

---

# Inductive Mutual Information Estimation: A Convex Maximum-Entropy Copula Approach

---

Yves-Laurent Kom Samo

KXY Technologies, Inc.

✉ y1@kxy.ai 📧 @Dr\_YLKS 🐦 @Dr\_YLKS

San Jose, California, USA

## Abstract

We propose a novel estimator of the mutual information between two ordinal vectors  $\mathbf{x}$  and  $\mathbf{y}$ . Our approach is inductive (as opposed to deductive) in that it depends on the data generating distribution solely through some nonparametric properties revealing associations in the data, and does not require having enough data to fully characterize the true joint distributions  $P_{\mathbf{x}, \mathbf{y}}$ . Specifically, our approach consists of (i) noting that  $I(\mathbf{y}; \mathbf{x}) = I(\mathbf{u}_y; \mathbf{u}_x)$  where  $\mathbf{u}_y$  and  $\mathbf{u}_x$  are the *copula-uniform dual representations* of  $\mathbf{y}$  and  $\mathbf{x}$  (i.e. their images under the probability integral transform), and (ii) estimating the copula entropies  $h(\mathbf{u}_y)$ ,  $h(\mathbf{u}_x)$  and  $h(\mathbf{u}_y, \mathbf{u}_x)$  by solving a maximum-entropy problem over the space of copula densities under a constraint of the type  $\alpha_m = E[\phi_m(\mathbf{u}_y, \mathbf{u}_x)]$ . We prove that, so long as the constraint is feasible, this problem admits a unique solution, it is in the exponential family, and it can be learned by solving a convex optimization problem. The resulting estimator, which we denote MIND, is marginal-invariant, always non-negative, unbounded for any sample size  $n$ , consistent, has MSE rate  $O(1/n)$ , and is more data-efficient than competing approaches.

## 1 Introduction

Mutual information plays a key role in statistical learning. It is directly related to the highest  $R^2$ , the highest true log-likelihood per observation, the lowest

root mean square error and the highest classification accuracy that can be achieved by using explanatory variables  $\mathbf{x}$  to predict categorical or continuous output(s)  $\mathbf{y}$ . It also plays an important role in representation learning (Brown et al. (1992); Bell and Sejnowski (1995); Tishby et al. (2000); Tishby and Zaslavsky (2015); Chen et al. (2016); Higgins et al. (2018)) and reinforcement learning (Pathak et al. (2017); Oord et al. (2018)).

Virtually every mutual information estimator in the literature implicitly assumes that we have a number  $n$  of i.i.d. samples  $(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_n, \mathbf{y}_n)$  that is large enough to characterize the underlying distribution  $P_{\mathbf{x}, \mathbf{y}}$ . We will refer to this scenario as the *deductive approach*. Examples include quantizing (Paninski (2003)) or hashing (Noshad et al. (2019))  $\mathbf{x}$  and  $\mathbf{y}$  and computing the mutual information between the resulting discrete distributions using sample frequencies. Other approaches approximate the pdfs using kernel density estimators (Moon et al. (1995); Kwak and Choi (2002)), using local geometric properties based on  $k$  nearest neighbors (Kraskov et al. (2004); Gao et al. (2015)) and using Edgeworth approximation (Hulle (2005)). Another perspective has been to learn lower bounds based on variational characterizations of the mutual information (Nguyen et al. (2010); Belghazi et al. (2018)) using M-estimators.

The *deductive approach* is fraught with limitations. When relying on discrete approximations, the mutual information can never be greater than the mutual information between fully dependent uniform distributions, namely  $\log n$  in the case of quantization (Paninski (2003)) and  $\log F$  in the case of hashing (Noshad et al. (2019)), where  $F < n$  is the number of distinct hashes and  $n$  the sample size. McAllester and Stratos (2020) extended this result empirically to various continuous mutual information estimators relying on the deductive approach, including variational estimators (Nguyen et al. (2010); Belghazi et al. (2018)). Their work touches on the core issue: if we require properly

---

Proceedings of the 24<sup>th</sup> International Conference on Artificial Intelligence and Statistics (AISTATS) 2021, San Diego, California, USA. PMLR: Volume 130. Copyright 2021 by the author(s).

characterizing the joint pdf nonparametrically from  $n$  i.i.d. samples in order to estimate a mutual information, then  $n$  ought to be large, otherwise we will not see enough tail events, and we will fail to account for tail dependency. Unfortunately, the alternative proposed by [McAllester and Stratos \(2020\)](#), namely approximating the density  $p(\mathbf{x})$  and  $p(\mathbf{x}|\mathbf{y})$  parametrically using deep neural networks is data inefficient and tends to overshoot on small sample sizes. Along the same line, [Song and Ermon \(2020\)](#) and [Poole et al. \(2019\)](#) reported that the variance of variational estimators MINE ([Belghazi et al. \(2018\)](#)) and NWJ ([Nguyen et al. \(2010\)](#)) could be very large, and even grow exponentially with the true mutual information ([Song and Ermon \(2020\)](#)), due to the need for a large number  $n$  of samples to accurately estimate an expectation of the form  $E[e^{T(\mathbf{x}, \mathbf{y})}]$  under  $P_{\mathbf{x}} \otimes P_{\mathbf{y}}$ .

Another major limitation of the deductive approach is the unnecessary need to accurately model the marginals of  $P_{\mathbf{x}, \mathbf{y}}$ , directly or implicitly, as a prerequisite for estimating the mutual information, which could be data and compute intensive, even though the mutual information does not depend on marginal distributions.

The *inductive approach* we introduce in this paper is structured in two stages. First, we measure a few nonparametric properties of the data generating distribution that serve as marginal-invariant proxies revealing associations between coordinates of  $\mathbf{x}$  and/or  $\mathbf{y}$ . Then we estimate the mutual information in the spirit of the maximum-entropy principle ([Jaynes \(1957a,b\)](#)), by being consistent with all observed properties, while remaining as uninformative as possible about any property we haven't observed. Intuitively, we would expect that the more expressive the properties we measure get, the closer we should get to the true mutual information. Indeed, we propose a family of nonparametric properties that give rise to a consistent estimator of the true mutual information.

The rest of the paper is structured as follows. In Section [2](#) we recall some results relating copulas and mutual information. In Section [3](#) we further motivate our inductive approach and we present core theoretical results pertaining to maximum-entropy inference of copulas. Our theoretical contribution builds on the study of the  $I$ -divergence geometry of probability distributions developed by [Csiszár \(1975\)](#), which generalizes the minimum discrimination information theorem of [Kullback et al. \(1966\)](#). In Section [4](#) we propose a convex program for solving maximum-entropy copula problems under linear constraints, and we discuss practical considerations. Finally, in Section [5](#) we illustrate that our estimator outperforms the state of the art on large and small mutual information problems on

synthetic data, we illustrate that our work can be used to mitigate mode collapse in GANs, and we apply our approach to the estimation of the highest performance achievable in a Kaggle competition.

## 2 Background

We begin by recalling that the *mutual information* between two random vectors  $\mathbf{x}$  and  $\mathbf{y}$  is defined as

$$I(\mathbf{y}; \mathbf{x}) := \int_{\mathcal{X} \times \mathcal{Y}} \log \frac{dP_{\mathbf{x}, \mathbf{y}}}{dP_{\mathbf{x}} \otimes P_{\mathbf{y}}} dP_{\mathbf{x}, \mathbf{y}}$$

where  $P_{\mathbf{x}, \mathbf{y}}$  (resp.  $P_{\mathbf{x}}, P_{\mathbf{y}}$ ) is the (joint) probability measure of  $(\mathbf{x}, \mathbf{y})$  (resp.  $\mathbf{x}, \mathbf{y}$ ), and  $\frac{dP_{\mathbf{x}, \mathbf{y}}}{dP_{\mathbf{x}} \otimes P_{\mathbf{y}}}$  is the Radon-Nikodym derivative of the joint probability measure with respect to the product measure of  $P_{\mathbf{x}}$  and  $P_{\mathbf{y}}$ . Friendlier expressions depending on whether the random vectors have continuous and/or categorical coordinates are provided in Table [2](#) in the appendix.

While the mutual information is the canonical approach for *quantifying* associations between random variables, copulas are the canonical tool for *modeling* associations between random variables. We recall some basic definitions and properties, and link the two notions.

**Definition 2.1.** A **copula distribution** is any probability distribution supported on  $[0, 1]^d$  whose marginals are uniform. A **copula** (resp. **copula density**) is any function that is the cdf (resp. pdf) of a copula distribution.

The following theorem shows that every distribution with pdf is uniquely associated to a copula density that fully captures its dependence structure, independently from marginals.

**Theorem 2.1.** (*Sklar's Theorem*) Any pdf  $f : \mathcal{Z} \subset \mathbb{R}^d \rightarrow \mathbb{R}^+$  whose marginal pdfs are  $f_1(z_1), \dots, f_d(z_d)$  with associated cdfs  $F_1(z_1), \dots, F_d(z_d)$  can be uniquely decomposed as

$$f(z_1, \dots, z_d) = c(F_1(z_1), \dots, F_d(z_d)) \prod_{i=1}^d f_i(z_i),$$

where  $c : [0, 1]^d \rightarrow \mathbb{R}^+$  is a copula density. We refer to  $\mathbf{u}_z := (F_1(z_1), \dots, F_d(z_d))$  as the **copula-uniform dual representation** of  $\mathbf{z} := (z_1, \dots, z_d)$ , and to  $\mathbf{z}$  as a **primal representation** of  $\mathbf{u}_z$ .

Interestingly, the entropy of a copula-uniform dual representation is invariant by continuous 1-to-1 primal feature transformations.

**Proposition 2.1.** *If  $\mathbf{u}_z$  is the copula-uniform dual representation of  $\mathbf{z} := (z_1, \dots, z_d)$  and  $\mathbf{u}_{g(\mathbf{z})}$  is the copula-uniform dual representation of  $g(\mathbf{z}) := (g_1(z_1), \dots, g_d(z_d))$  where the functions  $g_1, \dots, g_d$  are continuous 1-to-1 functions, then*

$$h(\mathbf{u}_z) = h(\mathbf{u}_{g(\mathbf{z})}).$$

See Appendix [B.1](#) for the proof.

The following entropy decomposition is a direct consequence of Sklar’s theorem.

**Proposition 2.2.** *If  $\mathbf{u}_z$  is the copula-uniform dual representation of  $\mathbf{z} := (z_1, \dots, z_d) \in \mathcal{Z} \subset \mathbb{R}^d$  and  $\mathbf{z}$  admits a pdf, then the differential entropy of  $\mathbf{z}$  can be decomposed as*

$$h(\mathbf{z}) = h(\mathbf{u}_z) + \sum_{i=1}^d h(z_i), \quad (1)$$

so long as all marginal entropies exist.  $h(\mathbf{u}_z)$  is the entropy of the associated copula distribution, and we refer to it as the **copula entropy** of  $\mathbf{z}$ .

This entropy decomposition implies that the mutual information between two continuous random vectors is the same as that of their copula-uniform dual representations:

$$\begin{aligned} I(\mathbf{y}; \mathbf{x}) &= h(\mathbf{y}) + h(\mathbf{x}) - h(\mathbf{y}, \mathbf{x}) \\ &= h(\mathbf{u}_y) + h(\mathbf{u}_x) - h(\mathbf{u}_y, \mathbf{u}_x) \\ &= I(\mathbf{u}_y; \mathbf{u}_x). \end{aligned}$$

The identity  $I(\mathbf{y}; \mathbf{x}) = I(\mathbf{u}_y; \mathbf{u}_x)$  extends to all ordinal random vectors by noting that copula-uniform dual representations are well defined for ordinal random vectors, are in a 1-to-1 relationship with their primal representations, and that the mutual information is invariant by 1-to-1 maps. In general, we may use [Table 2](#) to conclude that estimating any mutual information boils down to estimating copula entropies, and possibly one-dimensional primal entropies (when the problem involves categorical and non-ordinal coordinates that we choose not to ordinally encode).

Another important property of the mutual information we will rely on is that it is stable by addition of redundant information.

**Proposition 2.3.** *Let  $\mathbf{y} \in \mathcal{Y}$  and  $\mathbf{x} \in \mathcal{X}$  be two random vectors, and  $f$  a function defined on  $\mathcal{X}$ . Then we have:*

$$I(\mathbf{y}; \mathbf{x}, f(\mathbf{x})) = I(\mathbf{y}; \mathbf{x}) + \underbrace{I(\mathbf{y}; f(\mathbf{x}) | \mathbf{x})}_{=0} = I(\mathbf{y}; \mathbf{x}).$$

Going forward, and without loss of generality, we will focus on the estimation of the copula entropy of

a continuous random vector  $h(\mathbf{u}_z)$ , with the understanding that the mutual information is recovered as  $I(\mathbf{y}; \mathbf{x}) = h(\mathbf{u}_y) + h(\mathbf{u}_x) - h(\mathbf{u}_y, \mathbf{u}_x)$ .

### 3 Inductive Mutual Information Estimation

We consider estimating the copula entropy  $h(\mathbf{u}_z)$  of a continuous random vector  $\mathbf{z} \in \mathcal{Z} \subset \mathbb{R}^d$ , where  $\mathbf{u}_z$  is the copula-uniform dual representation of  $\mathbf{z}$ , which we assume admits a pdf.

#### 3.1 Motivation and Roadmap

The *deductive approach* to learning the copula entropy  $h(\mathbf{u}_z)$  requires assuming that we have gathered enough samples to fully characterize the pdf  $p(\mathbf{u}_z)$ , estimating the pdf, and then estimating the copula entropy as the entropy of the estimated pdf. This is both data and compute inefficient. If we partition  $[0, 1]^d$  into small hypercubes of side length  $\delta$ , then we need to observe at least one sample per hypercube for a small enough  $\delta$  to properly characterize the pdf non-parametrically. This requires  $n \propto \delta^{-d}$  samples. How small  $\delta$  needs to be depends on how quickly the true pdf varies on  $[0, 1]^d$ . Regardless,  $n$  would grow exponentially with the input dimension, and so would the associated compute requirement.

This inefficiency can be alleviated by assuming that the pdf belongs to a specific parametric family, at the expense of model misspecification. If the parametric family has sufficient statistics  $T(\mathbf{u}_1, \dots, \mathbf{u}_n)$  for  $n$  i.i.d. observations, then pdfs in the family should be maximum-entropy among all pdfs with the same statistics. If this is not the case, then the parametric family would be violating Occam’s razor as the learned pdf would be encoding more structure than evidenced by the data, and the estimated entropy would overshoot. We also note that the pdf  $p(\mathbf{u}_z)$  should have uniform marginals, which makes finding an appropriate parametric family even more difficult.

In the absence of any empirical evidence, Occam’s razor suggests that the most appropriate distribution for  $\mathbf{u}_z$  is the uniform distribution on  $[0, 1]^d$  as it is the least informative (or maximum-entropy) of all distributions supported on  $[0, 1]^d$ . Instead of choosing a rigid parametric family and inheriting its sufficient statistics, we could construct more expressive copula densities by first choosing how to reveal departure from the standard uniform distribution from the data, and then finding the least informative copula density among all copula densities satisfying the observed constraints.

The *inductive approach* we propose consists of revealing the dependence structure in  $\mathbf{u}_z$  by estimating an

expectation of the form  $\alpha_m = E_{P_{\mathbf{u}_z}} [\phi_m(\mathbf{u}_z)]$ , for a vector-valued statistics function  $\phi_m : [0, 1]^d \rightarrow \mathbb{R}^q$ , and estimating  $h(\mathbf{u}_z)$  as the highest copula entropy among all copulas satisfying the constraint  $\alpha_m = E_P[\phi_m(\mathbf{u}_z)]$ . When  $\phi_m$  is given and  $E_{P_{\mathbf{u}_z}}[\phi_m(\mathbf{u}_z)]$  is all the data scientist can reliably observe about the structure of the data (e.g. we are only given pairwise Spearman rank correlations), our estimator is the only estimator consistent with Occam’s razor. Crucially, if we may choose  $\phi_m$ , then we may approximate the copula entropy with arbitrary precision. Specifically, we show that, so long as  $(\phi_m)_m$  are universal approximators of continuous functions on  $[0, 1]^d$ , the solution to this maximum-entropy problem is a consistent estimator of the true copula entropy  $h(\mathbf{u}_z)$  (see Theorem 4.1). Equally important is Corollary 3.1 that states that we may perfectly recover the true mutual information using a finite dimensional statistics function that is not expressive enough to fully characterize the true copula distribution  $P_{\mathbf{u}_{x,y}}$ .

We note that  $\alpha_m$  can be efficiently estimated from  $n$  i.i.d. primal samples  $\mathbf{z}_1, \dots, \mathbf{z}_n$  as

$$\hat{\alpha}_{m,n} = \frac{1}{n} \sum_{i=1}^n \phi_m \left( \frac{\text{rg}(\mathbf{z}_i)}{n+1} \right), \quad (2)$$

where  $\text{rg}(\mathbf{z}_i)$  is the vector of coordinatewise ranks of  $\mathbf{z}_i$  among  $\mathbf{z}_1, \dots, \mathbf{z}_n$ . It follows from the weak convergence of the empirical copula process to the true copula that  $\hat{\alpha}_{m,n}$  is a consistent and asymptotically normal estimator of  $\alpha_m$  (Ruschendorf (1976)). Thus, our *inductive approach* to mutual information estimation truly does not require learning marginal distributions. We show that, so long as the constraint  $E_P[\phi_m(\mathbf{u}_z)] = \hat{\alpha}_{m,n}$  is feasible, the associated maximum-entropy problem admits a unique solution, and it is a consistent estimator<sup>1</sup> of  $h(\mathbf{u}_z)$  when  $(\phi_m)_m$  are universal approximators of continuous functions on  $[0, 1]^d$ . Finally, we introduce a convex optimization problem whose minimizer is the maximizer of our maximum-entropy problem.

### 3.2 Maximum-Entropy Copulas

Let  $\mathcal{C}_d$  be the space of all  $d$ -dimensional copula distributions with pdf,  $h(P)$  the differential entropy of the probability distribution  $P$ , and  $\phi_{m,d} : [0, 1]^d \rightarrow \mathbb{R}^{q(d)}$  a vector-valued function whose coordinate functions are not linearly dependent.

We define the following properties of  $\phi_{m,d}$ : (P1) the first coordinate of  $\phi_{m,d}$  is the constant 1, (P2) each coordinate of  $E_P[\phi_{m,d}(\mathbf{u})]$  captures a way in which  $P$  departs from the uniform distribution, (P3) the fam-

ily  $(\phi_{m,d})_m$  is a universal approximator of continuous functions defined on  $[0, 1]^d$  for any  $d$ , and (P4)  $q(l) + q(k) \leq q(l+k)$  for every  $l, k > 0$  and all coordinates of  $\phi_{m,l}$  and  $\phi_{m,k}$  are also included in  $\phi_{m,l+k}$ . Going forward, we will use  $\phi_m$  and  $q$  in-lieu-of  $\phi_{m,d}$  and  $q(d)$  for ease of notation when the input dimension is unambiguous.

We consider the following optimization problem:

$$\begin{cases} \max_{P \in \mathcal{C}_d} h(P) \\ \text{s.t. } E_P[\phi_m(\mathbf{u})] = \alpha_m \end{cases} \quad (\text{MIND})$$

The theorem below, which we prove in Appendix B.2 states that the solution of (MIND) is an exponential family distribution with sufficient statistics  $\phi_m$ , and with base measure the product of  $d$  measures on  $[0, 1]$  that are absolutely continuous with respect to the standard uniform on  $[0, 1]$ .

**Theorem 3.1.** *Let  $\phi_m$  satisfy (P1). If there is any copula distribution  $P$  with finite differential entropy and satisfying  $E_P[\phi_m(\mathbf{u})] = \alpha_m$ , then the maximum-entropy problem (MIND) admits a unique solution, and the maximizer  $P_M$  is the only copula distribution whose density takes the form*

$$p_M(\mathbf{u}; \phi_m, \alpha_m) = e^{\theta^T \phi_m(\mathbf{u})} \prod_{i=1}^d f_i(u_i), \quad (3)$$

and that satisfies the constraint  $E_{P_M}[\phi_m(\mathbf{u})] = \alpha_m$  for some constant  $\theta$ , and  $d$  non-negative univariate functions  $f_i$ ,  $\log f_i \in L_1([0, 1])$ . Moreover, for any copula distribution  $P \in \mathcal{C}_d$  satisfying  $E_P[\phi_m(\mathbf{u})] = \alpha_m$ ,

$$h(P_M) - h(P) = KL(P||P_M). \quad (4)$$

The practical challenge with applying Theorem 3.1 is the need to learn the free functions  $f_1, \dots, f_d$ . These functions ensure that the maximizer is a copula distribution—i.e. has uniform marginals. We now consider relaxing this requirement and solving the maximum-entropy problem over the space  $\mathcal{D}_d \supset \mathcal{C}_d$  of all continuous probability distributions supported on  $[0, 1]^d$ :

$$\begin{cases} \max_{P \in \mathcal{D}_d} h(P) \\ \text{s.t. } E_P[\phi_m(\mathbf{u})] = \alpha_m \end{cases} \quad (\text{A-MIND})$$

As we later show in Theorem 4.1, (P3) guarantees that, despite this relaxation, the solution to (A-MIND) converges to the true copula entropy as  $m$  goes to infinity. For a given  $m$ , to control how close to uniform the maximizer’s marginals are, we use the moment characterization of the standard uniform  $\forall j$ ,  $E(u^j) = 1/(1+j)$ ,

<sup>1</sup>Consistency here is jointly in  $n$  and  $m$ .



and we match the first  $k$  moments of marginals of candidate distributions to those of the standard uniform. To do so, we introduce the property (P5):  $\phi_m$  has the form

$$\phi_m^k(\mathbf{u}) = (1, u_1, \dots, u_1^k, \dots, u_d, \dots, u_d^k, \psi_m(\mathbf{u})),$$

with associated

$$\boldsymbol{\alpha}_m^k = (1, 1/2, \dots, 1/(1+k), \dots, 1/2, \dots, 1/(1+k), \boldsymbol{\beta}_m).$$

The solution of (A-MIND) is provided by the following theorem, which we prove in Appendix B.3

**Theorem 3.2.** *Let  $\phi_m^k$  satisfy (P5). If there is any distribution  $P$  supported on  $[0, 1]^d$ , with finite differential entropy and satisfying  $E_P[\phi_m^k(\mathbf{u})] = \boldsymbol{\alpha}_m^k$ , then the maximum-entropy problem (A-MIND) admits a unique solution of the form*

$$h_{AM}(\mathbf{u}; \phi_m^k, \boldsymbol{\beta}_m) = -\boldsymbol{\theta}^T \boldsymbol{\alpha}_m^k, \quad (5)$$

and the maximizer  $P_{AM}$  is the only distribution supported on  $[0, 1]^d$ , whose pdf takes the form

$$p_{AM}(\mathbf{u}; \phi_m^k, \boldsymbol{\beta}_m) = e^{\boldsymbol{\theta}^T \phi_m^k(\mathbf{u})} \quad (6)$$

for some constant  $\boldsymbol{\theta}$ , and that satisfies the constraint  $E_{P_{AM}}[\phi_m^k(\mathbf{u})] = \boldsymbol{\alpha}_m^k$ . Moreover, for any distribution  $P \in \mathcal{D}_d$  satisfying  $E_P[\phi_m^k(\mathbf{u})] = \boldsymbol{\alpha}_m^k$ ,

$$h(P_{AM}) - h(P) = KL(P||P_{AM}). \quad (7)$$

**Corollary 3.1.** *Let  $\mathbf{x} \in \mathcal{X}$  and  $\mathbf{y} \in \mathcal{Y}$  be two continuous random variables with mutual information  $I(\mathbf{y}; \mathbf{x})$  and true individual and joint copula distributions  $P_{\mathbf{u}_x}$ ,  $P_{\mathbf{u}_y}$ , and  $P_{\mathbf{u}_{x,y}}$ . If  $P_{AM}(\mathbf{u}_x)$ ,  $P_{AM}(\mathbf{u}_y)$  and  $P_{AM}(\mathbf{u}_{x,y})$  are the solutions to three (A-MIND) problems whose constraints are satisfied by the true copula distributions, and*

$$I_{AM}(\mathbf{y}; \mathbf{x}) := h(P_{AM}(\mathbf{u}_x)) + h(P_{AM}(\mathbf{u}_y)) - h(P_{AM}(\mathbf{u}_{x,y}))$$

is the associated mutual information estimator, then

$$I_{AM}(\mathbf{y}; \mathbf{x}) - I(\mathbf{y}; \mathbf{x}) = KL[P_{\mathbf{u}_x}||P_{AM}(\mathbf{u}_x)] + KL[P_{\mathbf{u}_y}||P_{AM}(\mathbf{u}_y)] - KL[P_{\mathbf{u}_{x,y}}||P_{AM}(\mathbf{u}_{x,y})]. \quad (8)$$

Equation (8) in Corollary 3.1 lays out the theoretical ground for favoring our *inductive approach* over the traditional *deductive approach*. Indeed, it shows that it is not necessary to accurately learn the true data generating distributions  $P_{\mathbf{x}}$ ,  $P_{\mathbf{y}}$ , and  $P_{\mathbf{x}, \mathbf{y}}$  or their copulas in order to accurately learn the mutual information  $I(\mathbf{y}; \mathbf{x})$ . The error made by (A-MIND) in estimating the joint copula entropy  $h(\mathbf{u}_y, \mathbf{u}_x)$  can offset the errors made estimating the individual copula entropies  $h(\mathbf{u}_y)$  and  $h(\mathbf{u}_x)$ , so that we may perfectly estimate the mutual information with a finite dimensional statistics function  $\phi_m$ , without accurately learning the copula distributions.

### 3.3 Iterative MIND

When either (MIND) or (A-MIND) are sparse in the sense that each coordinate of  $\phi_m^k$  is a function of some but not all coordinates of  $\mathbf{u}$ , they can be broken down into smaller, cheaper, more robust and cacheable problems.

**Theorem 3.3.** *Let us assume that  $\psi_m(\mathbf{u})$  takes the form  $\psi_m(\mathbf{u}) = (\eta_1(w, \mathbf{v}_1), \dots, \eta_q(w, \mathbf{v}_q))$  where  $w$  is a coordinate of  $\mathbf{u}$  and the vectors  $\mathbf{v}_1, \dots, \mathbf{v}_q$  are made of coordinates of  $\mathbf{u}$  but share no common coordinate and do not include  $w$ . If the constraint  $E_P[\psi_m(\mathbf{u})] = \boldsymbol{\beta}_m := (\boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_q)$  is feasible, then we have*

$$p_M(\mathbf{u}; \phi_m^k, \boldsymbol{\alpha}_m^k) = \prod_{i=1}^q p_M(w, \mathbf{v}_i; \eta_i, \boldsymbol{\beta}_i), \quad (9)$$

$$h_M(\mathbf{u}; \phi_m^k, \boldsymbol{\alpha}_m^k) = \sum_{i=1}^q h_M(w, \mathbf{v}_i; \eta_i, \boldsymbol{\beta}_i). \quad (10)$$

The proof is provided in Appendix B.4. Theorem 3.3 is useful in single-output problems when blocks of explanatory variables are known or assumed to be independent conditional on the output.

More generally, when each coordinate of  $\phi_m^k$  depends on some but not all coordinates of  $\mathbf{u}$ , the problem (A-MIND) can be broken down into smaller problems that can be solved iteratively.

**Theorem 3.4.** *Let  $\mathbf{u} = (\mathbf{v}_1, \dots, \mathbf{v}_q)$  and let us assume that  $\psi_m(\mathbf{u})$  takes the form  $\psi_m(\mathbf{u}) = (\eta_1(\mathbf{v}_1), \dots, \eta_q(\mathbf{v}_q), \gamma(\mathbf{u}))$  where all  $\eta_i$  and  $\gamma$  satisfy (P1). If the constraint  $E_P[\psi_m(\mathbf{u})] = \boldsymbol{\beta}_m := (\boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_q, \boldsymbol{\beta})$  is feasible, then the maximizer of (A-MIND) has density of the form*

$$p_{AM}(\mathbf{u}; \phi_m^k, \boldsymbol{\beta}_m) = e^{\boldsymbol{\theta}^T \phi_m^k(\mathbf{u})} \prod_{i=1}^q p_{AM}(\mathbf{v}_i; \eta_i, \boldsymbol{\beta}_i) \quad (11)$$

where  $\boldsymbol{\theta}$  is the only constant such that  $E_{P_{AM}}[\phi_m^k(\mathbf{u})] = \boldsymbol{\alpha}_m^k$ .

Moreover, the solution reads

$$h_{AM}(\mathbf{u}; \phi_m^k, \boldsymbol{\beta}_m) = -\boldsymbol{\theta}^T \boldsymbol{\alpha}_m^k + \sum_{i=1}^q h_{AM}(\mathbf{v}_i; \eta_i, \boldsymbol{\beta}_i) - \sum_{i=1}^q KL[p_{AM}(\mathbf{v}_i; \phi_m^k, \boldsymbol{\beta}_m) || p_{AM}(\mathbf{v}_i; \eta_i, \boldsymbol{\beta}_i)],$$

where  $p_{AM}(\mathbf{v}_i; \phi_m^k, \boldsymbol{\beta}_m)$  is the marginal of the maximizer of the full (A-MIND) problem and  $p_{AM}(\mathbf{v}_i; \eta_i, \boldsymbol{\beta}_i)$  the maximizer of the (A-MIND) problem pertaining to  $\mathbf{v}_i$ . Furthermore,

$$-\boldsymbol{\theta}^T \boldsymbol{\alpha}_m^k \leq \sum_{i=1}^q KL[p_{AM}(\mathbf{v}_i; \phi_m^k, \boldsymbol{\beta}_m) || p_{AM}(\mathbf{v}_i; \eta_i, \boldsymbol{\beta}_i)],$$

and the equality holds if and only if  $\theta = 0$ , condition satisfied if  $\forall \mathbf{u}$ ,  $\gamma(\mathbf{u}) = 1$ .

The proof is provided in Appendix B.5. Essentially, to solve a full (A-MIND) problem, we may partition  $\mathbf{u}$  into  $q$  blocks, solve the  $q$  (A-MIND) problems in parallel using within-block constraints, determine whether  $\prod_{i=1}^q p_{\text{AM}}(\mathbf{v}_i; \eta_i, \beta_i)$  satisfies the between-blocks constraints to an acceptable tolerance, and if not solve the full (A-MIND) problem using the parameters of  $\prod_{i=1}^q p_{\text{AM}}(\mathbf{v}_i; \eta_i, \beta_i)$  as initial parameters.

## 4 Estimation

We now turn to estimating the parameters of the solutions of the problems (A-MIND) and (MIND). With  $\phi_m^k$  satisfying (P5), let us consider the problem:

$$\min_{\theta} -\theta^T \alpha_m^k + \int_{[0,1]^d} e^{\theta^T \phi_m^k(\mathbf{u})} d\mathbf{u}. \quad (\text{CVX-MIND})$$

### 4.1 Convex Estimation

**Lemma 4.1.** *The optimization problem (CVX-MIND) is strictly convex.*

**Lemma 4.2.** *The minimizer of problem (CVX-MIND) is the maximizer of problem (A-MIND).*

**Theorem 4.1.** *Let  $\phi_m^k$  satisfy (P3) and (P5). If  $\hat{\beta}_{m,n}$  is a consistent estimator of  $\beta_m := E_{P_{\mathbf{u}_z}}[\psi_m(\mathbf{u})]$ , then for every  $m > 0$*

$$h_{\text{AM}}(\mathbf{u}; \phi_m^k, \hat{\beta}_{m,n}) \xrightarrow[k, n \rightarrow \infty]{} h_M(\mathbf{u}; \psi_m, \beta_m), \quad (\text{A})$$

and for every  $k > 0$

$$h_{\text{AM}}(\mathbf{u}; \phi_m^k, \hat{\beta}_{m,n}) \xrightarrow[m, n \rightarrow \infty]{} h(\mathbf{u}_z). \quad (\text{B})$$

See Appendix B.6 for the proof of Lemma 4.1, Appendix B.7 for the proof of Lemma 4.2, and Appendix B.8 for the proof of Theorem 4.1

### 4.2 Choice of $k$ and $\psi_m$

It follows from Theorem 4.1 that  $k$  controls how close to uniform marginals of  $p_{\text{AM}}$  are, while  $m$  controls how close  $h_{\text{AM}}(\mathbf{u}_z)$  is to  $h(\mathbf{u}_z)$ . When the object of study is to learn the copula itself not just its entropy,  $k$  should be as large as necessary. Note however that the maximum-entropy problem (A-MIND) inherently favors distributions that are as close to uniform as allowed by empirical evidence, so that a large  $k$  might not be needed in practice.

An example family  $(\psi_m)_m$  that satisfies (P3) are polynomials of degree  $m$ , thanks to the Stone-Weierstrass

theorem (Rudin (1973)). With this choice of  $\psi_m$ , our approach can be regarded as a maximum-entropy Taylor expansion of the true log copula density  $\log p(\mathbf{u}_z)$ . Recalling that  $\rho_{ij} = 12E[u_i u_j] - 3$  is the population version of the Spearman rank correlation between associated primal variables (Nelsen (2007)), it follows that with  $m$  as small as 2, we can capture all smooth 1-to-1 associations<sup>2</sup> between any two coordinates of  $\mathbf{z}$ .

Monomials with degree  $m > 2$  grow combinatorially in number, but do not provide as much insights per term as  $\psi_2$ . Thus, before considering higher degree polynomials, we suggest leveraging Proposition 2.3 in combination with  $\psi_2$  to incorporate specific types of smooth but non-1-to-1 associations. For instance, using  $f_{\mu}(\mathbf{x}) := |\mathbf{x} - \mu|$ , where the absolute value is coordinatewise, and solving the problem (MIND) with  $m = 2$  to estimate  $I(\mathbf{y}; \mathbf{x})$  through  $I(\mathbf{y}; \mathbf{x}, f_{\mu}(\mathbf{x}))$ , allows us to reveal any possible associations of the type ‘a coordinate of  $\mathbf{y}$  tend(s) to be monotonically related to the departure of coordinates of  $\mathbf{x}$  from some baseline values’, as well as all smooth 1-to-1 associations between coordinates of  $\mathbf{x}$  and  $\mathbf{y}$ . Good examples for  $\mu$  are the sample median or mean of  $\mathbf{x}$ . To capture departures from a standard range of values rather than a single one,  $f_{\mu}$  can be passed through an  $\epsilon$ -insensitive loss function. When output  $y_j$  is a quasi-periodic function of  $x_i$  (e.g.  $x_i$  is time and output  $y_j$  is seasonal),  $f_{\pi}(x_i) = x_i - \lfloor x_i / \pi_i \rfloor \pi_i$  allows us to capture seasonality-adjusted effects.

### 4.3 Handling Categorical Data

Categorical and non-ordinal variables should be ordinarily encoded as customary, and ordinal data should be treated as continuous variables. The only practical requirements for the validity of this approach are i) to use a ranking function that assigns different ranks to all inputs including ties (e.g. `scipy`’s ‘rankdata’ function with method ‘ordinal’), and ii) to avoid encoding methods that may result in linearly dependent coordinates (e.g. one-hot-encoding on a binary non-ordinal categorical variable). When a suitable ranking function is not available a small random jitter may be added to ordinal variables to remove ties.

This approach is mathematically valid thanks to the quantization characterization of the mutual information (Cover (1999), Definition 8.54). See Appendix A.2 for more details.

<sup>2</sup>A continuous function of one variable is 1-to-1 if and only if it is either decreasing or increasing, both of which are captured by Spearman’s rank correlation.

#### 4.4 Properties

We summarize some key properties of our mutual information estimator  $I_{AM}$ .

**Non-negativity:**  $I_{AM}$  is always non-negative thanks to requirement (P4) and Theorem 3.4

**Unboundedness for every  $n$ :** Because  $I_{AM}$  only depends on estimated expected statistics, not i.i.d. samples themselves, and because said estimated expected statistics may take extreme values for any  $n$ ,  $I_{AM}$  cannot be upper-bounded by a function of  $n$ .

**Marginal-Invariance:** Our entire approach does not depend on marginal distributions. Additionally, our copula entropy estimator is invariant by any increasing univariate feature transformation (it depends on the data solely through ranks), and any smooth 1-to-1 univariate feature transformation for  $k$  large enough (Proposition 2.1 and Theorem 4.1A).

**Low Variance:** Our approach depends on the data generating distribution solely through the expected statistics constraints, which are estimated using Equation (2) with  $O(1/n)$  MSE rate (Ruschendorf (1976)). By the delta method, both the associated natural parameters  $\theta$  and the corresponding copula entropies  $-\theta^T \hat{\alpha}_{m,n}$ , and therefore  $I_{AM}$ , have MSE rate  $O(1/n)$ .

**Consistency:** The fact that  $I_{AM}$  is a consistent estimator of the true mutual information is a direct consequence of the consistency of the rank estimator Equation (2) and of the individual copula entropy estimators as a result (see Theorem 4.1B).

**Low Complexity:** With our choice of  $\phi_m$ , pre-optimization complexity is dominated by the computation of ranks, which scales in  $\mathcal{O}(dn \log n)$ , while optimization can scale in  $\mathcal{O}(d^2)$  using gradient descent, where  $d$  is the number of inputs and output(s). Calculating the integral over  $[0, 1]^d$  is only required while solving (CVX-MIND) to compute the gradient and possibly the Hessian; it is not needed to calculate the optimal entropy itself. Thus, a crude approximation using naive Monte Carlo at every learning step is good enough; the resulting algorithm, which can be regarded as mini-batch stochastic gradient descent on a convex objective (Bottou (2010)), will converge to the right solution even for large  $d$ .

## 5 Applications

We begin by applying our approach to clarifying a common misconception.

### 5.1 The Multivariate Gaussian is Highly Structured

The use of multivariate Gaussian variables is often justified by the fact that they are maximum-entropy (or the least informative of all distributions supported on  $\mathbb{R}^d$ ) under Pearson correlation constraints. Such a choice is equivalent to assuming that marginals are Gaussian and the copula is the Gaussian copula. It might surprise the reader to know that the Gaussian copula is in fact highly structured/informative. For




Figure 1: Two bivariate pdfs with standard normal marginals, and the same Spearman rank correlation  $-0.7$ . The copula is Gaussian on the right, and maximum-entropy under the Spearman rank correlation constraint on the left.

instance, it can be seen in Figure 1 that the Gaussian copula posits that tails are much more tightly coupled than the corresponding<sup>3</sup> maximum-entropy copula.

Additionally, as illustrated in Figure 4 in the Appendix, the bivariate Gaussian pdf is on average about 10% off, and up to 50% off from the pdf with the same marginals and copula the least informative copula with the same Spearman correlation<sup>4</sup>. When applying the maximum-entropy principle in the primal space, constraints should always be broken down into constraints that solely apply to the copula and constraints that solely apply to marginals (if any). If this is not done, marginal entropies will tend to dominate the copula entropy in the entropy decomposition of Proposition 2.2, and the copula will tend to be low entropy. When constraints are so separable, Equation (1) allows us to break down the optimization problem into two, one maximum-entropy problem about the

<sup>3</sup>With the same Spearman correlation structure.

<sup>4</sup>The Spearman rank correlation  $\rho_s$  of a bivariate Gaussian with Pearson correlation  $\rho$  reads  $\rho_s = \frac{6}{\pi} \arcsin\left(\frac{\rho}{2}\right)$  (Kruskal (1958)).

copula, which our approach allows solving, and one about marginals. This is not possible using covariance matrices as maximum-entropy constraints given that Pearson’s correlation is not a functional of the copula: it *does* depend on marginals. As expected, marginal entropies in this case do dominate the copula entropy in the maximum-entropy problem, which explains why univariate Gaussians are indeed high entropy but the Gaussian copula is low entropy/highly structured.




Figure 2: Estimation of the mutual information between two 128-dimensional vectors  $\mathbf{x} = (x_1, \dots, x_d)$  and  $\mathbf{y} = (y_1, \dots, y_d)$  from a draw of 1000 i.i.d. samples.  $(x_i, y_i)$  are i.i.d. Gaussians with mean zero, unit marginal variance, and correlation  $\rho$ . For models using deep neural networks, we run the experiment 10 times and report for each batch number the estimate that is the closest to the ground truth. The ground truth is the black solid line labeled  $I(\mathbf{x}; \mathbf{y})$ .

## 5.2 MIND is Far More Data-Efficient Than Competing Approaches

Next, we illustrate that our approach is far more data-efficient than all alternatives, in both low and high mutual information settings.

We repeat the experiment of [McAllester and Stratos \(2020\)](#), and estimate the mutual information between two  $d$ -dimensional vectors  $\mathbf{x} = (x_1, \dots, x_d)$  and  $\mathbf{y} = (y_1, \dots, y_d)$ , where  $(x_i, y_i)$  are i.i.d. Gaussians with mean zero, unit marginal variance, and correlation  $\rho$ . The true mutual information in this case is  $I(\mathbf{y}; \mathbf{x}) = -\frac{d}{2} \log(1 - \rho^2)$ . We reuse the exact same settings as [McAllester and Stratos \(2020\)](#), except for one simple change. Rather than drawing a fresh mini-batch from the true data generating distribution, which is equivalent to using 384000 i.i.d. samples in total, we generate 1000 i.i.d. samples used by all experiments and from which mini-batches are sampled. We run the experiment in a high ( $\rho = 0.5$ ) and a low ( $\rho = 0.01$ ) mutual information setting. We use the code provided

$I(\mathbf{y}; \mathbf{x})$	MIND	DoE	MINE	NWJ	KSG
18.41	<b>18.36</b>	53.88	6.92	7.22	2.98
0.01	<b>0.08</b>	37.49	5.83	1.67	2.89

Table 1: Estimates of the mutual information between two 128-dimensional vectors from 1000 i.i.d. samples in low and high mutual information settings using various models. The ground truth is the column  $I(\mathbf{y}; \mathbf{x})$  and the closest model to it is in bold.

by the authors of [McAllester and Stratos \(2020\)](#) at <https://github.com/karlstratos/doi> for all models but KSG ([Kraskov et al. \(2004\)](#)) and MIND. For MIND, we use second order polynomials as  $\phi_m$ . As it can be seen in Figure 2 and in Table 1, MIND is the only model able to come anywhere close to the ground truth in high or low mutual information settings. DoE models clearly overshoot in both settings. If we refer to Figure 2 of [McAllester and Stratos \(2020\)](#), we may conclude that DoE models need about  $2000 \times 128$  i.i.d. samples to converge to the ground truth in this experiment, which is 256 times more than what MIND requires. Variational models in the primal space overshoot in low mutual information settings and seem to be upper-bounded by  $O(\log n)$  in high mutual information settings. As for the nonparametric KSG estimator, it struggles with large input dimensions.

## 5.3 Copula Entropy Regularized Generative Adversarial Networks (CER-GANs)

Finally, we illustrate that our approach may be used to prevent mode collapse in GANs.

We recall that GANs ([Goodfellow et al. \(2014\)](#)) are very effective generative models made of two modules: a generator whose state is represented by a function  $G : \mathcal{Z} \subset \mathbb{R}^q \rightarrow \mathcal{X}$ , and a discriminator whose state is represented by a function  $D : \mathcal{X} \rightarrow [0, 1]$ . The aim is to learn a function  $G$  that maps a simple noise or code distribution  $P_{\mathbf{z}}$  supported on  $\mathcal{Z}$  to a true data generating distribution of interest  $P_{\mathbf{x}}$  supported on  $\mathcal{X}$ , so that we may draw samples from the true distribution  $P_{\mathbf{x}}$  (e.g. realistic images) simply as  $G(\mathbf{z})$  with  $\mathbf{z} \sim P_{\mathbf{z}}$ . To do so, a GAN alternates between two steps. The first step consists of learning a discriminating function  $D$  that, as the predictive probability of a binary classifier, is effective at telling draws from the true distribution  $P_{\mathbf{x}}$  apart from draws of the form  $G(\mathbf{z})$  (i.e. that are fake). This is done by maximizing the likelihood-like objective

$$\mathcal{L}_D := E_{P_{\mathbf{x}}} [\log D(\mathbf{x})] + E_{P_{\mathbf{z}}} [\log(1 - D(G(\mathbf{z})))]$$

over functions  $D$  induced by a deep neural network, for a given  $G$ . The second step consists of updating the state  $G$  of the generator so as to trick the discriminator






Figure 3: Illustration of GANs trained on the 25 Gaussians dataset (left) without regularization (middle), and with copula entropy regularization (right). Both models use the same discriminator and generator architectures and were stopped after 500 epochs.

into thinking that fake samples are real. This is done by minimizing the objective

$$\mathcal{L}_G := E_{P_{\mathbf{z}}} [\log (1 - D(G(\mathbf{z})))]$$

over functions  $G$  induced by another deep neural network. When  $P_{\mathbf{x}}$  is multi-modal, as is often the case in real-life applications, if the generator becomes good at sampling from a mode of  $P_{\mathbf{x}}$ , then it will keep generating samples near the same mode, as no term in the objectives  $\mathcal{L}_D$  and  $\mathcal{L}_G$  incentivizes the generator to keep exploring beyond a mode. This pathology of GANs is known as mode collapse (Che et al. (2016)). Belghazi et al. (2018) identified as possible solution regularizing  $\mathcal{L}_G$  with the entropy of  $G(\mathbf{z})$  so as to foster exploration, but the authors considered this solution intractable. Instead they followed Chen et al. (2016) and focused on cases where, in addition to  $\mathbf{z}$ , the generator uses meta-data  $\mathbf{c}$  that implicitly identify modes of the true distribution  $P_{\mathbf{x}}$  (e.g. the digit in the case of MNIST). The authors realized that a low mutual information  $I(G(\mathbf{z}, \mathbf{c}); \mathbf{c})$  between fake samples and the associated codes reveals mode collapse, and consequently proposed regularizing  $\mathcal{L}_G$  with the negative of the foregoing mutual information term:

$$\tilde{\mathcal{L}}_G = \mathcal{L}_G - \beta I(G(\mathbf{z}, \mathbf{c}); \mathbf{c}), \beta > 0.$$

We propose an alternative that does not rely on meta-data.

It follows from Sklar’s theorem (Theorem 2.1) that to prevent the generator from collapsing to a mode, it suffices to prevent the copula of  $G(\mathbf{z})$  from collapsing to a mode, which can be done by regularizing the generator’s loss function with the entropy of the copula of  $G(\mathbf{z})$ :

$$\bar{\mathcal{L}}_G = \mathcal{L}_G - \beta h(\mathbf{u}_{G(\mathbf{z})}), \beta > 0.$$

Using the problem (CVX-MIND), we may write the regularized batch training step of the generator as

$$\min_G \frac{1}{b} \sum_{i=1}^b \log [1 - D(G(\mathbf{z}_i))] + \beta \left[ \boldsymbol{\theta}^T \phi(\hat{\mathbf{u}}_i) - e^{\boldsymbol{\theta}^T \phi(\mathbf{u}_i)} \right]$$

where  $\beta > 0$ ,  $\phi$  is the constraint function,  $\mathbf{u}_i \sim \mathcal{U}([0, 1]^d)$  are standard uniforms, and  $\hat{\mathbf{u}}_i = \frac{\text{rg}(G(\mathbf{z}_i))}{b+1}$ , the rank being understood as within-batch. The generator step is now followed by the batch copula step

$$\min_{\boldsymbol{\theta}} \frac{1}{b} \sum_{i=1}^b -\boldsymbol{\theta}^T \phi(\hat{\mathbf{u}}_i) + e^{\boldsymbol{\theta}^T \phi(\mathbf{u}_i)}.$$

We call this model Copula Entropy Regularized GANs (CER-GANs). Figure 3 illustrates the efficacy of CER-GANs on the 25-Gaussians dataset (Belghazi et al. (2018)).

## 6 Conclusion

We propose a novel approach for solving maximum-entropy copula problems under flexible linear constraints as a convex optimization problem, and we apply our finding to estimating the mutual information between two random vectors. Our approach is inductive in that it relies on the data generating distribution solely through some association-revealing nonparametric properties; it does not assume we have enough data to fully characterize the underlying true joint distribution. This allows the resulting estimator, which we denote MIND, to be considerably more data efficient than all competing models. We show that MIND can accurately estimate the mutual information even when the sample size  $n$  is not large enough to fully characterize the true data generating distribution. For large  $n$  settings, we show that MIND is a consistent estimator of the true mutual information and has MSE rate  $O(1/n)$ . Beyond mutual information estimation, we show that mode collapse in GANs can be mitigated by adding a regularizing term that maximizes the copula entropy of the generator using MIND.

An implementation of MIND can be accessed from the Python package ‘kxy’ available from Pypi (by running ‘pip install kxy’) or GitHub (<https://github.com/kxytechnologies/kxy-python>).

## References

- Belghazi, M. I., Baratin, A., Rajeshwar, S., Ozair, S., Bengio, Y., Courville, A., and Hjelm, D. (2018). Mutual information neural estimation. In *International Conference on Machine Learning*, pages 531–540.
- Bell, A. J. and Sejnowski, T. J. (1995). An information-maximization approach to blind separation and blind deconvolution. *Neural computation*, 7(6):1129–1159.
- Bottou, L. (2010). Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010*, pages 177–186. Springer.
- Brown, P. F., Della Pietra, V. J., Desouza, P. V., Lai, J. C., and Mercer, R. L. (1992). Class-based n-gram models of natural language. *Computational linguistics*, 18(4):467–480.
- Che, T., Li, Y., Jacob, A. P., Bengio, Y., and Li, W. (2016). Mode regularized generative adversarial networks. In *5th International Conference on Learning Representations, ICLR 2017*.
- Chen, X., Duan, Y., Houthoofd, R., Schulman, J., Sutskever, I., and Abbeel, P. (2016). Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in neural information processing systems*, pages 2172–2180.
- Cover, T. M. (1999). *Elements of information theory*. John Wiley & Sons.
- Csiszár, I. (1975). I-divergence geometry of probability distributions and minimization problems. *The annals of probability*, pages 146–158.
- Dumoulin, V., Belghazi, I., Poole, B., Mastropietro, O., Lamb, A., Arjovsky, M., and Courville, A. (2016). Adversarially learned inference. *arXiv preprint arXiv:1606.00704*.
- Gao, S., Ver Steeg, G., and Galstyan, A. (2015). Efficient estimation of mutual information for strongly dependent variables. In *Artificial intelligence and statistics*, pages 277–286.
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative adversarial networks. *arXiv preprint arXiv:1406.2661*.
- Higgins, I., Amos, D., Pfau, D., Racaniere, S., Matthey, L., Rezende, D., and Lerchner, A. (2018). Towards a definition of disentangled representations. *arXiv preprint arXiv:1812.02230*.
- Hulle, M. M. V. (2005). Edgeworth approximation of multivariate differential entropy. *Neural computation*, 17(9):1903–1910.
- Jaynes, E. T. (1957a). Information theory and statistical mechanics. *Physical review*, 106(4):620.
- Jaynes, E. T. (1957b). Information theory and statistical mechanics. ii. *Physical review*, 108(2):171.
- Kraskov, A., Stögbauer, H., and Grassberger, P. (2004). Estimating mutual information. *Physical review E*, 69(6):066138.
- Kruskal, W. H. (1958). Ordinal measures of association. *Journal of the American Statistical Association*, 53(284):814–861.
- Kullback, S., Khairat, M., et al. (1966). A note on minimum discrimination information. *The Annals of Mathematical Statistics*, 37(1):279–280.
- Kwak, N. and Choi, C.-H. (2002). Input feature selection by mutual information based on parzen window. *IEEE transactions on pattern analysis and machine intelligence*, 24(12):1667–1671.
- Lin, Z., Khetan, A., Fanti, G., and Oh, S. (2018). Pacgan: The power of two samples in generative adversarial networks. *Advances in neural information processing systems*.
- McAllester, D. and Stratos, K. (2020). Formal limitations on the measurement of mutual information. In *International Conference on Artificial Intelligence and Statistics*, pages 875–884.
- Metz, L., Poole, B., Pfau, D., and Sohl-Dickstein, J. (2016). Unrolled generative adversarial networks. *arXiv preprint arXiv:1611.02163*.
- Moon, Y.-I., Rajagopalan, B., and Lall, U. (1995). Estimation of mutual information using kernel density estimators. *Physical Review E*, 52(3):2318.
- Nelsen, R. B. (2007). *An introduction to copulas*. Springer Science & Business Media.
- Newey, K. and McFadden, D. (1994). Large sample estimation and hypothesis. *Handbook of Econometrics, IV, Edited by RF Engle and DL McFadden*, pages 2112–2245.
- Nguyen, T. D., Le, T., Vu, H., and Phung, D. (2017). Dual discriminator generative adversarial nets. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 2667–2677.

- Nguyen, X., Wainwright, M. J., and Jordan, M. I. (2010). Estimating divergence functionals and the likelihood ratio by convex risk minimization. *IEEE Transactions on Information Theory*, 56(11):5847–5861.
- Noshad, M., Zeng, Y., and Hero, A. O. (2019). Scalable mutual information estimation using dependence graphs. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2962–2966. IEEE.
- Oord, A. v. d., Li, Y., and Vinyals, O. (2018). Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Paninski, L. (2003). Estimation of entropy and mutual information. *Neural computation*, 15(6):1191–1253.
- Pathak, D., Agrawal, P., Efros, A. A., and Darrell, T. (2017). Curiosity-driven exploration by self-supervised prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 16–17.
- Poole, B., Ozair, S., Van Den Oord, A., Alemi, A., and Tucker, G. (2019). On variational bounds of mutual information. In *International Conference on Machine Learning*, pages 5171–5180.
- Rudin, W. (1973). *Functional analysis*, mcgraw-hill series in higher mathematics.
- Ruschendorf, L. (1976). Asymptotic distributions of multivariate rank order statistics. *The Annals of Statistics*, pages 912–923.
- Saatchi, Y. and Wilson, A. G. (2017). Bayesian gan. *arXiv preprint arXiv:1705.09558*.
- Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., and Chen, X. (2016). Improved techniques for training gans. *arXiv preprint arXiv:1606.03498*.
- Shohat, J. A. and Tamarkin, J. D. (1943). *The problem of moments*. Number 1. American Mathematical Soc.
- Song, J. and Ermon, S. (2020). Understanding the limitations of variational mutual information estimators. In *International Conference on Learning Representations*.
- Tishby, N., Pereira, F. C., and Bialek, W. (2000). The information bottleneck method. *arXiv preprint physics/0004057*.
- Tishby, N. and Zaslavsky, N. (2015). Deep learning and the information bottleneck principle. In *2015 IEEE Information Theory Workshop (ITW)*, pages 1–5. IEEE.