \quad (4)$$

where  $F(\epsilon)$  is some function of  $\epsilon$ . The above imposes a constraint where the distribution is evenly distributed across the support of  $P$ . Next, we consider a slightly modified notion of probability density  $P_{supp}(X)$ , where  $P_{supp}(X_0) = \lim_{vol(V) \rightarrow 0} \int_{X \in V} P(X) dX / vol(V)$ , where  $V$  is the  $k$ -dimensional hypercube around  $X_0$  and  $k$  is the ID at  $X_0$ . Thus  $P_{supp}(X)$  measures how the probability is distributed w.r.t the local distribution support. Then, as  $\epsilon \rightarrow 0^+$ ,

$$\mathbb{E}_X [\log(P_{supp}(X))] = \log(\epsilon) \mathfrak{h}(X) + \log \frac{K}{F(\epsilon)}, \quad (5)$$

for some fixed scalar  $K$ .

**Remark 4.** Note that there are two separate conditions for the result in Proposition 1 to hold, (a) the equi-probability condition indirectly enforces that among all distributions which are only non-zero valued in the support of  $P(X)$ ,  $P(X)$  is of the largest differential entropy, and (b) the re-imagining of probability density with respect to the volume of the support of  $P(X)$ , rather than the hypercube in  $\mathbb{R}^d$ , which is usually the case. For example, if the data manifold in  $\mathbb{R}^d$  is a 2D plane, the probability density is the total probability divided by the total surface area of the plane.

### Perspectives on ID-Entropy

In this section, we provide additional results pertaining to *ID-Entropy* and its variants, to shed more perspectives on the nature of information encoded by the metric.

#### ID-Entropy: New Bottleneck Criteria

We find that *ID-Entropy* can be used to formulate information bottleneck criteria similar to (Tishby and Zaslavsky 2015), for both auto-encoders and classifiers. Let  $T$  represent the hidden layer of a network.  $\mathfrak{h}(T)$  is non-zero and finite, unlike  $I(X; T)$  which is actually infinite when  $T$  is a deterministic function of  $X$ . We outline the two bottleneck criteria in the following definitions, which are later also supported by the theoretical results in the following section.

**Definition 8. (Bottleneck for Auto-Encoders)** Consider the class of auto-encoders which can be represented sequentially as  $X \rightsquigarrow T \rightsquigarrow \tilde{X}$ , where  $T$  is the encoding of  $X$  as some latent feature representation. When  $X = \tilde{X}$ , we have that  $\mathfrak{h}(X) = \mathfrak{h}(T) \leq \mathfrak{h}^\epsilon(T) \leq \dim(T)$ . Thus, we outline the condition for a relaxed information bottleneck criterion as follows.

$$\min \mathfrak{h}^\epsilon(T) \text{ s.t. } \tilde{X} = X. \quad (6)$$

**Definition 9. (Bottleneck for Classifiers)** Consider the class of feedforward supervised architectures which can be represented sequentially as  $X \rightsquigarrow T \rightsquigarrow \tilde{Y}$ , where  $T$  represents any hidden layer within the architecture. Given this, we outline the conditions for a relaxed information bottleneck criterion as follows.

$$\min \mathfrak{h}^\epsilon(T) \text{ s.t. } \tilde{Y} = Y. \quad (7)$$

**Remark 5.** Note that unlike conventional bottleneck measures which work with the MI between the features and inputs, or features and labels, here, the bottleneck principle for ID-Entropy, in both scenarios, concerns with the  $\epsilon$ -ID-Entropy of the features. Specifically, for these bottlenecks, the objective is to minimize  $\epsilon$ -ID-Entropy of the features, under the condition that the network has a perfect fit on the output, for both scenarios.

### ID-Entropy: Relevance to Generalization

In this section, we provide two theoretical results that connect ID-Entropy of the input and the hidden layers to generalization error, for both auto-encoders and classifiers. We first give some necessary definitions, after which we provide the bounds on generalization error for both cases.

**Definition 10. (Label Generating Function (LGF):)** For any given binary classification task, where the input data  $X \in \mathbb{R}^d$  and  $X \sim P$ , and labels  $Y \in \{-1, 1\}$ , we define the ground truth label generating function  $g_{true}$  as the function which generates the true labels on all datapoints  $X \sim P$ .

**Definition 11. (Probabilistically Lipschitz:)** Any function  $f : \mathbb{R}^d \rightarrow \{-1, 1\}$  is said to be probabilistically  $C_p$ -Lipschitz, if it satisfies the following. For any  $X, X' \in \mathbb{R}^d$ ,

$$P(f(X) \neq f(X')) \leq C_p \|X - X'\|. \quad (8)$$

**Theorem 1. (Auto-Encoders)** We consider a scenario with an auto-encoder network. Let  $S = \{X_1, X_2, \dots, X_m\}$ , be the training data points. Also, let  $\widehat{err}_{MSE}(S)$  represent the training mean-squared reconstruction error and  $err_{MSE}(g(f_T(X)), X) = \mathbb{E}_S[\widehat{err}_{MSE}(S)]$  be the generalization error, where  $f_T(X)$  represents the function at any layer  $T$ , and  $g(\cdot)$  represents the reconstruction function. Let  $g_{true}(\cdot)$  be the ground truth reconstruction function. We consider architectures which yield continuous  $f_T$  and  $g$ . Given that  $X \sim P$ , let  $d_{max} = \max_{X \sim P, i} \|X - X_i\|$ . Then, as  $m \rightarrow \infty$ , for Lipschitz continuous  $g$ ,  $g_{true}$  and  $f_T$ , and for some  $C > 0$  we have that

$$err_{MSE}(g(f_T(X)), X) \leq \widehat{err}_{MSE}(S) + \frac{C d_{max}^2 \mathfrak{h}(T)}{\mathfrak{h}(T) + 2}, \quad (9)$$

where  $C$  depends on Lipschitz constants of  $g$ ,  $g_{true}$  and  $f_T$ .

**Corollary 1.1. (Classifiers)** We are given training data points and labels  $S = \{(X_1, y_1), (X_2, y_2), \dots, (X_m, y_m)\}$ . Also, let  $\widehat{err}(S)$  represent the 0-1 loss on the training data and  $err(g(f_T(X)), y)$  represent the generalization error, where  $f_T(X)$  represents the function at any layer  $T$ , and  $g(f_T(X))$  yields the network output. Let  $g_{true}(\cdot)$  be the ground truth label generating function. We consider architectures which yield continuous  $f_T$  and  $g$ . Then, as  $m \rightarrow \infty$ , for probabilistically Lipschitz continuous  $g$  and  $g_{true}$  and Lipschitz continuous  $f_T$ , we have that

$$err(g(f_T(X)), y) \leq \widehat{err}(S) + \frac{C_p d_{max} \mathfrak{h}(T)}{\mathfrak{h}(T) + 1}, \quad (10)$$

where  $C_p$  depends on the probabilistically Lipschitz constants of  $g$  and  $g_{true}$ , and the Lipschitz constant of  $f_T$ .

**Remark 6.** Note that these theoretical results support the bottleneck criteria defined in the previous section. Taken together, Theorem 1 and Corollary 1.1 both indicate that a larger ID-Entropy of the feature layers will likely lead to a larger generalization gap. Note that these results are derived for the asymptotic case of infinite training examples, but we should expect them to hold for large training datasets as well, and to a certain extent in smaller training datasets. In our experiments we indeed find that for moderately sized training datasets in MNIST and CIFAR-10, ID-Entropy of feature layers indeed grow with the generalization gap.

**Remark 7.** Note that these results are not intended to be computable bounds, but rather to provide insights into the overall trends on how ID-Entropy can potentially relate to the generalization gap, particularly the asymptotic behaviour of the gap. Note, as  $m \rightarrow \infty$ , we have  $d_{max} \rightarrow 0$ .

### ID-Entropy: Connections to Causality

The properties of ID-Entropy indicate that it could potentially encode the underlying causal factors under a specific set of assumptions for cause-effect relationships. The following proposition shows this for the case where  $k$  independent, continuous valued causes generate the data.

**Proposition 2.** Consider any arbitrary causal diagram (directed acyclic graph) which has an observable RV  $X \in \mathbb{R}^d$ , and a set of hidden variables (potentially causes) in  $C = \{C_1, C_2, \dots, C_n\}$  where  $C_i \in \mathbb{R} \forall i$ . Furthermore, assume that if  $(X_1, X_2, \dots, X_p)$  are the parent nodes of  $Y$ , then we have that  $(X_1, X_2, \dots, X_p) \rightsquigarrow Y$ . Given this, we have

$$\mathfrak{h}(X) \leq \min k \text{ s.t. } (C_{a(1)}, C_{a(2)}, \dots, C_{a(k)}) \rightsquigarrow X, \quad (11)$$

where  $1 \leq a(i) \leq n$ . Thus, the ID-Entropy of  $X$  is less than the minimum number of variables in  $C$  such that  $X$  is  $f$ -related to them.

**Remark 8.** The result in proposition 2 holds for causal models where, if  $X_1, \dots, X_k$  are the sole causes of  $Y$ , then  $(X_1, \dots, X_k) \rightsquigarrow Y$ , i.e.,  $Y$  is some continuous function of  $X_1, \dots, X_k$ . In this setting, we find that the ID-Entropy of observable variables are internally related to the minimal number of causes that functionally relate to the observed variable. As ID-Entropy is invariant to the shape and statistics of the distribution (Remark 3), in the context of proposition 2 it implies that ID-Entropy can capture structural information pertaining to the hidden causes, rather than purely statistical information.

### Related Work

With regard to our proposed variants of ID-entropy measure, their various entropy-like properties and the relevance to generalization, we did not find much directly related work. However, there are various bodies of work in literature which have individually studied the various aspects of our proposed measure. First, we note that as one of our objectives is to find a suitable measure of entropy that is invariant to homeomorphisms, this rules out the well-known variants of divergence measures in the literature. They include, KL-divergence, Bregman-divergence, and the f-divergences

(Basseville 2013). We note that these measures are not topological invariants, i.e., they are not invariant to the set of all homeomorphisms. Similarly, other topological properties of a distribution such as the Hausdorff dimension and the fractal dimensions are not topological invariants (Fernández-Martínez 2016). We also note that intrinsic dimensionality has been studied in previous works (Birdal et al. 2021; Nakada and Imaizumi 2020), not from the perspective of an entropy measure, however. Rather, the use of homeomorphisms to  $\mathbb{R}^n$  in our work in order to arrive at the local dimensionality is more aligned with the concept of topological dimension (Coornaert 2015). However, the topological dimension is a global metric whereas our proposed measures are expected values of their local counterparts. There has been a body of work which connects dimensionality and entropy (Wu and Verdú 2010; Zmeskal, Dzik, and Vesely 2013), however these approaches are different to our formulation of entropy, and can be characterized through the fractal dimension of the distribution, which are not topological invariants. This includes the information dimension (Rényi 1959), which was found to relate to ID under a similar expression in (Romano et al. 2016), only under the finite entropy constraint, thus not being a true topological invariant. Apart from the aforementioned, there are works that propose alternative entropy definitions motivated by the structure of data (Giraldo, Rao, and Principe 2014; Yu and Principe 2019; Yu, Yu, and Príncipe 2021), but are not ID-based and are not topological invariants. Note that here we define ID in a specific manner, encompassing both discrete and continuous variables (1), that extends to distributions that aren't locally Euclidean (e.g., intersecting manifolds). Lastly, we additionally discuss its relevance to a broad range of concepts such as causality, information bottlenecks and generalization error. With regard to generalization error, recent studies in (Ansuini et al. 2019; Birdal et al. 2021; Nakada and Imaizumi 2020; Latorre et al. 2021) also note that low data ID yields a lower generalization error. However, (i) these studies mainly consider the global ID and (ii) a theoretical study of how the expected local ID asymptotically relates to generalization error for both classifier and auto-encoder architectures is lacking. Other work also has studied the relevance of ID (global) to generalization error such as in (Birdal et al. 2021), however, the ID there is estimated for the network function, rather than the input and hidden layers.

## Experiments and Discussions

Here we showcase two experiments, where we compute the ID-Entropy of the feature layers of *classifier* and *auto-encoder* architectures on MNIST and CIFAR-10, and contrast it with generalization performance. For all experiments we estimated the ID-entropy using Algorithm 1, with  $n = 2000$  and  $k = 100$ . We used FisherS intrinsic dimensionality estimator in (Bac and Zinovyev 2020) as the global ID estimator  $f_{ID}$ , as we found it to be a robust choice. Following the results of Theorem 1 and Corollary 1.1, higher ID should point to worse generalization performance and vice-versa. Experiment details and further results are given in the Appendix, available in (Ghosh and Motani 2023), and code is available at: <https://github.com/kentridgeai/ID-Entropy>.

## ID-Entropy Computation: Classifiers

**Generalization Gap and ID-Entropy:** With a fixed 4-layer CNN architecture for MNIST and a ResNet-44 for CIFAR-10, we repeat the training routine with different choices of the training data size and random network initializations. This leads to significant changes in test accuracy across varying training data sizes, with the training accuracy of the resulting networks being  $\approx 100\%$  in each case. Subsequently, we report the average ID-Entropy for each choice of training dataset size across 5 separate runs and their standard deviation, and plot them alongside the average test accuracy for the corresponding training dataset size in Figure 1 (a). Furthermore, we also report the corresponding estimates of global ID using the FisherS estimator in each case. We see that ID-Entropy of the feature layer clearly reduces as the test accuracy increases, in agreement with the implications of Corollary 1.1, whereas we find that the global ID also yields similar trends for CIFAR-10, but not for MNIST.

**Label-Noise Addition:** For both MNIST (4-layer CNN) and CIFAR-10 (ResNet-44), we add label noise by randomly changing the label of a training data point with a certain probability. It is well known that adding label noise makes the problem harder (Karimi et al. 2020), significantly impacting test performance, increasing the generalization gap. We summarize the results in Figure 1 (c), where we plot the trajectories of average training error, average test error, and average ID-Entropy of the CNN feature layer (second to last), for various choices of the label noise probability  $p = \{0, 0.1, 0.2, 0.3, 0.4, 0.5\}$ . All reported measures are computed by averaging over five runs. The results show that with greater label noise, the test accuracy declines and the ID-Entropy of the features increases. Furthermore, in each case, we also find that the training time trend of ID-Entropy follows an opposite trend compared to the network's test performance in most cases. Lastly, we see that for the no-noise case  $p = 0$ , ID-Entropy shows a sharp decline, after which it either stabilizes or slowly increases with training.

**Information Bottleneck Discussion:** With respect to our proposed bottleneck principle, the ID-Entropy of the last hidden layer  $T$ ,  $\mathcal{H}(T)$ , is also the mutual information  $\mathcal{I}(X; T)$ , due to P7. As our proposed bottleneck principle dictates that ideally  $\mathcal{H}(T)$  must be minimized, we verify this in our experiments. As the results in Figure 1 (c) and 1 (f) show,  $\mathcal{H}(T)$  has a consistent trend during training. Furthermore, in contrast to  $I(X; T)$  in (Tishby and Zaslavsky 2015), which predicts two distinct phases in the respective order: fitting ( $I(X; T)$  and  $I(Y; T)$  increase) and compression ( $I(X; T)$  decreases), we find that the phases are *reversed* for  $\mathcal{I}(X; T)$ . From the perspective of ID-Entropy, as seen in the label noise plots, the network can undergo up to two distinct phases: (i) *simultaneous compression and fitting* (ID-Entropy decreases) and (ii) *decompression* (ID-Entropy increases), and only in the latter phase the network potentially can overfit the dataset. This agrees with observations in (Stephenson et al. 2021) which finds that networks learn generalizable features in early epochs and only memorize later. When  $p = 0$  (no label noise), we find that the network simply minimizes  $\mathcal{H}(T)$  as training progresses, and thus the decompression phase isn't observed.

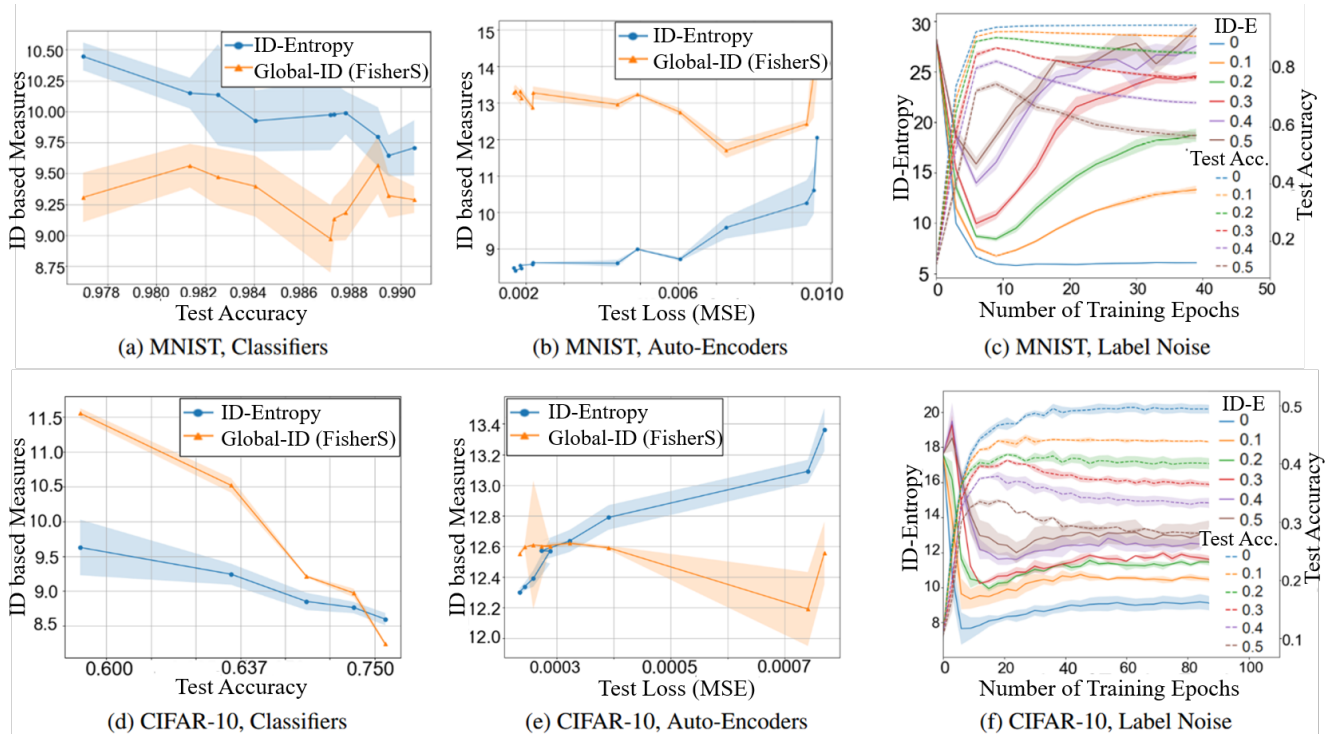


Figure 1: Figures showing the various trends of ID-Entropy and the model performance, for models trained on MNIST and CIFAR-10. (a) and (d) show the average ID-Entropy and global ID (scaled) of the last hidden layer for a 4-layer CNN trained on MNIST and a Resnet-44 trained on CIFAR-10, for various choices of training data size, each leading to a different average test accuracy, (b) and (e) show the dependence between test loss (mean-squared error) and the ID-Entropy and global ID of the encoded feature layer, for a 4 layer Convolutional auto-encoder trained on MNIST and CIFAR-10, and (c) and (f) show the progression of average ID-Entropy of the last hidden layer for a 4-layer CNN trained on MNIST and a ResNet-44 trained on CIFAR-10 for different levels of label-noise in the training samples (from 0 to 0.5).

### ID-Entropy Computation: Auto-Encoders

Like before, we plot the ID-Entropy of a 4-layer Convolutional auto-encoder (CAE) trained on the MNIST and CIFAR-10 datasets. In Figure 1 (b) and (e), we plot the ID-Entropy and global ID of the feature layer, which is also the input layer for the decoder part of the CAE, and plot the average test loss (mean-squared error) during training. All reported measures are computed by averaging over five runs. Here we find that ID-Entropy correlates to the observed test loss in both datasets, which agrees with the implications of the Corollary 1. Furthermore, in both cases, we find that the global ID does not adhere to any such clear trends. This indicates that a local ID based measure such as ID-Entropy more accurately tracks generalization error in the case of auto-encoders, compared to global ID.

### Reflections

Through its properties, we find that the proposed measure of *ID-Entropy* encodes more *structural* information regarding the distribution, rather than statistical. Furthermore, the main differentiator between *ID-Entropy* and conventional differential entropy is that *ID-Entropy* preserves the intrinsic dimensionality information of the data manifold, whereas,

for differential entropy and other cardinality-driven variants, this information can be lost, due to Cantor’s argument.

*ID-Entropy* also seems to relate to causal factors that may be responsible for generating the data (as seen in Proposition 2), when we assume that cause-effect relationships are only via continuous functions, which yield  $f$ -relations between causes and effects. As such,  $f$ -relatedness also represents a fundamental property between two variables, and is independent of the probability distributions of the causal variables, i.e., changing the distribution of the cause  $X$  through interventions still preserves the property that  $X \rightsquigarrow Y$ . Also, as in the definition of *ID-Entropy*, note that only when  $X \leftrightarrow Y$ ,  $X$  and  $Y$  can be said to be homeomorphic.

The results relating ID-Entropy to generalization showcase the relevance of ID-Entropy to generalization error both in the context of classifier and auto-encoder architectures. Empirically we see this to be the case as well. An interesting empirical finding was that for datasets that are primarily *harder* to generalize, such as the ones with label-noise tested here, the ID-Entropy of the hidden layers shows a significant increase. Thus, to a certain extent, ID-Entropy showcases the level of memorization in the features learnt by the network, as noisy labels forces a neural network to memorize the labels, leading to poor test performance (Zhang et al. 2021).

## Acknowledgements

This research is supported by A\*STAR, CISCO Systems (USA) Pte. Ltd and the National University of Singapore under its Cisco-NUS Accelerated Digital Economy Corporate Laboratory (Award I21001E0002). Additionally, we would like to thank the members of the Kent-Ridge AI research group at the National University of Singapore for helpful feedback and interesting discussions on this work.

## References

- Ansuini, A.; Laio, A.; Macke, J. H.; and Zoccolan, D. 2019. Intrinsic dimension of data representations in deep neural networks. *Advances in Neural Information Processing Systems*, 32.
- Bac, J.; and Zinovyev, A. 2020. Local intrinsic dimensionality estimators based on concentration of measure. In *2020 International Joint Conference on Neural Networks (IJCNN)*, 1–8.
- Basseville, M. 2013. Divergence measures for statistical data processing—An annotated bibliography. *Signal Processing*, 93(4): 621–633.
- Birdal, T.; Lou, A.; Guibas, L. J.; and Simsekli, U. 2021. Intrinsic dimension, persistent homology and generalization in neural networks. *Advances in Neural Information Processing Systems*, 34.
- Cantor, G. 1874. On a Property of the Class of all Real Algebraic Numbers. In *Crelle's Journal for Mathematics*, volume 77, 258–262.
- Coornaert, M. 2015. Topological Dimension. In *Topological Dimension and Dynamical Systems*, 3–25. Springer.
- Cover, T. M. 1999. *Elements of information theory*. John Wiley & Sons.
- Fernández-Martínez, M. 2016. A survey on fractal dimension for fractal structures. *Applied Mathematics and Nonlinear Sciences*, 1(2): 437–472.
- Ghosh, R.; and Motani, M. 2023. Local Intrinsic Dimensional Entropy. arXiv:2304.02223.
- Giraldo, L. G. S.; Rao, M.; and Principe, J. C. 2014. Measures of entropy from data using infinitely divisible kernels. *IEEE Transactions on Information Theory*, 61(1): 535–548.
- Goldfeld, Z.; and Polyanskiy, Y. 2020. The information bottleneck problem and its applications in machine learning. *IEEE Journal on Selected Areas in Information Theory*, 1(1): 19–38.
- Karimi, D.; Dou, H.; Warfield, S. K.; and Gholipour, A. 2020. Deep learning with noisy labels: Exploring techniques and remedies in medical image analysis. *Medical image analysis*, 65: 101759–101759.
- Latorre, F.; Dadi, L. T.; Rolland, P.; and Cevher, V. 2021. The Effect of the Intrinsic Dimension on the Generalization of Quadratic Classifiers. *Advances in Neural Information Processing Systems*, 34.
- Lee, J. 2010. *Introduction to topological manifolds*, volume 202. Springer Science & Business Media.
- Nakada, R.; and Imaizumi, M. 2020. Adaptive Approximation and Generalization of Deep Neural Network with Intrinsic Dimensionality. *J. Mach. Learn. Res.*, 21: 174–1.
- Rényi, A. 1959. On the dimension and entropy of probability distributions. *Acta Mathematica Academiae Scientiarum Hungarica*, 10(1): 193–215.
- Rényi, A. 1960. On Measures of Information and Entropy. *Proceedings of the 4th Berkeley Symposium on Mathematics, Statistics and Probability*, 20: 547–561.
- Romano, S.; Chelly, O.; Nguyen, V.; Bailey, J.; and Houle, M. E. 2016. Measuring dependency via intrinsic dimensionality. In *2016 23rd international conference on pattern recognition (ICPR)*, 1207–1212. IEEE.
- Saxe, A. M.; Bansal, Y.; Dapello, J.; Advani, M.; Kolchinsky, A.; Tracey, B. D.; and Cox, D. D. 2018. On the Information Bottleneck Theory of Deep Learning. In *International Conference on Learning Representations*.
- Shannon, C. E. 1948. A Mathematical Theory of Communication. *Bell System Technical Journal*, 27(3): 379–423.
- Stephenson, C.; Padhy, S.; Ganesh, A.; Hui, Y.; Tang, H.; and Chung, S. 2021. On the geometry of generalization and memorization in deep neural networks. *arXiv preprint arXiv:2105.14602*.
- Tishby, N.; and Zaslavsky, N. 2015. Deep learning and the information bottleneck principle. In *2015 IEEE Information Theory Workshop (ITW)*, 1–5.
- Wikipedia. 2022. Vector field — Wikipedia, The Free Encyclopedia. <http://en.wikipedia.org/w/index.php?title=Vector%20field&oldid=1119577269>. [Online; accessed 02-December-2022].
- Wu, Y.; and Verdú, S. 2010. Rényi information dimension: Fundamental limits of almost lossless analog compression. *IEEE Transactions on Information Theory*, 56(8): 3721–3748.
- Yu, S.; and Principe, J. C. 2019. Understanding autoencoders with information theoretic concepts. *Neural Networks*, 117: 104–123.
- Yu, X.; Yu, S.; and Principe, J. C. 2021. Deep deterministic information bottleneck with matrix-based entropy functional. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 3160–3164. IEEE.
- Zhang, C.; Bengio, S.; Hardt, M.; Recht, B.; and Vinyals, O. 2021. Understanding Deep Learning (Still) Requires Rethinking Generalization. *Commun. ACM*, 64(3): 107–115.
- Zmeskal, O.; Dzik, P.; and Vesely, M. 2013. Entropy of fractal systems. *Computers & Mathematics with Applications*, 66(2): 135–146.